

Construction and analysis of non-Poisson stimulus-response models of neural spiking activity

Riccardo Barbieri ^{a,*}, Michael C. Quirk ^b, Loren M. Frank ^b, Matthew A. Wilson ^b,
Emery N. Brown ^a

^a Department of Anesthesia and Critical Care, Neuroscience Statistics Research Laboratory, Massachusetts General Hospital,
Division of Health Sciences and Technology, Harvard Medical School/MIT, 55 Fruit Street, Clinics 3, Boston, MA 02114-2698, USA

^b Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Received 28 July 2000; received in revised form 18 September 2000; accepted 13 October 2000

Abstract

A paradigm for constructing and analyzing non-Poisson stimulus-response models of neural spike train activity is presented. Inhomogeneous gamma (IG) and inverse Gaussian (IIG) probability models are constructed by generalizing the derivation of the inhomogeneous Poisson (IP) model from the exponential probability density. The resultant spike train models have Markov dependence. Quantile–quantile (Q–Q) plots and Kolmogorov–Smirnov (K–S) plots are developed based on the rate-rescaling theorem to assess model goodness-of-fit. The analysis also expresses the spike rate function of the neuron directly in terms of its interspike interval (ISI) distribution. The methods are illustrated with an analysis of 34 spike trains from rat CA1 hippocampal pyramidal neurons recorded while the animal executed a behavioral task. The stimulus in these experiments is the animal's position in its environment and the response is the neural spiking activity. For all 34 pyramidal cells, the IG and IIG models gave better fits to the spike trains than the IP. The IG model more accurately described the frequency of longer ISIs, whereas the IIG model gave the best description of the burst frequency, i.e. ISIs ≤ 20 ms. The findings suggest that bursts are a significant component of place cell spiking activity even when position and the background variable, theta phase, are taken into account. Unlike the Poisson model, the spatial and temporal rate maps of the IG and IIG models depend directly on the spiking history of the neurons. These rate maps are more physiologically plausible since the interaction between space and time determines local spiking propensity. While this statistical paradigm is being developed to study information encoding by rat hippocampal neurons, the framework should be applicable to stimulus-response experiments performed in other neural systems. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Hippocampal place cells; Interspike interval distributions; Markov models; Quantile–quantile plots; Kolmogorov–Smirnov plots; Inhomogeneous gamma process; Inhomogeneous inverse Gaussian process; Inhomogeneous Poisson process; Rate-rescaling theorem

1. Introduction

Accurate statistical descriptions of neural spiking activity are crucial for understanding neural information encoding. Simple and inhomogeneous Poisson (IP) processes are the statistical models most frequently used in simulation studies and experimental analyses of neural spike trains (Tuckwell, 1988; Rieke et al., 1997; Gabbiani and Koch, 1998; Reich et al., 1998). Al-

though the Poisson models have provided theoretical predictions and statistical methods for experimental data analysis, spiking activity in many neural systems cannot be completely described by these processes (Gabbiani and Koch, 1998; Shadlen and Newsome, 1998). Statistical analyses of several different neural systems under approximately stationary experimental conditions have described interspike interval (ISI) distributions as unimodal, right skewed probability densities such as the gamma, inverse Gaussian and lognormal (Pfeiffer and Kiang, 1965; Tuckwell, 1988; Iyengar and Liao, 1997; Gabbiani and Koch, 1998). Analytic and simulation studies of elementary stochastic biophysical models of single neurons are consistent

* Corresponding author. Tel.: +1-617-7241061; fax: +1-617-7268410.

E-mail address: barbieri@srlb.mgh.harvard.edu (R. Barbieri).

with these empirical findings (Tuckwell, 1988; Iyengar and Liao, 1997; Gabbiani and Koch, 1998). Because the exponential probability density is the probability model associated with a simple Poisson counting process, the counting processes associated with these elementary models cannot be Poisson. These elementary ISI models based on unimodal right-skewed probability densities are renewal models, and as such, do not allow for dependence in the spiking pattern of the neurons.

Stimulus-response experiments are widely used neuroscience protocols that require statistical analyses for correct interpretation of experimental data. In these experiments a natural or artificial stimulus is given and the response, the spiking activity of one or a set of neurons, is measured. Examples include motion stimulation of the fly H1 neuron (Bialek et al., 1991), natural sound stimulation of the bullfrog eighth nerve (Rieke et al., 1995), position stimulation of pyramidal (place) cells in the rat hippocampus (O'Keefe and Dostrovsky, 1971), and wind stimulation of the cricket (Miller et al., 1991) and cockroach (Rinberg and Davidowitz, 2000) cercal systems. In addition to the applied stimulus, neural spiking activity can also be modified by background variables not associated with the stimulus, such as the theta rhythm in the case of the hippocampal place cells (O'Keefe and Recce, 1993; Skaggs et al., 1996). A strong appeal of using the IP model to analyze stimulus-response data is that the effects of stimuli and background variables on spiking activity can be simply modeled by making the Poisson rate an explicit function of these variables (Brown et al., 1998b; Zhang et al., 1998). Despite this analytic convenience, goodness-of-fit assessments have shown that the IP model does not completely describe the stochastic structure of spike trains from stimulus-response experiments (Van Steveninck et al., 1997; Brown et al., 1998b; Fenton and Muller, 1998; Reich et al., 1998).

More accurate statistical description of neural spike train activity should result if models that allow dependence in the spike train and include the effects of stimuli and background variables on the spiking activity are considered. Development of these new models also requires appropriate goodness-of-fit procedures to assess model agreement with experimental data.

Here, inhomogeneous neural spike train models are derived by combining a renewal model with a one-to-one transformation function to relate the variable of the renewal probability density to the spike times and the stimulus. The transformation used will generate inhomogeneous models in which the spike trains have Markov dependence. This construction is used to build inhomogeneous gamma (IG) and inverse Gaussian (IIG) probability models and present goodness-of-fit methods using Akaike's Information Criterion (AIC), Bayesian Information Criterion (BIC), quantile–quantile (Q–Q) plots, and Kolmogorov–Smirnov (K–S)

plots, to evaluate model agreement with experimental data. The models and the diagnostic methods are applied to the analysis of spike trains from place cells in the CA1 region of the rat hippocampus collected simultaneously with path data from rats foraging in an open circular environment. In these experiments the stimulus is the animal's position in the environment and the response is the neural spiking activity.

2. Model construction

2.1. Inhomogeneous interspike interval probability models

The objective is to model the ISI distribution of a neuron as a function of spiking history, stimulus inputs and background variables. This is done by combining a probability model that defines the stochastic structure of the neural spike train with a one-to-one transformation that relates the spike times, the stimulus, and covariates to the random variable of the probability density. Because the construction uses the intensity-rescaling transformation, the history dependence in the spike train will be Markov.

We begin by defining notation. Let $(0, T]$ denote the observation interval of the experiment and $0 \leq t_0 < t_1 < t_2, \dots, t_k < t_{k+1}, \dots, t_K \leq T$ be the spikes recorded from a given neuron in that interval. Let $f_z(z)$ be a renewal process probability density defined for z on $(0, \infty)$, let t be the time variable defined on (t_a, ∞) , and $z = g(t)$ be a one-to-one differentiable transformation from (t_a, ∞) to $(0, \infty)$, for $t_a \geq 0$. By the change-of-variables formula (Port, 1994), the probability density of t can be computed from $f_z(z)$ and $g(t)$ and is defined as

$$f_t(t) = \left| \frac{dg}{dt} \right| f_z(g(t)). \quad (1)$$

To illustrate the construction, the standard gamma probability density is considered as the first renewal model defined as

$$f_z(z) = \frac{1}{\Gamma(\gamma)} z^{\gamma-1} e^{-z}, \quad (2)$$

where $z > 0$, $\gamma > 0$ is the shape parameter and $\Gamma(\gamma)$ is the gamma function (Johnston and Kotz, 1971). For the standard gamma probability density one has $E(z) = \gamma$ and $\text{Var}(z) = \gamma$. The gamma probability density in Eq. (2) has the exponential probability $f_z(z) = e^{-z}$ as the special case $\gamma = 1$. As γ increases, the shape of the probability density changes from that of an exponential density to one that is more peaked and symmetric (Gabbiani and Koch, 1998). If $\gamma < 1$ ($\gamma > 1$) then the spike train data are more (less) variable than the simple exponential or Poisson counting models. The parameter γ is also the signal-to-noise ratio, and $\gamma^{-1/2}$

is the coefficient-of-variation. The gamma probability model has been suggested as an alternative to the exponential probability model for describing ISI probability densities because the additional parameter γ gives a more flexible range of shapes (Bishop et al., 1964; Nakahama et al., 1968; Correia and Landolt, 1977; Gabbiani and Koch, 1998). The Erlang probability density is a special case of the gamma probability density in which γ is a positive integer greater than 1. The Erlang probability density is the ISI model obtained when a non-leaky stochastic integrate-and-fire model with a fixed threshold is driven by Poisson excitatory inputs with a constant rate parameter (Tuckwell, 1988).

The intensity-rescaling transformation is defined as

$$z = g(t) = \gamma \int_{t_a}^t \lambda(u) du, \quad (3)$$

where $\lambda(t)$ is a strictly positive intensity function. Applying Eq. (1) with $f_z(z)$ given in Eq. (2), $g(t)$ given in Eq. (3), and setting $t = t_k$, and $t_a = t_{k-1}$, yields the inhomogeneous gamma (IG) probability density

$$f_t(t_k|t_{k-1}) = \frac{\gamma \lambda(t_k)}{\Gamma(\gamma)} \left[\gamma \int_{t_{k-1}}^{t_k} \lambda(u) du \right]^{\gamma-1} \exp \left\{ -\gamma \int_{t_{k-1}}^{t_k} \lambda(u) du \right\}, \quad (4)$$

since

$$\left. \frac{dz}{dt} \right|_{t_k} = \gamma \lambda(t_k),$$

(Berman, 1981; Brown et al., 1998a). The probability density in Eq. (4) has the IP probability density as the special case of $\gamma = 1$, defined as

$$f_t(t_k|t_{k-1}) = \lambda(t_k) \exp \left\{ -\int_{t_{k-1}}^{t_k} \lambda(u) du \right\}. \quad (5)$$

As a second example the inverse Gaussian probability density is taken as the renewal probability model

$$f_z(z) = \left(\frac{1}{2\pi z^3} \right)^{1/2} \exp \left\{ -\frac{1}{2} \frac{(z - \psi)^2}{\psi^2 z} \right\}, \quad (6)$$

where $z > 0$ and $\psi > 0$ is the location parameter and $E(z) = \psi$ (Chhikara and Folks, 1989). The inverse Gaussian probability density has been suggested as an empirical model for ISI probability densities (Tuckwell, 1988). The inverse Gaussian has also been suggested because it is the ISI probability density that is derived when a non-leaky integrate-and-fire model with a fixed threshold is driven by Wiener process inputs (Tuckwell, 1988; Iyengar and Liao, 1997).

Applying Eq. (1) with $f_z(z)$ given in Eq. (6), $z = g(t) = \int_{t_a}^t \lambda(u) du$, $dz/dt|_{t_k} = \lambda(t_k)$, $t = t_k$ and $t_a = t_{k-1}$, yields the IIG probability density

$$f_t(t_k|t_{k-1}) = \frac{\lambda(t_k)}{\left[2\pi \left[\int_{t_{k-1}}^{t_k} \lambda(u) du \right]^3 \right]^{1/2}} \times \exp \left\{ -\frac{1}{2} \frac{\left(\int_{t_{k-1}}^{t_k} \lambda(u) du - \psi \right)^2}{\psi^2 \int_{t_{k-1}}^{t_k} \lambda(u) du} \right\}. \quad (7)$$

The inhomogeneous models in Eqs. (4) and (7) should provide better descriptions of stimulus-response ISI data than the IP model. Unlike in the case of the IP model, the counting process models associated with Eqs. (4) and (7) cannot be expressed in closed form. The specification of $\lambda(t)$ used in the analyses of hippocampal place cells is described in Section 4.2.

2.2. Relation between ISI distribution and the spike rate function

For the IP model the intensity function $\lambda(t)$ is also the rate function of the spiking process. This is not true in general for point process models. The rate function, also termed the hazard function in survival analysis, can be defined for any ISI probability density as

$$r(t) = \frac{f_t(t|t_{k-1})}{1 - \int_{t_{k-1}}^t f_t(u|t_{k-1}) du} \quad (8)$$

for $t > t_{k-1}$ (Daley and Vere-Jones, 1988). The rate function defines the probability of a spike in $[t, t + \Delta t)$, given that the last spike occurred at t_{k-1} , as $r(t)\Delta t$. In survival analysis, the hazard $r(t)\Delta t$ may be interpreted as the risk of a failure in $[t, t + \Delta t)$ given that the system has survived up to time t . In Appendix A (Eq. (A.4)) it is shown that $r(t) = \lambda(t)$ for the IP model. It is also shown in Appendix A that any ISI probability density may be written explicitly in terms of its rate or hazard function (Kalbfleisch and Prentice, 1980; Brillinger, 1988; Daley and Vere-Jones, 1988). Eq. (8) is the basis for the goodness-of-fit techniques developed in Section 3 using the rate-rescaling theorem.

3. Goodness-of-fit

An essential component of the statistical analysis paradigm is the assessment of goodness-of-fit, i.e. the evaluation of how well the spike train data are described by a given inhomogeneous probability model. Four goodness-of-fit procedures are considered. Two of these are the well-known model comparison statistics, AIC (Box et al., 1994) and BIC (Shwartz, 1978). The other two are the Q–Q plots and K–S plots, derived specifically for point process data analysis.

AIC and BIC give single number summaries of the efficiency with which a given model describes a set of data by measuring the trade-off between the fit of the model to the data and the number of parameters required to achieve that fit. These indices are helpful for identifying a good model or set of models for a data series. However, they do not describe what characteristics in the data the model does or does not describe. For this purpose Q–Q (Wilks et al., 1962) and K–S (Ogata, 1988) plots are developed. Both of these diagnostic techniques are based on the rate-rescaling theorem (Meyer, 1969) described below.

3.1. Rate-rescaling theorem

Let $R(t) = \int_0^t r(u) du$ define the rate-rescaling transformation function. Because $r(t)$ is nonnegative, $R(t)$ is a monotonically increasing function, and the rate-rescaling transformation is a one-to-one mapping of the t_k into $R(t_k)$. We let

$$\tau_k = R(t_k) - R(t_{k-1}). \quad (9)$$

If the inhomogeneous model from which $R(t)$ is constructed is correct, then by the rate-rescaling theorem (Meyer, 1969; Daley and Vere-Jones, 1988; Ogata, 1988) the τ_k s are independent, identically distributed exponential random variables with unit rate. Otherwise stated, the τ_k s are distributed as a stationary Poisson process with unit rate. The rate-rescaling theorem states that any regular point process can be transformed into a Poisson process with unit rate.

3.2. Q–Q plots

For any postulated ISI model the rate-rescaling theorem can be used to measure agreement between the model and the spike train by constructing a Q–Q plot. Assume that there are K spikes in the recorded spike train. Let $F_\tau(\tau)$ be the cumulative distribution function (CDF) corresponding to the unit exponential probability density in Eq. (2). If the corresponding inhomogeneous model is correct, then the transformed spike times τ_k s are independent and identically distributed with exponential probability density $e^{-\tau}$. Let $\tau_{(k)}$ be the set of τ_k s ordered from smallest to largest and let $b_k = (k - 1/2)/K$, for $k = 1, \dots, K$. The $\tau_{(k)}$ s are the K empirical quantiles of the transformed spike times for a given model and set of parameter estimates. One can also estimate the quantiles of the transformed data from the model by computing

$$\tilde{\tau}_k = F_\tau^{-1}(b_k) = -\log(1 - b_k), \quad (10)$$

where F_τ^{-1} is the inverse of the CDF of $e^{-\tau}$. A plot of $\tau_{(k)}$ against $\tilde{\tau}_k$, i.e. the empirical quantiles versus the model derived quantiles, is termed a Q–Q plot. If the model is correct, then the Q–Q plot should be a

straight line through the origin with slope 1 (Wilks et al., 1962). In most applications the Q–Q plots assess agreement between a data series and a given probability density model without the need to transform the data. Because in this analyses the transformation between the τ_k s and $R(t_k)$ is one-to-one, there is close agreement between the inhomogeneous model and the place cell spike train series if and only if there is close agreement between the τ_k s and the unit exponential probability model. Hence, Q–Q plots on the transformed scale may be used to assess goodness-of-fit of the inhomogeneous model. The deviations from the 45° line show where the model fails to agree with the data. An important advantage of the Q–Q plots is that they use all the data points and require no binning.

3.3. K–S plots

The K–S plot is another graphical measure of goodness-of-fit constructed based on the rate-rescaling theorem. If the model generating the τ_k s is correct, then the $u_k = 1 - e^{-\tau_k}$ are independent, identically distributed uniform random variables on $[0, 1)$. Therefore, the ordered quantiles $u_{(k)}$, plotted against $b_k = (k - 1/2)/K$ should be a 45° line (Ogata, 1988). Confidence bounds based on the distribution of the K–S statistic can be constructed for the plot to test formally the agreement between a proposed model and the data (Ogata, 1988). Like the Q–Q plots, the K–S plots use all the data points and do not require binning.

4. Data analysis: hippocampal place cells

To illustrate the model construction and goodness-of-fit methods developed in Section 3, in this section the IP, IG and IIG models using a Gaussian spatio-temporal intensity function are presented. The model fits are used in two applications: an analysis of burst activity and a study of temporal and spatial rate maps.

4.1. Experimental protocol

The experimental methods have been previously reported in detail (Brown et al., 1998b). Briefly, a Long Evans rat was implanted with microdrive arrays housing 12 tetrodes (four wire electrodes) (Wilson and McNaughton, 1993). Recordings of the animal's position (stimulus), spiking activity (response) and EEG theta rhythm were made while the animal foraged for chocolate pellets scattered approximately uniformly in a black cylindrical environment 70 cm in diameter with 30 cm high walls and a fixed visual cue (Muller et al., 1987).

The spiking activity of 34 place cells from the animal recorded during the first 14 min of foraging was ana-

lyzed. The number of spikes per cell ranged from 379 to 2254.

4.2. A spatio-temporal model of place cell spiking activity

The form of the intensity function $\lambda(t)$ specified for the place cell data is the one proposed by Brown et al. (1998b). The essential features of the place field model are summarized here. To define the model let $x(t) = [x_1(t), x_2(t)]'$ be the 2×1 vector denoting the animal's position at time t , and let $\theta(t)$ be the phase of the theta rhythm at time t . Assume that the place field position component for a cell is modeled as a Gaussian function defined as

$$\lambda(t|x(t), \xi_x) = \exp\left\{\alpha - \frac{1}{2}(x(t) - \mu)'W^{-1}(x(t) - \mu)\right\}, \quad (11)$$

where $\mu = [\mu_1, \mu_2]'$ is the 2×1 vector whose components are the x_1 and x_2 coordinates of the place field center, α is the location intensity parameter,

$$W = \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix} \quad (12)$$

is a scale matrix whose scale parameters in the x_1 and x_2 directions are σ_1^2 and σ_2^2 , respectively, and $\xi_x = [\alpha, \mu, W]$.

Eq. (11) defines the intensity function of the stimulus-response model and represents the function's dependence on the animal's position in the environment as a two-dimensional Gaussian surface. The stimulus (position) is related to the response (neural spiking activity) by substituting Eq. (11) into Eqs. (4), (5) and (7). For the IP, IG and IIG models the parameters estimated by maximum likelihood (Casella and Berger, 1990) for each of the 34 place cells are respectively $\xi_{IP} = \xi_x$, $\xi_{IG} = (\xi_x, \gamma)$ and $\xi_{IIG} = (\xi_x, \psi)$.

For the IP analyses the effect of theta phase modulation on the spike train activity is also included. The theta phase component of the cell is modeled as a cosine function defined as

$$\lambda_\theta(t|\theta(t), \xi_\theta) = \exp\{\beta \cos(\theta(t) - \theta_0)\}, \quad (13)$$

where β is a modulation factor, θ_0 is the theta phase of the maximum of $\lambda_\theta(t|\theta(t), \xi_\theta)$ and $\xi_\theta = [\beta, \theta_0]$. Therefore, the intensity function for the IP model with theta modulation is the product of the intensity functions in Eqs. (11) and (13). This model was also fit by maximum likelihood to the spike train data of each of the 34 place cells.

4.3. Goodness-of-fit

4.3.1. AIC and BIC analysis

AIC and BIC give identical results when comparing

goodness-of-fit by cell of the IP, IP with theta, IG, and IIG models. According to both indices, the IG model fits better than both the IP model and the IP model with theta for all the cells. By AIC and BIC, the IIG model gives a better fit than the IP model with and without theta for all the cells with the exception of one (Cell 8). The differences between the respective AIC and BIC values for Cell 8 for the IIG, IP and IP model with theta were small, but nevertheless, favored slightly the two Poisson models. Overall, the AIC/BIC criteria suggested that 11 cells (including Cell 8) were best fit by the IG model, whereas the remaining 23 were best fit by the IIG model.

4.3.2. Q–Q plot analysis

The quantiles of the rate-rescaled probability densities vary across two or more orders of magnitude. Therefore, to evaluate accurately model agreement and disagreement with the experimental data the Q–Q plots of the place cells are displayed in three different ranges: 0–30th percentile (Fig. 1A); 0–95th percentile (95% range, Fig. 1B); and 0–100th percentile (Fig. 1C). Ten representative cells are shown to summarize the findings. It is shown below, in Section 4.4.1, that the τ_k 's up to the 30th percentile correspond approximately to ISI's ≤ 20 ms. The study of model goodness-of-fit in this range will be important for interpreting the bursts analysis in that section.

In the 30% range (Fig. 1A), the IIG model fits show the best agreement with the data as their Q–Q plots in this range are all close to the 45° line. The IP model fits appreciably underestimate the probability densities of the data, whereas the IG model fits initially overestimate the probability densities of the data in the lower quantiles and then underestimate the larger quantiles in this range. In the 95% range (Fig. 1B), the IG model fits give the strongest agreement with the probability densities of the data. Beyond approximately the 30th percentile (small black square in Fig. 1B), the IIG model fits overestimate the probability densities of the data. In contrast, all the IP model fits underestimate the probability densities of the data up to approximately the 50th percentile, and then overestimate them from the 50th to the 95th percentile. Both the IP and IG models overestimate the data probability densities in the 95th to 100th percentile (Fig. 1C). Although all the IIG model fits in this range underestimate the probability densities of the data, they are closer to the 45° line than those of the other two models.

To present the Q–Q plot goodness-of-fit analysis for all the cells, the slopes of the Q–Q plots were computed for each cell for the three models. These findings were summarized as boxplots in the 30% (Fig. 1D), the 95% (Fig. 1E), and the 100% (Fig. 1F) ranges. The slope of the Q–Q plot provides a quantitative measure of how

close a given model fit is to the data. A slope of 1 suggests close agreement between the model and the data.

The boxplot summaries across all the cells agree completely with the findings from the graphical analyses in Fig. 1A–C. The median slope of the IP model in the 30% range was 0.14 (Fig. 1D), consistent with this model underestimating the probability density of the data in this range (Fig. 1A). The median slope of the IG model in the 30% range was 0.54 (Fig. 1D), because the IG model fits initially overestimate, then underestimate the probability densities of the data in this range (Fig. 1A). In contrast, the median slope of the IIG model fits in the 30% range was 0.79 (25th–75th boxplot percentile ranging from 0.67 to 0.98) suggesting as in Fig. 1A, that the IIG model has the best agreement in this range.

For the 95% range (Fig. 1E) the median slope of the IP model Q–Q plots is 1.19 (0.96–1.23), consistent with the IP model fits underestimating the data probability

densities up to approximately the 50th percentile, and overestimating them from this percentile upward (Fig. 1B). Because the degree of overestimation exceeds the degree of underestimation, the median slope exceeds 1. Nearly all the IIG Q–Q plot slopes in the 95% range are greater than 1, with a median of 1.17. The median slope of the IG model fits in the 95% range was 0.95 (0.89–0.99) in agreement with Fig. 1B. For all the cells, the IG fits give the best agreement with the probability densities of the data in the 95% range.

In the 100% range, the slopes of the IIG model are closer to 1 than those of either the IP or IG (Fig. 1F). The median Q–Q plot slopes were 1.03 for the IIG model, 1.19 for the IG model, and 1.88 for the IP model. Overall, the IIG model fits are closer to the probability densities of the data because they agree well with the data in the 30% range (Fig. 1A), are close in the 50–95% range (Fig. 1D), and deviate least in the 95th to 100th percentile (Fig. 1C).

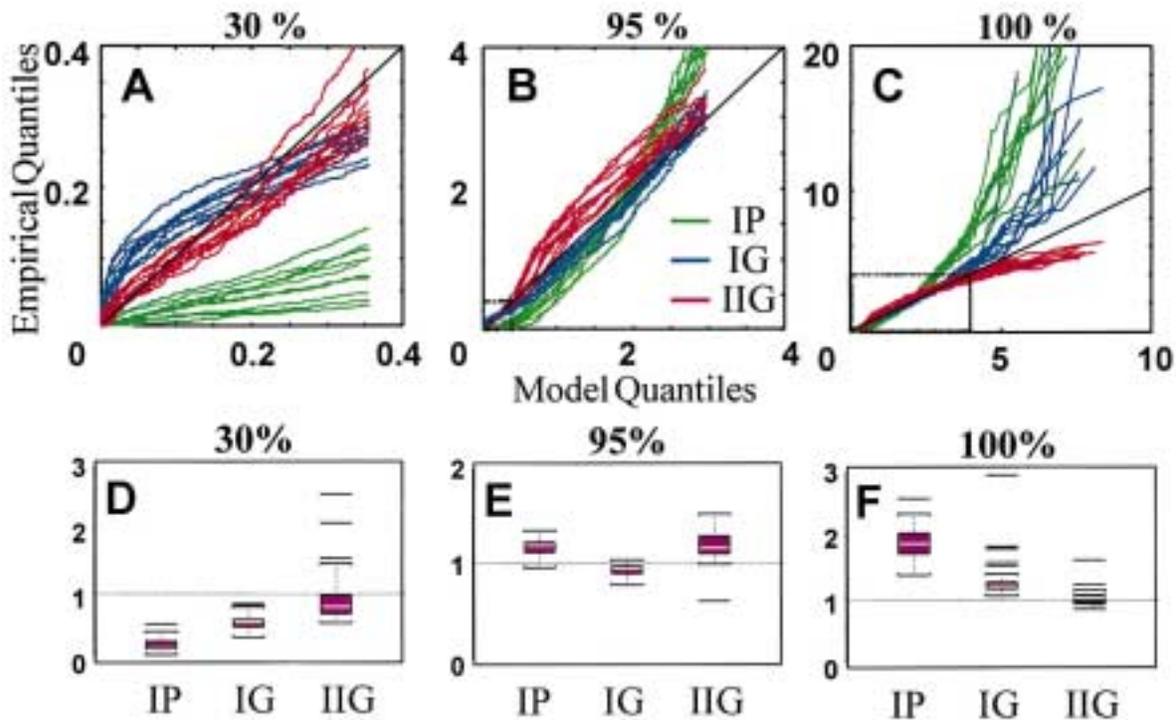


Fig. 1. (A–C) Quantile–quantile (Q–Q) plots of the inhomogeneous Poisson (IP) (green lines), inhomogeneous gamma (IG) (blue lines), and inverse Gaussian (IIG) (red lines) model fits to the spike trains of ten of the 34 place cells analyzed. The 45° line represents exact agreement between the model and the spike train data. The Q–Q plots are displayed for 0–30th percentile (A); 0–95th percentile (B); and 0–100th percentile (C). The dotted squares in (B) and (C) indicate respectively the 30th and the 95th percentiles. The IIG model fits the data best in both the 30% range and the 95–100% percentile range, whereas the IG model fits best in the 30–95% range. (D–F) Box and whiskers plot summaries of the Q–Q plot slopes for the IP, IG and IIG model fits of all 34 cells for the 30% (D), the 95% (E), and the 100% (F) ranges. The lower border of the box is the 25th percentile of the distribution and the upper border is the 75th percentile. The white bar within the box is the median of distribution. The distance between the 25th and 75th percentiles is the interquartile range (IQR). The lower (upper) whisker is at $1.5 \times$ the IQR below (above) the 25th (75th) percentile. All the black bars below (above) the lower (upper) whiskers are far outliers. A slope of 1 for an approximately linear relation between the empirical and model quantiles suggests close agreement between the model and the data. The boxplot summaries are in strong agreement with the Q–Q plots in A, B, and C. D and F, show that the IIG model fits describe best the data in the 30 and the 95–100% ranges, respectively. E shows that the IG model gives the best fits in the 95% range.

4.3.3. K–S plot analysis

The K–S plots provide an alternative graphical summary for comparing the model fits with the data. A rigorous application of the K–S tests to the fits of the IP, IG and IIG models rejects the hypothesis that any of these three models agrees with the spike data of any of the 34 cells. However, because the K–S plot provides an explicit graphical summary of where the model fails to agree with the data, we measured the lack-of-fit between each model and the data taking as the K–S distance the average distance between the K–S plot and the 45° line.

The K–S distances agreed closely with the AIC/BIC findings regarding which models best described the data. The AIC/BIC criteria identified 11 cells as best fit by the IG model, 23 cells as best fit by the IIG model, and no cell as best fit by the IP model. The K–S statistics identified 21 of the same 23 cells as best fit by the IIG model, the remaining 13 as best fit by the IG model and none as best fit by the IP model. The two cells for which the AIC/BIC and the K–S distance classifications did not agree had the closest AIC/BIC values of all the cells when comparing IG and IIG.

Fig. 2A (B) shows the K–S plot of Cell 1 (Cell 9) which is best fit by the IG (IIG) model according to both the AIC/BIC and K–S distance criteria. The close agreement of the IG model with the spike train of Cell 1 is evidenced by the fact that its K–S plot lies almost entirely within the 95% confidence bounds (Fig. 2A). Although the overall fit of the IG model is best for this cell, the IIG model also shows excellent agreement with these data through the lower quantiles. Similarly Cell 9, which is best fit by the IIG model, has its K–S plot almost entirely within the 95% confidence bound (Fig. 2B).

The K–S distances for each cell for the fits of each of the three models are summarized with boxplots in the 30% (Fig. 2C), the 95% (Fig. 2D), and the 100% (Fig. 2E) ranges. As indicated by the Q–Q plot analysis the IIG models fits agree most closely with the probability densities of the data in the 30% range (Fig. 2C). The median K–S distance in the 30% range was 0.12 for the IP model, 0.038 for IG, and 0.028 for IIG. Twenty-nine of the cells had smaller K–S distances in this range for the IIG model fits. The remaining five cells had smaller K–S distances for the IG model fits.

In the 95% range (Fig. 2D) the median K–S distance was 0.18 for the IP model, 0.072 for the IG, and 0.074 for the IIG. Both the IG and IIG models gave better fits than the IP in this range. However, unlike in the Q–Q plot analyses, there was only a small difference in the median value of the K–S statistic for the IG model compared with the IIG model. The difference between the IG and the IIG models is less pronounced than in the Q–Q plot analysis because the K–S plots are normalized to lie between 0 and 1. In the 100% range

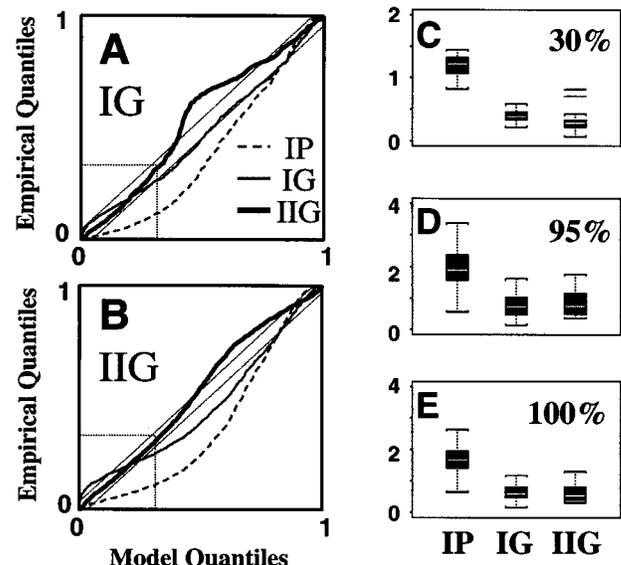


Fig. 2. (A, B) Kolmogorov–Smirnov (K–S) plots for the inhomogeneous Poisson (IP) (dashed lines), inhomogeneous gamma (IG) (thin solid lines), and inverse Gaussian (IIG) (bold solid lines) model fits of two place cells (A, Cell 1; B, Cell 9). The dotted squares in the lower left corners indicate the 30th percentile. Two sided 95% confidence error bounds of the K–S statistics are displayed for each cell (dotted 45° lines). The IG or IIG label in each panel corresponds to the model that gave the best fit in the Akaike’s Information Criterion (AIC)/Bayesian Information Criterion (BIC) and K–S distance analyses. Because the K–S plots are constructed from a uniform distribution on the interval (0, 1), the 30th, 95th, and 100th percentiles correspond respectively to quantiles 0.3 (dotted line), 0.95 and 1.0 on both the x - and y -axes. Cell 1 was best fit by the IG model based on the AIC/BIC criteria, and the IG K–S plots of these cells are closer to the 95% confidence bounds than either those of the IP or IIG models. Cell 9 was best fit by the IIG model based on the AIC/BIC analysis, and the IIG K–S plots for these cells are the closest of the three models to the 95% confidence bounds. The K–S plots, like the Q–Q plots, show that the IIG model fits best in the lower quantiles, the IG fits best in the intermediate quantiles and the IP model fits give the poorest agreement with the data of the three models. (C–E) Boxplot summaries of the K–S distances for the 30% (C), the 95% (D), and the 100% (E) range. These summaries suggest that both the IG and the IIG models give better fits than the IP model. The IIG model fits agree best with the probability densities of the data in the 30% range.

(Fig. 2E) the median K–S distance was 0.16 for IP, 0.059 for IG, and 0.052 for IIG. These findings are consistent with the Q–Q plot analysis findings, which showed that the IIG model fits deviates least from the probability models of the data in the 95th–100th percentile range. Here again, the normalization of the K–S plot makes the difference between the IG and IIG models less pronounced.

4.4. Applications

4.4.1. Burst analysis

By construction, the rate-rescaling transformation maps a certain fraction of short ISIs into a certain

fraction of short τ_k s. The closer agreement between the IIG model and the probability densities of the data found in both the Q–Q and K–S plot analyses for the 30% range may be due to a more accurate description by this model of the frequency of short ISIs or bursts in the spike trains. Therefore, the relation between short ISI and the τ_k s defining a short ISI or burst ($\text{ISI} \leq 20$ ms) was analyzed (McNaughton et al., 1983).

For each place cell the maximum value of τ_k associated with the ISIs ≤ 20 ms was computed, and this value was termed the τ_{burst} quantile. Each τ_{burst} quantile defined a different percentile of the τ_k distribution for each place cell, for each probability model. These percentiles were approximately around the 30th percentile, which is why the 0–30% range was used in the Q–Q and K–S plot analyses. Overall, the median fraction of spike train ISIs that were bursts was 0.39 with a range of 0.15–0.58 (0.38 for Cell 9 in Fig. 3). By construction every burst had $\tau_k < \tau_{\text{burst}}$. However, there were also ISIs longer than 20 ms which had $\tau_k < \tau_{\text{burst}}$ (Fig. 3). Therefore, to measure how strongly bursts were associated with small τ_k s, we computed for each place cell, for each model, the fraction of ISIs that were bursts ($\text{ISIs} \leq 20$ ms) given that the τ_k values were less than that cell's τ_{burst} .

The median fraction of ISIs satisfying this burst condition for the IP and IG model fits was 0.85 with a range of 0.66–0.94. The median fraction of ISIs satisfying this condition for the IIG model fits was 0.97 with a range of 0.78–1. The IIG values were significantly closer to 1 than the ones for the IP or IG ($P < 10^{-8}$, sign test). Because the IIG model describes well the small τ_k s and, the small τ_k s are highly associated with bursts, it can be concluded that a considerable improvement of the IIG over the IP model is due to a more accurate description of the place cell bursting activity.

Similarly, if one defines the non-burst range as the τ_k s $> \tau_{\text{burst}}$ we found that the median fraction of ISIs longer than 20 ms that lie in this range is 0.91 (0.78–0.95). This analysis suggests the improvement of the IG model over the IP is due primarily to a more accurate description of longer τ_k s associated with ISIs longer than 20 ms. This finding is consistent with the Q–Q plot analyses, which showed that the IG models gave the best fits in the 95% range (Fig. 1B).

4.4.2. The rate function

To study further the implications of the model framework the temporal and spatial rate functions derived for each model are compared for Cell 9. Because Eq. (11) is a spatio-temporal model, the rate function in Eq. (8) can be plotted for a given model as either a function of time (Fig. 4A, C, F) or space (Fig. 4B, E, H). For the IP model the rate function and the intensity function are equivalent. As a consequence, the rate does not depend on the history of the spiking activity (Fig. 4A).

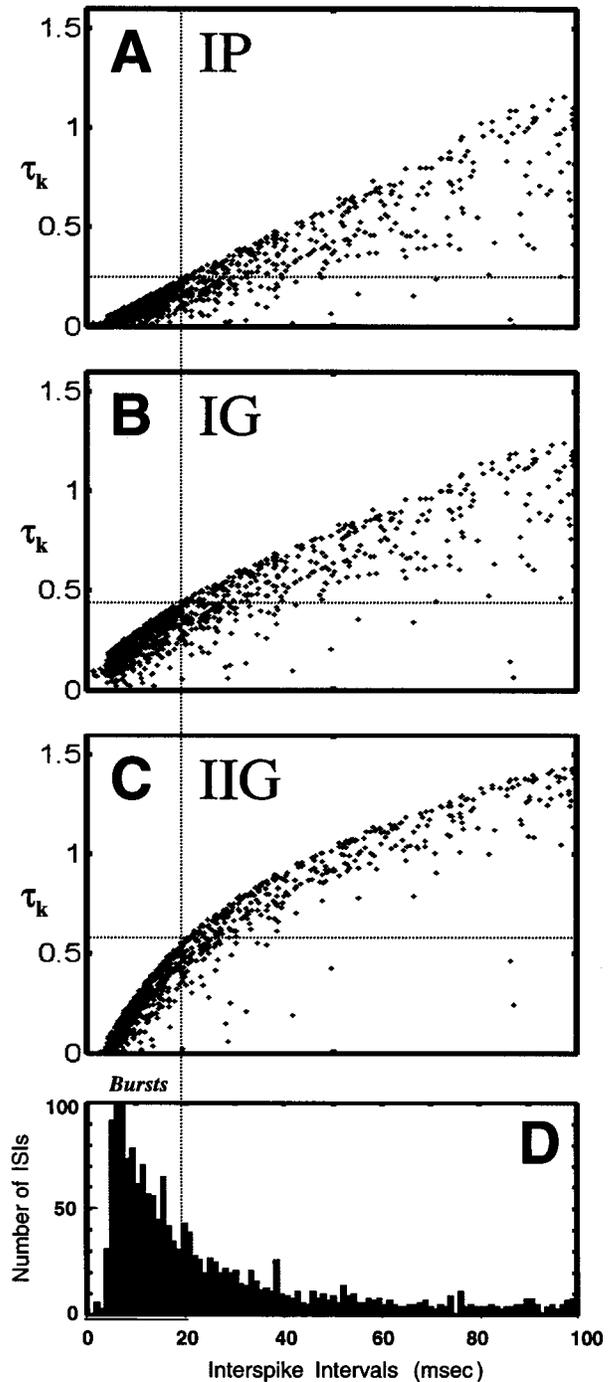


Fig. 3. (A) Representative examples of the relation between the estimated τ_k and the interspike interval (ISIs) from the inhomogeneous Poisson (IP) (A), inhomogeneous gamma (IG) (B), and inverse Gaussian (IIG) (C) from Cell 9. The τ_k s from quantiles 0–75 are plotted along with their corresponding ISIs. The bursts correspond to ISIs ≤ 20 ms (vertical dashed line). The value of τ_{burst} (horizontal dashed line), defined for each cell as the largest value of τ_k for ISIs ≤ 20 ms. The τ_{burst} value for each cell is used to define the range of τ_k corresponding to burst ISIs. The greater the correspondence between bursts and τ_k s for a given model, the greater the number points in the lower left rectangle relative to the number in the lower right rectangle in panels A–C. (D) The ISI histogram of Cell 9. Thirty-eight percent of the ISIs of this cell are bursts.

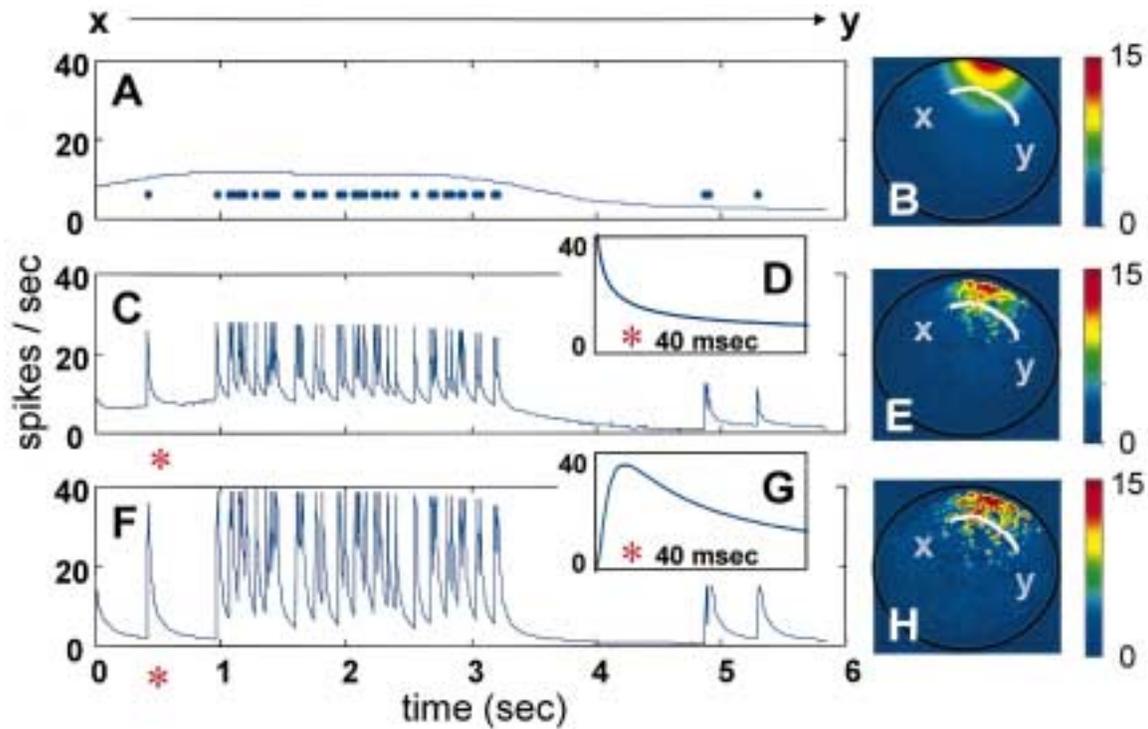


Fig. 4. Spiking activity, model temporal and spatial rate function estimates for Cell 9 during 6 s of the experiment. (A) The spiking activity (dots) and the temporal rate function for the inhomogeneous Poisson (IP) model (solid line). (B) Pseudocolor map of the IP spatial rate function estimate and the trajectory of the animal as it moved from point x to point y (white curve). (C) The temporal rate function for the inhomogeneous gamma (IG) model. (D) Detail (40 ms) of the IG temporal rate function at the time location indicated by the red asterisk in (C). (E) Pseudocolor map of the IG spatial rate function estimate. (F) Temporal rate function of the inverse Gaussian (IIG) model. (G) Detail (40 ms) of the IIG temporal rate function at the time location indicated by the red asterisk in (F). (H) Pseudocolor map of the IIG spatial rate function estimate. Because the IP rate function is the intensity function defined in Section 4.2, it does not depend on spike train history, whereas the rate functions of the IG and IIG models have Markov dependence. In contrast to the IP rate function, if the animal does not visit a particular region of space the IG and IIG rate functions are not defined at that location. The Akaike's Information Criterion (AIC)/Bayesian Information Criterion (BIC) and Kolmogorov–Smirnov (K–S) plot analyses suggest that this cell's spike train data are best described by the IIG model. Hence, the IIG rate maps offer the most accurate description of this cell's local spiking propensity.

The IG (Fig. 4C) and IIG (Fig. 4F) rate functions depend on the history of the spiking activity and are discontinuous at each spike time. The two rate functions have the same general shape, but with important marked local differences in their behavior immediately following a spike. For all the cells the IG model fits had estimated values of γ less than 1, consistent with the spike train having more variation than a Poisson model as discussed in Section 2. Because for $\gamma < 1$ the gamma probability density in Eq. (2) is infinite at the origin, the IG rate function is infinite immediately after the time of the previous spike, then decreases until the next spike arrives (Fig. 4D inset). For the IIG model, the rate function (Fig. 4G inset) is zero at the previous spike, increases rapidly, and then decreases—faster in these analyses than the IG rate—until the next spike. The IIG rate function is the more physiologically plausible, since its values are always near zero immediately following a spike. Despite these differences in local behavior of the IG and IIG rate functions immediately following a spike, both follow the dynamics of the spike train data more closely than the IP rate function. The

improved agreement is especially evident during periods of high spiking activity such as the intervals between 1 and 3 s in Fig. 4A, C, and F. The improvement is also evident between 0.5 and 1 s where spiking activity is absent. In this interval, the IP model predicts a rate increase because the animal's trajectory passes near the maximum of the Poisson rate map, whereas both the IG and IIG models predict declining rates which are more consistent with the observed data.

The IP spatial rate function is again the spatial component of the intensity function described in Eq. (11) and Fig. 4B. It depends only on space and is independent of the spike train history. Because the IG and IIG models depend on the history of spiking activity, trajectories through an identical location do not carry the same rate information. In contrast to the IP rate function, if the animal does not visit a particular region of space, the IG and IIG rate functions are not defined at that location. Although the IG and IIG models use a Gaussian intensity function, the spatial rate maps for these two models are non-Gaussian due to the dependence of spiking activity on history. Be-

cause Cell 9 is best fit by the IIG model, its temporal and spatial rate map estimates give the most accurate characterization of local spiking propensity for this neuron.

5. Discussion

Our paradigm provides a flexible framework for constructing and analyzing non-Poisson statistical models of neural spike train activity collected in stimulus-response experiments. The paradigm consists of maximum likelihood fitting an inhomogeneous model and measuring of its goodness-of-fit using AIC, BIC, Q–Q plots and K–S plots. The framework is illustrated in an analysis of hippocampal place cells.

The inhomogeneous stimulus-response models are constructed by generalizing the way in which the intensity-rescaling transformation is used with the exponential probability density to derive the inhomogeneous Poisson model. The choice of renewal probability density and the one-to-one transformation between the spike train and the random variable of the renewal process is arbitrary. Hence, they may be selected to give the best statistical description of a particular neural system based on either biophysical and/or empirical evidence. Inclusion of covariates, such as the theta rhythm for hippocampal place cells, as a component of the intensity function along with the stimulus is equally as easy for these models as it is for the inhomogeneous Poisson. By construction the current models assume Markov dependence in the spike trains. As suggested in Appendix A, the framework extends easily to model more complex dependence between current spiking activity, and the history of the spike train, the stimulus, and covariates. In addition, the paradigm shows explicitly how to define the rate of a spike train given the specifications of its ISI probability model. Because Eq. (8) relates the ISI distribution of a neural spike train directly to the spike rate function, it has important implications for analyzing the general question of temporal versus rate coding in neural systems.

Most non-Poisson statistical analyses of ISI distributions have used renewal models with no stimulus or covariates inputs. These analyses have included gamma, lognormal and inverse Gaussian (Rodieck et al., 1962; Tuckwell, 1988; Gabbiani and Koch, 1998) and generalized inverse Gaussian models (Iyengar and Liao, 1997). The gamma and inverse Gaussian renewal probability models are members of the generalized inverse Gaussian family studied by Iyengar and Liao. These authors assessed goodness-of-fit using a test statistic similar to AIC. Brillinger (1988) reported a maximum likelihood approach using a generalized linear model to estimate a discretized random threshold model in which the triggering stimuli were spike trains from the same or

different neurons. Kass and Ventura (in preparation) have recently studied Markov dependent ISI models that can be fit by discretizing time and using the generalized linear model. Because our rate functions are not linear functions of the stimulus and the covariates, our models cannot be fit to experimental data using the generalized linear model framework.

The Q–Q and K–S plots provide visual assessments of agreement between a model and a spike train without binning. For the Q–Q and K–S plots the reference probability model is, because of the rate-rescaling theorem, the exponential probability density. As a consequence, these goodness-of-fit techniques are broadly applicable. When using either the Q–Q or K–S plot to compare different models for the same spike train data set, the analysis can be displayed on the same graph. In addition, the K–S plots provide confidence bounds and allow formal testing of agreement between the data and any proposed model. Different percentile ranges of the Q–Q plots were plotted to facilitate graphical analysis of model goodness-of-fit.

Berman (1983) and Ogata (1988) first proposed graphical analyses based on rescaling transformations to assess agreement between inhomogeneous models and point process data arising from the study of seismic data. As suggested by Ogata (1988), multiple goodness-of-fit measures were applied to assess agreement between the models and the spike trains. The current study may represent the first application of point process Q–Q and K–S plots in neural data analysis. Previous methods for assessing goodness-of-fit for inhomogeneous models of neural spike trains have been devoted primarily to Poisson models (Brown et al., 1998b; Fenton and Muller, 1998). None of these approaches extends to goodness-of-fit assessments for inhomogeneous spike train models. Reich et al. (1998) used a rate-rescaling transformation and a test statistic, called the power ratio, to assess agreement between neural spike trains and models based on the peristimulus time histogram. Use of the power ratio entails computing the Fourier transform of the spike train and performing a Monte Carlo simulation to evaluate the statistical significance of the discrepancy between the model and the spike train series. By using the rate-rescaling theorem to construct the Q–Q and K–S plots, our analysis obviates the power ratio's need for additional Fourier and Monte Carlo computations to assess model agreement with the spike train. Moreover, the power ratio analysis does not provide a graphical assessment of goodness-of-fit comparable to the Q–Q and K–S plots.

These methods are part of our research on analyzing spatial information encoding by neurons in the rat hippocampus. In this example all the goodness-of-fit measures showed that both the IG and IIG models provide significant improvement over the IP models

with and without theta rhythm modulation previously described by Brown et al. (1998b). According to both AIC/BIC and K–S plot analyses, 11 of the 34 cells were best described by the IG model, and 21 were best described by the IIG model. For the two cells classified differently by the AIC/BIC and the K–S plot criteria, the respective values of the statistics for the two models were quite close. In no case was a form of the Poisson model superior to both the IG and IIG models. The Q–Q plot analysis was completely consistent with the AIC/BIC and K–S plot analyses. Without including theta phase, the IG and IIG models are superior to the IP plus theta phase model for all the cells. This suggests that an improved model of the intrinsic mechanism explains better the spiking activity than one that includes an exogenous variable. This result is compelling because it holds for all 34 of the place cells analyzed. It has previously been shown that inclusion of theta phase made only a small improvement in the fit of the IP model because in open field experiments theta modulation is weaker (Skaggs et al., 1996; Brown et al., 1998a) than on the linear track. All the improvements in fit of the IG and IIG models to the data are due to giving a better description of the intrinsic spiking activity of the neurons.

The burst analysis suggested that the aspect of the intrinsic mechanism which the IIG model describes better are the short ISIs, or bursts, whereas the IG model describes most accurately the longer ISIs. This finding suggests that the bursts are a significant component of the stochastic structure in the place cell spike trains even when other factors, such as position and theta phase, are taken into account. It supports the hypothesis that these short ISIs may carry important information about features of the position stimulus (Buzsaki et al., 1996).

The rate map analysis shows that unlike the IP model the IG and IIG models give spatial and temporal rate maps that are history dependent. The fact that the IG and IIG models give better statistical fits to the place cell data than the IP model suggests that the rate maps from the former models provide more accurate descriptions of these local spiking properties of the neuron. History-dependent spatial rate maps are more physiologically plausible because space and time must interact to define the firing rate of these cells. For example, the expected rate of spiking during burst activity in the center of a place field should be different from that in the periphery of the field. To the authors' knowledge these results are the first presentation of history dependent rate maps.

Because neither of the new models completely describes all the structure in the place cell spiking activity, it remains to define what are the other properties of the place cells the analysis should characterize. In the immediate neighborhood of the origin, the lack-of-fit of

the IG model is due in part to the infinite mode of the standard gamma probability model for $\gamma < 1$. Because the IG and IIG models describe better different sets of ISI's, the optimal model of place cell spiking activity may combine the properties of the IG and IIG models. Such a model could be constructed by developing an inhomogeneous model that combines the intensity-rescaling transformation with the generalized inverse Gaussian probability model. The gamma and inverse Gaussian models are special cases of the generalized inverse Gaussian probability density (Iyengar and Liao, 1997). We are currently studying generalized inhomogeneous Gaussian models as a basis for constructing more accurate models of the place-specific firing activity of individual hippocampal neurons (Barbieri et al., 2000).

The remaining lack-of-fit for the IG and IIG models may be due to inaccuracy of the Gaussian surface as a model of the place field spatial structure, as well as omission from the current model of covariates such as the animal's direction of motion, running velocity (Zhang et al., 1998), and the phase precession effect (O'Keefe and Recce, 1993). Another possibility is that the history dependence in the place cell spiking activity may be of a higher order than Markov. The importance of these factors is also being investigated systematically using the current framework.

Although our paradigm was illustrated with an analysis of hippocampal place cells, it may be applied to any experiment in which spike train activity is recorded simultaneously with input stimuli and relevant covariates. Therefore, the methods should offer a useful set of generally applicable tools for neural spike train data analysis.

Acknowledgements

We thank the anonymous referees for suggestions that helped improve the accuracy and presentation in the manuscript. Support was provided in part by National Institutes of Mental Health grants MH59733 and MH61637, National Science Foundation grant IBN-0081548, DARPA, the Office of Naval Research, and the Center for Learning and Memory at MIT.

Appendix A

The joint probability density of a set of spike times for an arbitrary stimulus input over a fixed time interval is derived. This is termed the sample path probability density (Snyder and Miller, 1991). First, it is shown that any ISI probability density may be expressed in terms of the hazard or rate function (Eq. (8)). The hazard function generalizes for an arbitrary ISI distribution the Poisson rate function for the ISI probability

density in Eq. (5). The ISI probability density expressed in terms of the hazard function resembles the ISI distribution for the IP model. This derivation is based on Kalbfleisch and Prentice (1980) and Daley and Vere-Jones (1988). On an observation interval $(0, T]$, given n spike times $0 < t_1 < t_2 < \dots < t_n < T$, let H_k denote the history of the spikes up to t_{k-1} and the history of the stimulus up to time t_k or T , whichever is shorter. A more precise definition of H_k in terms of sigma fields is given in Daley and Vere-Jones (1988). The joint probability density of exactly n spikes in $(0, T]$

$$\begin{aligned} f(t_1, \dots, t_n \cap t_{n+1} > T) &= \prod_{k=1}^n f(t_k | t_1, \dots, t_{k-1}, x(t) \text{ on } (0, t_k]) \\ &\quad \times \Pr(\text{no spikes in } (t_n, T] | t_1, \dots, t_n, x(t) \text{ on } (0, T]) \\ &= \prod_{k=1}^n f(t_k | H_k) [1 - F(T | H_{n+1})], \end{aligned} \quad (\text{A.1})$$

where

$$[1 - F(T | H_{n+1})] = \int_{t_n}^T f(u | H_{n+1}) du.$$

From Eq. (8) one has for $t > t_{k-1}$

$$r(t | H_k) = \frac{f(t | H_k)}{1 - F(t | H_k)}, \quad (\text{A.2})$$

and

$$r^*(t) = \begin{cases} r(t) & 0 < t < t_1 \\ r(t | H_k) & t_{k-1} < t < t_k \end{cases} \quad (\text{A.3})$$

Eq. (A.3) is termed the hazard function, conditional intensity function or stochastic intensity function. To show that the hazard function and the intensity function are the same for the IP model Eq. (5) is substituted in Eq. (A.2) to obtain

$$\begin{aligned} r(t | H_k) &= \frac{\lambda(t) \exp\left\{-\int_{t_{k-1}}^t \lambda(u) du\right\}}{1 - F(t | H_k)} \\ &= \frac{\lambda(t) \exp\left\{-\int_{t_{k-1}}^t \lambda(u) du\right\}}{\exp\left\{-\int_{t_{k-1}}^t \lambda(u) du\right\}} \\ &= \lambda(t). \end{aligned} \quad (\text{A.4})$$

Integrating both sides of Eq. (A.2) from t_{k-1} to t gives

$$-\log[1 - F(t | H_k)] = \int_{t_{k-1}}^t r^*(u) du, \quad (\text{A.5})$$

and exponentiating both sides yields

$$[1 - F(t | H_k)] = \exp\left\{-\int_{t_{k-1}}^t r^*(u) du\right\}. \quad (\text{A.6})$$

Rearranging Eq. (A.2) and using Eq. (A.3) one has

$$\begin{aligned} f(t_k | H_k) &= r(t | H_k) [1 - F(t | H_k)] \\ &= r^*(t) \exp\left\{-\int_{t_{k-1}}^t r^*(u) du\right\}, \end{aligned} \quad (\text{A.7})$$

which is the first part of the desired result. Returning to Eq. (A.1) and using Eq. (A.7), the joint probability density of exactly n spikes in $(0, T]$ is

$$\begin{aligned} f(t_1, t_2, \dots, t_n \cap t_{n+1} > T) &= \prod_{k=1}^n f(t_k | H_k) [1 - F(T | H_{n+1})] \\ &= \prod_{k=1}^n r^*(t_k) \exp\left\{-\int_0^T r^*(u) du\right\}. \end{aligned} \quad (\text{A.8})$$

Eqs. (A.7) and (A.8) show that an arbitrary spike time probability density or joint probability density can be expressed in terms of its associated hazard (rate) function.

References

- Barbieri R, Frank LM, Quirk MC, Wilson MA, Brown EN. Construction and analysis of the inhomogeneous general inverse Gaussian probability model of place cell spiking activity. *Soc Neurosci Abstr* 2000;26:716.
- Berman M. Inhomogeneous and modulated gamma processes. *Biometrika* 1981;68:143–52.
- Berman M. Comment on ‘likelihood analysis of point processes and its applications to seismological data’ by Y Ogata. *Bull Int Stat Inst* 1983;50(3):412–8.
- Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D. Reading a neural code. *Science* 1991;252:1854–7.
- Bishop PO, Levick WR, Williams WO. Statistical analysis of the dark discharge of lateral geniculate neurones. *J Physiol* 1964;170:598–612.
- Box GEP, Jenkins GM, Reinsel GC. *Time Series: Forecasting and Control*, 3rd edition. Englewood Cliffs, NJ: Prentice Hall, 1994.
- Brillinger DR. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biol Cyber* 1988;59:189–200.
- Brown EN, Frank LM, Barbieri R, Quirk MC, Wilson MA. A time-dependent gamma distribution model of spike train activity in hippocampal place cells. *Soc Neurosci Abstr* 1998a;24:931.
- Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J Neurosci* 1998b;18:7411–25.
- Buzsaki G, Penttonen M, Nadasdy Z, Bragin A. Pattern and inhibition-dependent invasion of pyramidal cell dendrites by fast spikes in the hippocampus in vivo. *Proc Natl Acad Sci* 1996;93:9921–5.
- Casella G, Berger RL. *Statistical Inference*. Belmont, CA: Duxbury, 1990.
- Chhikara RS, Folks JL. *The Inverse Gaussian Distribution: Theory, Methodology and Applications*. New York: Marcell-Dekker, 1989.
- Correia MJ, Landolt JP. A point process analysis of the spontaneous activity of anterior semicircular canal units in the anesthetized pigeon. *Biol Cyber* 1977;27:199–213.
- Daley DJ, Vere-Jones D. *An Introduction to the Theory of Point Process*. New York: Springer-Verlag, 1988.

- Fenton AA, Muller RU. Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proc Natl Acad Sci USA* 1998;95:3182–7.
- Gabbiani F, Koch C. Principles of spike train analysis. In: Koch C, Segev I, editors. *Methods in Neuronal Modeling: From Ions to Networks*, 2nd edition. Cambridge MA: MIT, 1998:313–60.
- Iyengar S, Liao Q. Modeling neural activity using the generalized inverse Gaussian distribution. *Biol Cyber* 1997;77:289–95.
- Johnson NL, Kotz S. *Continuous Univariate Distributions — 1*. New York: Wiley, 1971.
- Kalbfleisch JD, Prentice RL. *The Statistical Analysis of Failure Time Data*. New York: Wiley, 1980.
- Kass RE, Ventura V. A spike train probability model, *Neural Computation*. In press.
- Meyer PA. Demonstration simplifiée d'un théorème de Knight, *Seminaire de Probabilités V, Lecture Notes in Mathematics*. New York: Springer, 1969:191–5.
- Miller JP, Jacobs GA, Theunissen FE. Representation of sensory information in the cricket cercal sensory system. I: Response properties of the primary interneurons. *J Neurophys* 1991;66:1680–703.
- Muller RU, Kubie JL, Ranck JB. Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J Neurosci* 1987;7:1935–50.
- Nakahama H, Suzuki H, Yamamoto M, Aikawa S, Nishioka S. A statistical analysis of spontaneous activity of central single neurons. *Physiol Behav* 1968;3:745–52.
- Ogata Y. Statistical models for earthquake occurrences and residual analysis for point processes. *J Am Stat Assoc* 1988;83:9–27.
- O'Keefe J, Dostrovsky J. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 1971;34:171–5.
- O'Keefe J, Recce ML. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 1993;3:317–30.
- Pfeiffer RR, Kiang NY-S. Spike discharge patterns of spontaneous and continuously stimulated activity in the cochlear nucleus of anesthetized cats. *Biophys J* 1965;5:301–16.
- Port SC. *Theoretical Probability for Applications*. New York: Wiley, 1994.
- Reich DS, Victor JD, Knight BW. The power ratio and the interval map: spiking models and extracellular recordings. *J Neurosci* 1998;18:10090–104.
- Rieke F, Bodnar D, Bialek W. Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory neurons. *Proc R Soc Lond Ser B* 1995;262:259–65.
- Rieke F, Warland D, de Ruyter van Steveninck RR, Bialek W. *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT, 1997.
- Rinberg D, Davidowitz H. Do cockroaches 'know' about fluid dynamics? *Nature* 2000;405:756.
- Rodieck RW, Kiang NY-S, Gerstein GL. Some quantitative methods for the study of spontaneous activity of single neurons. *Biophys J* 1962;2:351–68.
- Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 1998;18:3870–96.
- Shwartz G. Estimating the dimension of a model. *Ann Stat* 1978;6:461–4.
- Skaggs WE, McNaughton BL, Wilson MA, Barnes CA. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 1996;6:149–72.
- Snyder DL, Miller MI. *Random Point Processes in Time and Space*, 2nd edition. New York: Springer-Verlag, 1991.
- Tuckwell H. *Introduction to Theoretical Neurobiology*, vol. 2. New York: Cambridge University Press, 1988.
- Van Steveninck RRdR, Lewen GD, Strong SP, Koberle R, Bialek W. Reproducibility and variability in neural spike trains. *Science* 1997;275:1805–8.
- Wilks MB, Gnanadesikan R, Huyett MB. Probability plots for the gamma distribution. *Technometrics* 1962;4:1–20.
- Wilson MA, McNaughton BL. Dynamics of the hippocampal ensemble code for space. *Science* 1993;261:1055–8.
- Zhang KC, Ginzburg I, McNaughton BL, Sejnowski TJ. Interpreting neuronal population activity by reconstruction: a unified framework with application to hippocampal place cells. *J Neurophys* 1998;79:1017–44.