

Model based decoding of spike trains

Matthew C. Wiener¹, Barry J. Richmond*

Laboratory of Neuropsychology, National Institute of Mental Health, National Institutes of Health, Bldg 49, Rm 1B80, Bethesda, MD 20892, USA

Accepted 22 August 2002

Abstract

Reliably decoding neuronal responses requires knowing what aspects of neuronal responses are stimulus related, and which aspects act as noise. Recent work shows that spike trains can be viewed as stochastic samples from the rate variation function, as estimated by the time dependent spike density function (or normalized peristimulus time histogram). Such spike trains are exactly described by order statistics, and can be decoded millisecond-by-millisecond by iterative application of order statistics.

© 2002 Elsevier Science Ireland Ltd. All rights reserved.

Keywords: V1; Vision; Coding

1. Introduction

The amount of information a spike train conveys depends on how the spike train is read, or decoded. Some aspects of the spike train are considered important for carrying stimulus-related information and used for decoding, while other aspects are considered noise and ignored. Frequently the number of spikes arriving in a relatively long time window (spike count) is considered the information-carrying code, and spike timing is disregarded. However, it is known

that at least some aspects of the time course of firing, such as the variation in firing rate, also carry stimulus-related information not available from spike count alone (Richmond et al., 1987; Richmond and Optican, 1990; McClurkin et al., 1991; Victor and Purpura, 1996; Tovee et al., 1993). Fig. 1 shows an example in which the spike count distributions elicited by two stimuli are identical (and therefore, spike count can carry no information distinguishing the two stimuli) but the firing rate over time differs, and could be used to distinguish the two stimuli.

The example in Fig. 1 shows that modulation in the firing rate over time can carry stimulus-related information. However, the responses in the lower panels of Fig. 1 were low-pass filtered with a Gaussian ($\sigma = 5$ ms, 3 db point < 20 Hz), so this is a coarse representation of the timing in the responses. It has also been proposed that precise

* Corresponding author. Tel.: +1-301-496-5626; fax: +1-301-402-0046

E-mail address: bjr@ln.nimh.nih.gov (B.J. Richmond).

¹ Present address: Applied Computer Science and Mathematics Department, RY84-202 Merck Research Laboratories, Rahway, NJ 07065, USA.

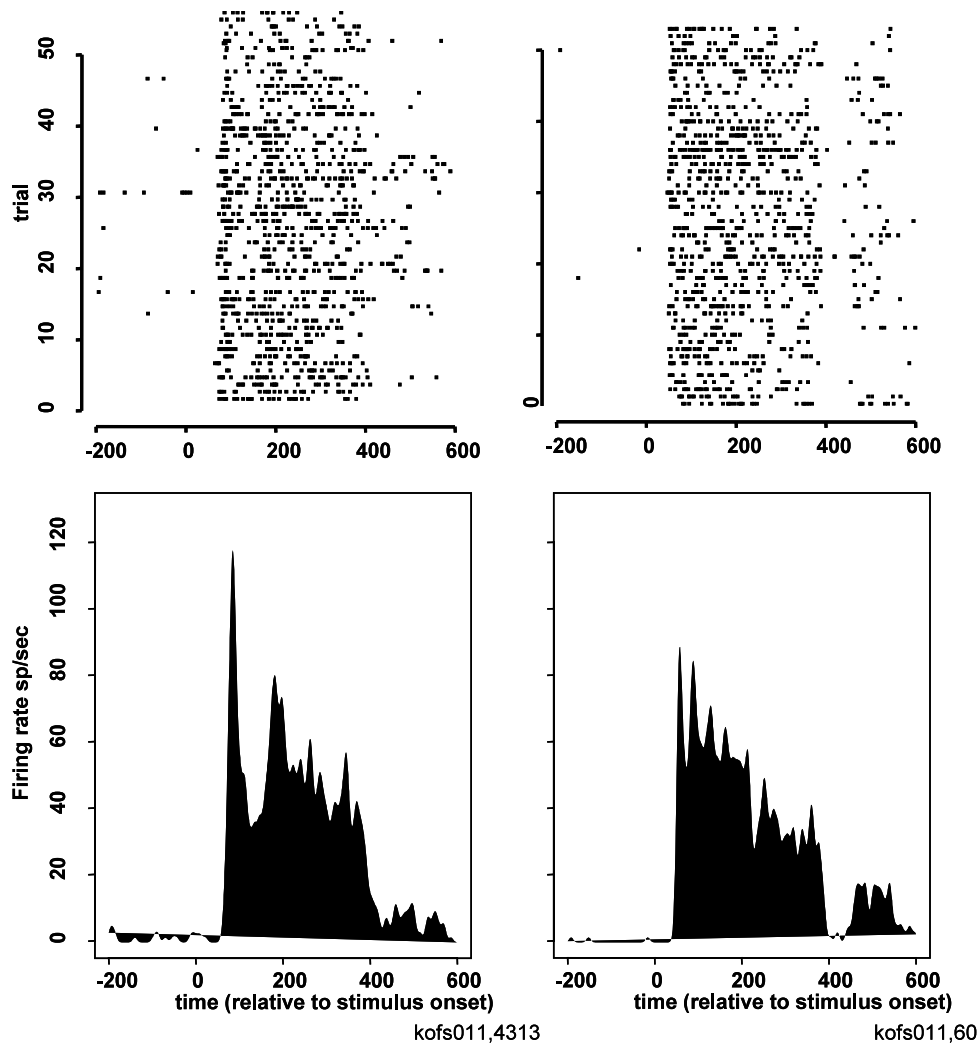


Fig. 1. Responses of a V1 complex cell to two different visual stimuli. The two responses are indistinguishable in spike count. The shapes of the spike densities are different, showing that the changes in rate over time could differentiate the stimuli.

(on the order of 1 ms) time relations among spikes within or between spike trains might also carry information. Specifically, it has been proposed that such relations might underlie higher functions (such as perception of whole objects) requiring large regions of cortex subserved by large numbers of neurons (von der Malsburg, 1995). If precisely timed spike patterns do carry information, we would expect them to occur at frequencies greater than those expected by chance. Determining whether they do occur at greater than chance

frequencies requires knowing how many to expect by chance.

A number of studies over approximately the past 10 years have modeled spike trains as non-homogeneous Poisson processes, and concluded that some patterns of three or four spikes (triplets or quadruplets) occur more frequently than expected by chance (Lestienne and Tuckwell, 1998; Abeles et al., 1993; Riehle et al., 1997; Aertsen et al., 1991; Prut et al., 1998; Martignon et al., 2000). If the inhomogeneous Poisson process model of

spike trains is correct, we would expect the overall distribution of spike counts to be a Poisson distribution. However, it has been shown that spike counts are frequently not Poisson-distributed (Baddeley et al., 1997; Gershon et al., 1998).

Recently Oram et al. (1999) studied whether triplets or quadruplets of spikes occur at levels other than those expected from stochastic sampling of the spike density function. The model used to evaluate how many patterns would occur by

chance was similar to the non-homogeneous Poisson models, but instead of assuming the spike count distribution was Poisson, the model forced the spike count distribution to match the distribution seen in the data. This ‘spike count matched’ model can be implemented by selecting an appropriate set of trials from a set generated by a non-homogeneous Poisson process, although the actual implementation is more efficient (Fig. 2). The spike count matched model generates simulated spike trains with the same spike count distribution and (smoothed) firing rate over time as seen in the data, but with stochastically determined spike times.

Using the spike count matched model for data from V1 and LGN, Oram et al. (1999) were able to generate simulated spike trains in which the numbers and types of triplets and quadruplets were indistinguishable from those observed in recorded spike trains. Similar results were found in primary motor cortex (Oram et al., 2001) and, in preliminary work in this lab, in areas TE and perirhinal cortex. Baker and Lemon (2000), using a related but distinct method, also found that spike patterns occurred at chance levels both within and between neurons. Because the occurrence of spike patterns can be predicted from the spike count distribution and the spike density function, the patterns cannot carry any information not already available from these primary features.

To summarize: the information-carrying aspects of single neuronal responses can be characterized by conventional and relatively easily obtainable measurements. When the spike count distribution, a severely low-pass filtered version of the spike density function, and to a small degree, the interspike interval distribution are taken into account, all of the more complex relations among individual spike times are predictable using a simple stochastic model, that is, the number and types of repeating patterns observed in data are consistent with the predictions of the stochastic model (Oram et al., 1999; Richmond et al., 1999; Oram et al., 2001). This provides a strong and exact null hypothesis for evaluating the results of further experiments.

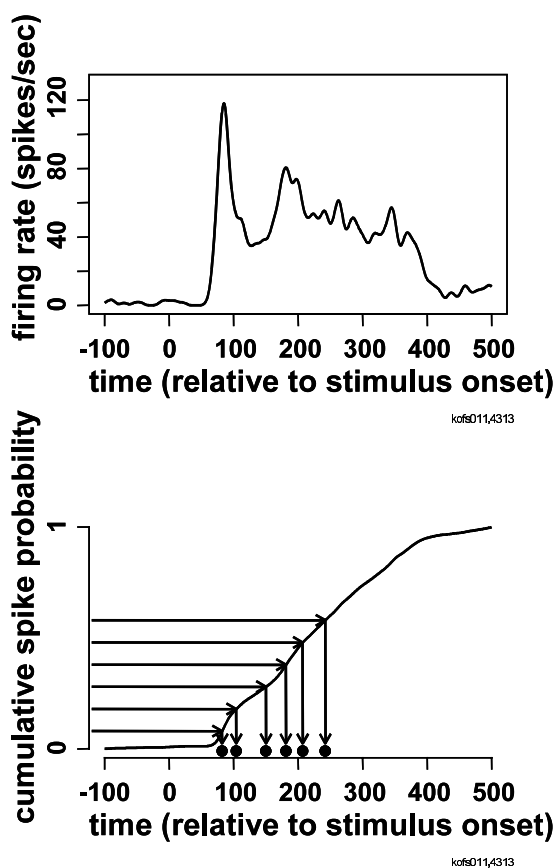


Fig. 2. Left: Spike count matched model. The PSTH above is normalized and integrated to provide the cumulative probability function over the interval (here 500 ms). Enough random numbers (range: 0–1) to match the number needed for the simulated spike train are chosen and mapped through the cumulative distribution (arrows). These new values (x -axis) represent the spike times in the simulated spike train. In this example, six equally spaced numbers between 0 and 1 lead to six spikes with different intervals between them.

2. Order statistics and decoding stochastic spike trains

The success of the spike count matched model shows that spike trains elicited by a stimulus can be described as a series of independent draws from the spike density function. This is exactly the situation described by order statistics (Arnold et al., 1992). Classically, order statistics describe stochastic sampling of continuous processes, in which two draws can be arbitrarily close together. In neuronal responses, spikes are prevented from falling too close together by the refractory period. Order statistics can be adjusted to allow for the refractory period that prevents spikes from falling too close to one another.

The fundamental equation of order statistics (Arnold et al., 1992) is:

$$h_{n,k,s}(t) = \left(\frac{n!}{(k-1)!(n-k)!} \right) F_s(t)^{k-1} f_s(t) \times [1 - F_s(t)]^{n-k} \quad (1)$$

where $h_{n,k,s}(t)$ is the probability of the k th spike occurring at time t in a train with n spikes; f_s , the spike density function; F_s , the corresponding cumulative spike density function; and the subscript s indexes the stimulus. The terms in the equation are a normalizing term for the number of combinations, a term for the number of spikes that have occurred up to time t , the term for the probability of a spike at time t , and a term for the number of spikes that will follow time t . This form is, however, unsuitable because it describes a continuous process. By conditioning on the spikes that have already appeared (actually just the immediately preceding spike because order statistics have a Markov property), we can recursively calculate subsequent order statistics for the ‘next first spike’. The equation then simplifies to:

$$h_{n,1,s}(t_1|t_0) = n[F_s(t_1|t_0)]^0 f_s(t_1|t_0)[1 - F_s(t_1|t_0)]^{n-1} \quad (2)$$

i.e.

$$h_{n,1,s}(t_1|t_0) = n f_s(t_1|t_0)[1 - F_s(t_1|t_0)]^{n-1} \quad (3)$$

This recursive procedure allows inclusion of a refractory period by multiplying the order statistic by a refractory function.

Calculating order statistics requires knowing how many spikes will be in a train. To avoid needing the decoder to know the future, we average the first order statistics $h_{n,1,s}$ for particular spike counts n according to distribution of spike counts, $h_{1,s} = \sum_n p(n)h_{n,1,s}(t)$. To avoid problems based on undersampling, we have used a truncated Gaussian fit to the measured distribution (Gershon et al., 1998; Wiener and Richmond, 1999). Using Bayes rule we can then calculate the probabilities of the stimuli from the spike trains on a trial-by-trial, spike-by-spike basis:

$$p(s|t_1) = \frac{p(s)h_{1,s}(t_1|t_0)}{\sum_{s'} p(s')h_{1,s'}(t_1|t_0)} \quad (4)$$

$p(s|t_1)$ becomes $p(s)$ needed for the next iteration. A similar equation gives the probability when no spike appears at time t , allowing decoding millisecond-by-millisecond, whether or not there is a spike (Fig. 3). These probabilities are those that would be arrived at by an ideal observer. Given a decision, one can then guess which stimulus




Fig. 3. Ms-by-ms decoding, in the sense of classifying among four stimuli using the order statistic model. The x -axis shows time relative to stimulus onset, with the spike train below; the y -axis shows stimulus probability. What appear to be thick lines in the representation of the spike train in fact represent two spikes occurring close together. This shows that, as more spikes arrive, the certainty of the decoding becomes greater, although, in this example, the correct stimulus would have been easily chosen at an early time. Not shown: the correct stimulus is the one with the largest response, which weighs heavily in this result. The a priori response probabilities must be known; we assume that they arise from genetics and/or learning.




Fig. 4. Information (upper panel) and probability of guessing right assuming a maximum likelihood decoder (lower panel) for spike count (squares) and the full spike count plus timing code (filled circles). This shows that for timing, information rises throughout this time period, whereas the other three sets of measures rise relatively quickly and stabilize.

elicited the response. We choose the maximum likelihood rule. Fig. 4 shows that considering the timing significantly increases the information extracted and the percent of trials correctly decoded.

3. Discussion

There is only one assumption needed for order statistics to apply. The spikes need to be describable as stochastic samples from spike probability functions set by the spike density and the spike count distribution. This assumption is likely to be true at some resolution; it is the timing resolution that must be determined. This resolution may well depend on brain region and experimental condition. For example, in the LGN neurons can follow flicker at rates up to 100 Hz (e.g. they follow the refresh rate of TV monitors), whereas in inferior temporal cortex there is no suggestion that neurons do that. In MT neurons are sensitive to some aspects of rapidly changing stimuli (Bair and Koch, 1996; Buracas et al., 1998), yet this sensitivity is modulated by other factors (Bair and Koch, 1996).

Any decoder needs to know something about the responses to be decoded. Our decoder relies on a particular stochastic model: it assumes that spike trains can be described as stochastic samples from the spike density function, with the number of samples in different trains set by the spike count distribution. The spike count distribution and spike density function must be estimated from data, but all other calculations in the decoding are exact. Thus to the extent this model is accurate, the decoding described here extracts all possible information from the spike trains (up to precision of the estimates of spike count distribution and spike density function).

The decoder also relies on a set of a-priori stimulus probabilities. In an experiment these are simply taken to be the relative frequencies of the various stimuli, and are under our control. In nature such priors might be genetically specified or learned. In our experiments we use static stimuli, and are able (because of the Markov property of order statistics) to ignore all spikes but the immediately preceding one. Brown et al. (1998), examining the responses of place fields in rat hippocampus, found that better decoding occurred by taking into account more of the preceding spike (Brown et al., 1998). We have not yet examined the order statistic method with such data.

Finally, it is clear that most brain functions rely on the activity of many neurons. Oram et al. (2001) extended the spike count matched model to investigate spike patterns in pairs of neurons, and found that the number of synchronous spikes seen in pairs of motor cortex neurons exceeded those expected by a constant scale factor. The scaling was such that the synchronous spikes carried no information beyond that available from the spike count distribution (including correlation between spike counts in the two neurons) and the two spike density functions. The synchrony acts as though it arises from an input pool shared by the two neurons, a conclusion similar to that reached by Bair et al. (2001). This suggests that in the future suitable manipulations of order statistics should describe the responses of neuronal populations, also.

References

- Abeles, M., Bergman, H., Margalit, E., Vaadia, E., 1993. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol.* 70, 1629–1638.
- Aertsen, A., Vaadia, E., Abeles, M., Ahissar, E., Bergman, H., Karmon, B., Lavner, Y., Margalit, E., Nelken, I., Rotter, S., 1991. Neural interactions in the frontal cortex of a behaving monkey: signs of dependence on stimulus context and behavioral state. *J. Hirnforsch.* 32, 735–743.
- Arnold, B.C., Balakrishnan, N., Nagaraja, H.N., 1992. *A First Course in Order Statistics*. Wiley, New York.
- Baddeley, R., Abbott, L.F., Booth, M.C., Sengpiel, F., Freeman, T., Wakeman, E.A., Rolls, E.T., 1997. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc. R. Soc. Lond. B Biol. Sci.* 264, 1775–1783.
- Bair, W., Koch, C., 1996. Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput.* 8, 1185–1202.
- Bair, W., Zohary, E., Newsome, W.T., 2001. Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J. Neurosci.* 21, 1676–1697.
- Baker, S.N., Lemon, R.N., 2000. Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance levels. *J. Neurophysiol.* 84, 1770–1780.
- 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. *J. Neurosci.* 18, 7411–7425.
- Buracas, G.T., Zador, A.M., DeWeese, M.R., Albright, T.D., 1998. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 20, 959–969.
- Gershon, E.D., Wiener, M.C., Latham, P.E., Richmond, B.J., 1998. Coding strategies in monkey V1 and inferior temporal cortices. *J. Neurophysiol.* 79, 1135–1144.
- Lestienne, R., Tuckwell, H.C., 1998. The significance of precisely replicating patterns in mammalian CNS spike trains. *Neuroscience* 82, 315–336.
- Martignon, L., Deco, G., Laskey, K., Diamond, M., Freiwald, W., Vaadia, E., 2000. Neural coding: higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comput.* 12, 2621–2653.
- McClurkin, J.W., Gawne, T.J., Richmond, B.J., Optican, L.M., Robinson, D.L., 1991. Lateral geniculate neurons in behaving primates. I. Responses to two-dimensional stimuli. *J. Neurophysiol.* 66, 777–793.
- Oram, M.W., Wiener, M.C., Lestienne, R., Richmond, B.J., 1999. Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J. Neurophysiol.* 81, 3021–3033.
- Oram, M.W., Hatsopoulos, N.G., Richmond, B.J., Donoghue, J.P., 2001. Excess synchrony in motor cortical neurons provides redundant direction information with that from coarse temporal measures. *J. Neurophysiol.* 86, 1700–1716.
- Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Slovin, H., Abeles, M., 1998. Spatiotemporal structure of cortical activity: properties and behavioral relevance. *J. Neurophysiol.* 79, 2857–2874.
- Richmond, B.J., Optican, L.M., 1990. Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. II. Information transmission. *J. Neurophysiol.* 64, 370–380.
- Richmond, B.J., Optican, L.M., Podell, M., Spitzer, H., 1987. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. I. Response characteristics. *J. Neurophysiol.* 57, 132–146.
- Richmond, B.J., Oram, M.W., Wiener, M.C., 1999. Response features determining spike times. *Neural Plasticity* 6, 133–145.
- Riehle, A., Grun, S., Diesmann, M., Aertsen, A., 1997. Spike synchronization and rate modulation differentially involved in motor cortical function [see comments]. *Science* 278, 1950–1953.
- Tovee, M.J., Rolls, E.T., Treves, A., Bellis, R.P., 1993. Information encoding and the responses of single neurons in the primate temporal visual cortex. *J. Neurophysiol.* 70, 640–654.
- Victor, J.D., Purpura, K.P., 1996. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J. Neurophysiol.* 76, 1310–1326.
- von der Malsburg, C., 1995. Binding in models of perception and brain function. *Curr. Opin. Neurobiol.* 5, 520–526.
- Wiener, M.C., Richmond, B.J., 1999. Using response models to estimate channel capacity for neuronal classification of stationary visual stimuli using temporal coding. *J. Neurophysiol.* 82, 2861–2875.