A Spike Train Probability Model

Robert E. Kass and Valérie Ventura
Department of Statistics and Center for the Neural Basis of Cognition
Carnegie Mellon University

Abstract

Poisson processes usually provide adequate descriptions of the irregularity in neuron spike times after pooling the data across large numbers of trials, as is done in constructing the peristimulus time histogram. When probabilities are needed to describe the behavior of neurons within individual trials, however, Poisson process models are often inadequate. In principle, an explicit formula gives the probability density of a single spike train in great generality, but without additional assumptions the firing-rate intensity function appearing in that formula can not be estimated. We propose a simple solution to this problem, which is to assume that the time at which a neuron fires is determined probabilistically by, and only by, two quantities: the experimental clock time and the elapsed time since the previous spike. We show that this model may be used successfully to fit neuronal data.

1 Introduction

Probability distributions of spike trains are used for a variety of theoretical and data-analytical purposes including determination of information content, Bayesian population coding, testing for spike patterns, and detecting synchrony (Barbieri *et al.*, 1999; Brown *et al.*, 1998; Zhang *et al.*, 1998; Sanger, 1996; Oram *et al.*, 1999; Optican and Richmond, 1987; Riehle *et al.*, 1998). The peristimulus time histogram (PSTH) provides an estimate of the time-varying firing rate (or "intensity" in the jargon of probability theory) and for Poisson processes the firing rate, in turn, determines the spike-train distribution. When large numbers of trials are combined, the resulting set of pooled spike times essentially follow a Poisson process, due to general limit theory (Daley and Vere-Jones, 1988, Theorem 9.2V). Thus, when reasoning only from data pooled across trials, or when the spike trains on individual trials do themselves follow a Poisson process, the PSTH (or smoothed versions of it) provides a complete summary of the data from which other probability-based calculations may be made. However, when neuron firing on individual trials is of interest and the process is non-Poisson, the joint distribution of the spikes can not be obtained from the PSTH; nor can it be obtained from the PSTH in combination with the distribution of interspike intervals (ISIs). The purpose of this note is to suggest a class of probability models for spike trains that allow recovery of the joint spike-train distributions within individual trials, and to report results from fitting these models to neuronal data.

2 Inhomogeneous Markov Interval Processes

Letting $N(t)$ represent the number of spikes that have already occurred at time t on a particular trial, a spike occurs at time t if $N(t+\Delta t) - N(t) = 1$ for all sufficiently small increments Δt. In the language of probability theory, the spike times follow a point process and $N(t)$ is the corresponding counting process (e.g., Daley and Vere-Jones, 1988, Chapter 13). If $s(t) = (s_1, s_2, \ldots, s_{N(t)})$ denotes the spike times up to time t, then the probability density of the spike train during an interval $[0, T]$ is

$$p(s_1, \ldots, s_n) = e^{-\int_0^T \lambda(u|s(u))du} \prod_{k=1}^{n} \lambda(s_k \mid s_1, \ldots, s_{k-1}) \tag{1}$$

where

$$\lambda(t \mid s(t)) = \lim_{\Delta t \to 0} \frac{P(N(t+\Delta t) - N(t) = 1 \mid s(t))}{\Delta t}$$

is the conditional intensity of the process and $N(T) = n$. In principle, this solves the problem of representing the spike trains in terms of a probability distribution. However, when n spikes are observed the conditional intensity is a function of $n+1$ variables $(t, s_1, \ldots, s_n)$, which poses an extremely difficult nonparametric estimation problem.

A tractable simplification is to take the conditional intensities to have the form

$$\lambda(t \mid s(t)) = \lambda(t, t - s_*(t)) \tag{2}$$

where $s_*(t)$ is the last spike time preceding t. We call the resulting processes inhomogeneous Markov interval (IMI) processes. When $\lambda(t, t - s_*(t))$ is a function only of its first argument, i.e., the firing rate does not depend on occurrence times of previous spikes, we obtain an inhomogeneous (time-varying) Poisson process: the beauty of a Poisson process is that the spike probabilities are determined solely by the time-varying firing rate, or intensity, which may be written as $\lambda(t)$ and is estimated by the smoothed

PSTH. The special case of a *homogeneous* Poisson process occurs when there is also no dependence on experimental clock time, meaning that the firing rate is constant across time. In this case, the ISIs are independent and all follow the same exponential distribution. A generalization, not requiring the Poisson assumption, is the class of renewal processes, in which the ISIs all follow the same probability distribution, but it is no longer required to be an exponential distribution. These processes fit into (2) when $\lambda(t, t - s_*(t))$ is a function only of its second argument, i.e., the firing rate does not depend on the experimental clock time.

An interesting sub-class of IMI processes are the time-rescaled inhomogeneous versions of renewal processes recently used by Barbieri, *et al.* (2000). These models begin with the fact that an inhomogeneous Poisson spike train is obtained from a sequence of exponential ISIs together with a rescaling of time.[1] A different ISI distribution, such as a gamma distribution, may be used in place of the exponential and time may again be transformed in the same way.

An alternative sub-class that may be treated nonparametrically (i.e., does not require the assumption of a specific ISI distribution) are the multiplicative IMI processes in which

$$\lambda(t, t - s_*(t)) = \lambda_1(t) \cdot \lambda_2(t - s_*(t)). \tag{3}$$

Here, the factor $\lambda_1(t)$ modulates the firing rate only as a function of experimental clock time while $\lambda_2(t - s_*(t))$ represents non-Poisson spiking behavior. Note, however, that the observed ISI distribution is based on both factors. The units of $\lambda(t, t - s_*(t))$ are those of firing rate: spikes per unit time. By convention, in our analysis of neuronal data reported below we will take $\lambda_1(t)$ also to have units of firing rate, leaving $\lambda_2(t - s_*(t))$ dimensionless.

3 Fitting IMI Processes to Data

In general, a spike train in the form of a sequence of "exact" spike times $(s_1, s_2, \ldots, s_n)$ may be approximated by discretizing time into small intervals of length Δt (e.g., $\Delta t = 1$ msec) and converting to a binary sequence of 0s and 1s, with each 1 indicating that a spike occurred within the corresponding time interval. As Brillinger (1988) observed, this is useful statistically

[1] Specifically, the probability density for the inhomogeneous Poisson is given by (1) when $\lambda(t|s(t)) = \lambda(t)$; the inhomogeneous Poisson reduces to a homogeneous Poisson after the change of variables in time from t to τ according to $\tau(t) = \int_0^t \lambda(u)du$.

because the binary sequence may be fitted using generalized linear models (McCullagh and Nelder, 1989), with maximum likelihood in this generalized form of regression playing a role analogous to least squares in ordinary linear regression. (Formally, it is not hard to show that the likelihood function from (1) may be approximated by the corresponding binary regression likelihood function.) This implies that IMI processes may be fitted with standard methods and software. In particular, to fit the multiplicative IMI processes in (3), we use the additive form

$$\log \lambda(t, t - s_*(t)) = \log \lambda_1(t) + \log \lambda_2(t - s_*(t)). \tag{4}$$

and define suitable functional representations for each of the components $\log \lambda_1(t)$ and $\log \lambda_2(t - s_*(t))$, which will play a role analogous to explanatory variables in ordinary linear regression. An especially flexible and tractable way to represent the components is with cubic splines. That is, we represent each component $\log \lambda_1(t)$ and $\log \lambda_2(t - s_*(t))$ as a piecewise cubic in its argument (respectively, t and $t - s_*(t)$), with the cubic pieces joined at "knots" in such a way that the resulting functions are twice continuously differentiable. One may choose the knots by preliminary examination of the data. Then, the coefficients of the spline basis elements are determined via maximum likelihood; we have used the function `gam` in the software S-PLUS (MathSoft Inc., Seattle), as described in Chapter 11 of Venables and Ripley (2000). The inputs to `gam` are the discretized binary sequence of spike times, the time arguments t and $t - s_*(t)$, the symbolic specification of the additive model in (4), and the knots. Further details on the use of spline-based regression methods (or "regression splines" in statistical parlance) for analyzing neuronal data may be found in Olson, *et al.* (2000) and Ventura *et al.* (2000), where they are applied in the case of inhomogeneous Poisson processes.

To fit the more general IMI processes defined in (2), terms formed by taking products ("interactions") of the spline representations for $\log \lambda_1(t)$ and $\log \lambda_2(t - s_*(t))$ may be included in the regression function. In addition, terms may be added to the statistical model that reflect dependence on timing of spikes prior to the most recent one.

We have not yet been explicit about the way the intensity (2) might vary from trial to trial. The simplest way to allow for excess trial-to-trial variability is to include multiplicative constants for each trial in (2), which become additive constants on the log scale.

4 Results

We examined data recorded from a neuron in the supplementary eye field of a Macaque monkey while he was carrying out 15 trials of a delayed eye movement task (neuron PK166a.1 from Olson *et al.*, 2000; we used the pattern condition). We fitted the IMI process intensity function (2) as described above. We included (i) spline basis elements for $\log \lambda_1(t)$, (ii) spline basis elements for $\log \lambda_2(t - s_*(t))$, (iii) cross-products between the spline basis elements for $\log \lambda_1(t)$ and those for $\log \lambda_2(t - s_*(t))$ (interaction terms), (iv) additional basis elements to represent functions of time since the spike $s_{**}(t)$ prior to $s_*(t)$ and time since the spike $s_{***}(t)$ prior to $s_{**}(t)$, and also (v) constants for each trial. Using (i) alone would result in fitting an inhomogeneous Poisson process, (ii) represents departures from a Poisson process, (iii) represents departures from (3) within the general framework (2), (iv) represents departures from (2), and (v) represents excess trial-to-trial variability. In addition to the overall constant, there were 4 degrees of freedom for (i) resulting from the use of 3 knots,[2] 3 for (ii) resulting from the use of 2 knots, 12 for (iii) (the coefficients for all products between terms in (i) and (ii)), 3 each for the two components in (iv) (each based on 2 knots), and 14 for (v) (because there were 15 trials).

For this neuron we found that the additive IMI process (3) without excess trial-to-trial variability clearly provided the best fit of the models considered: none of the components (iii)-(v) was statistically significant, but the components (i) and (ii) were highly significant ($p < .0005$). Figure 1 displays the firing behavior of this neuron. As is typically the case, the fit of a Poisson model (effectively applying (2) with only the first component $\lambda_1(t)$) to the data pooled across trials is good: in part (A) the fitted spline simply smooths the PSTH. Furthermore, when the 15 within-trial IMI fits are averaged across trials the result agrees very closely with the Poisson model. However, the effect of the highly significant non-Poisson component (ii), shown in part (B), may be understood when we consider the way the neuron behaves following a spike at $t = 50$ ms, as shown in part (C): there it may be seen that the neuron has a refractory period of roughly 10 ms during which it is less likely to fire again than predicted by the Poisson model. After the firing rate reaches its peak at roughly 100 ms, the effect of the non-Poisson component diminishes over time. The effect over the course of a single trial is illustrated in part (D). Immediately after each occurrence

[2]Essentially, we have linear, quadratic, and $k + 1$ cubic coefficients for k knots, and then lose degrees of freedom because of the differentiability constraints; technically, in S-PLUS we use the natural spline basis.

of a spike the neuron partially "resets," lowering its firing intensity below the average firing across trials predicted by the smoothed PSTH, then it increases to a rate higher than the PSTH-based prediction until the neuron fires again.[3]

5 Discussion

The main purpose of this note was to suggest the class of IMI processes defined in (2) as potentially useful representations of neuron spike trains. These statistical models generalize inhomogeneous Poisson processes, allowing computation of quantities defined by the spike train probability distribution once the IMI intensity function $\lambda(t, t - s_*(t))$ is estimated from the data. For example, it is now possible to generalize the test for synchrony used by Riehle *et al.* (1997), avoiding the possibly-objectionable Poisson assumption. We hope to report on this in a future communication. The "Markov Interval" terminology comes from Cox and Lewis (1972), who discussed the time-homogeneous case for coupled processes; we are currently investigating extensions of IMI processes for multiple neurons.

Estimation of $\lambda(t, t - s_*(t))$ could be carried out by various methods. We have here used cubic splines on the log scale estimated by maximum likelihood within the statistical framework of generalized linear models, a standard statistical methodology for which software is widely available. We found the multiplicative model (3) to have fit well for a particular neuron in the SEF, but we would like to emphasize that this multiplicative form may or may not be an appropriate simplification of (2) in other settings—this is an empirical issue that must be examined in each case. Substantively, as illustrated in part (D) of Figure 1, we found that this neuron has much more rapid fluctuations in its firing rate within trials than predicted by the Poisson model, which simply smooths the PSTH.

The fundamental utility of IMI processes is that they avoid the assumption that spike trains are Poisson processes, which fails to account for important effects such as the refractory period. Instead, they make a weaker and seemingly reasonable assumption that spike timing is determined by, and only by, experimental clock time and the elapsed time since the pre-

[3] An absolute refractory period, during which the neuron would not fire, would force λ_2 to zero as its argument $t - s_*(t)$ goes to zero. However, the data from Olson *et al.* (2000) were acquired only to 1 msec accuracy and there is no evidence here of an absolute refractory period longer than 1 msec. Thus, our fitted function λ_2 in part (B) does not vanish at the origin.

vious spike; there is no further restriction. Statistically, the main point is that the IMI assumption reduces the fitting problem posed in (1) to a very manageable two-variable nonparametric binary regression. We solved this problem using splines with knots fixed at suitable locations. We are currently developing a more fully automated procedure based on the method of DiMatteo, Genovese, and Kass (2000).

ACKNOWLEGMENTS. We are grateful to Emery Brown and John Lehoczky for many helpful conversations. Comments from the referees substantially improved the exposition. This work was supported in part by grants CA54652-08 from the National Institutes of Health and DMS-9803433 from the National Science Foundation.

References

Barbieri, R., Quirk, M.C., Frank, L.M., Wilson, M.A., and Brown, E.N. (2000) Non-Poisson stimulus-response models of neural spike train activity. (*Submitted*)

Brillinger, D.R. (1988) Maximum likelihood analysis of spike trains of interacting nerve cells. *Biol. Cyber.*, **59**, 189-200.

Brown, E.N., Frank, L.M., Tang, D., Quirk, M.C. and 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. *J. Neurosci.*, **18**, 7411-7425.

Cox, D.R. and Lewis, P.A.W. (1972) Multivariate point processes, *Proc. 6th Berkeley Symp. Math. Statist. Prob.* **3**, 401-448. University of California Press.

Daley, D.J. and Vere-Jones, D. (1988) *Introduction to the Theory of Point Processes*, New York: Springer.

DiMatteo, I., Genovese, C.R., and Kass, R.E. (2000) Bayesian curve fitting with free-knot splines. (*Submitted*)

McCullagh, P. and Nelder, J.A. (1989) *Generalized Linear Models*, Second Edition, New York: Chapman and Hall.

Olson, C.R., Gettner, S.N., Ventura, V., Carta, R. and Kass, R.E. (2000) Neuronal activity in Macaque supplementary eye field during planning of saccades in response to pattern and spatial cues. *J. Neurophys.*, **84**, 1369-1384.

Optican, L.M. and Richmond, B.J. (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J. Neurophysiol.*, **57**, 162-178.

Oram, M.W., Wiener, M.C., Lestienne, R., and Richmond, B.J. (1999) Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J. Neurophysiol.*, **81**, 3021-3033.

Riehle, A., Grün, S. Diesmann, M. and Aertsen, A. (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* **278,** 1950–1953.

Sanger, T.D. (1996) Probability density estimation for the interpretation of neural population codes. *J. Neurophysiol.*, **76**, 2790-2793.

Venables, W.N. and Ripley, B.D. (2000) *Modern Applied Statistics with S-PLUS*, Third Edition. New York: Springer.

Ventura, V., Carta, R., Kass, R.E., Gettner, S.N. and Olson, C.R. (2000) Statistical analysis of temporal evolution in single-neuron firing rates. *Submitted*

Zhang, K., Ginzburg, I., McNaughton, B.L. and Sejnowski, T.J. (1998) Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *J. Neurophysiol.*, **79**, 1017-1044.

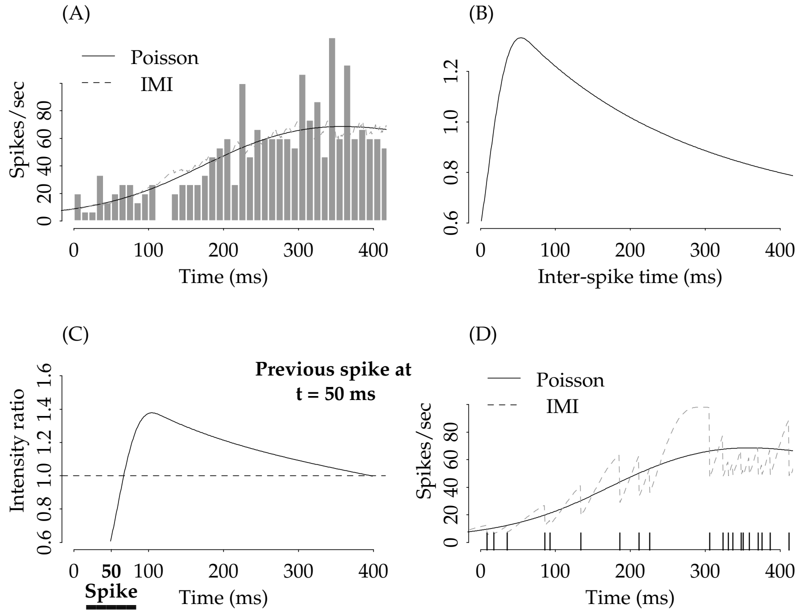

Figure 1: Firing behavior of the SEF neuron. In part (A) the PSTH is shown together with the fitted intensity assuming a Poisson process and the average across the 15 trials of the fitted intensities from the multiplicative IMI model (3). The fitted factor $\lambda_2(t - s_*(t))$, which modulates the time-varying firing rate to account for non-Poisson spiking behavior, is shown in part (B). If the spike trains followed a Poisson process, this intensity would be constant and equal to 1. Instead, the cell begins (during its refractory period) with a diminished firing intensity relative to what would be predicted by the Poisson assumption; this is followed by a fairly

rapid increase in firing intensity until about 50 ms at which time the intensity slowly declines. (To ease comparison with part (C) we have scaled $\lambda_2(t - s_*(t))$ in part (B) so that the value $\lambda_2(t - s_*(t)) = 1$ corresponds to the fitted Poisson intensity at $t = 50$ ms; equivalently, $\lambda_1(t)$ is equal to the fitted Poisson intensity at $t = 50$ ms.) The effect of $\lambda_2(t - s_*(t))$ on this neuron's firing rate is shown in parts (C) and (D). According to the fitted model (3) the IMI firing intensity is the product of the fitted factor $\lambda_2(t - s_*(t))$ shown in part (B) and the fitted factor $\lambda_1(t)$ (which is not shown); because of the former it depends on the time at which the previous spike occurred. Part (C) displays the ratio of IMI to Poisson fitted firing intensities when the previous spike occurred at $t = 50$ ms. As implied by part (B), during the first 10 ms following the spike at $t = 50$ ms the cell is less likely to fire again than predicted by the Poisson model. Part (D) shows the fitted IMI intensity for the first trial together with the fitted Poisson intensity based on pooling all 15 trials (as in part (A)). The bars along the horizontal axis indicate the times at which spikes occurred on this trial. After each spike the firing rate immediately drops, then climbs to a rate substantially above the predicted Poisson rate.