

Response Features Determining Spike Times

Barry J. Richmond,[†] Mike W. Oram, and Matthew C. Wiener

*Laboratory of Neuropsychology, National Institute of Mental Health, Bethesda, MD, USA;
School of Psychology, St. Andrews University, St. Andrews, Scotland*

ABSTRACT

Interpreting messages encoded in single neuronal responses requires knowing which features of the responses carry information. That the number of spikes is an important part of the code has long been obvious. In recent years, it has been shown that modulation of the firing rate with about 25 ms precision carries information that is not available from the total number of spikes across the whole response. It has been proposed that patterns of exactly timed (1 ms precision) spikes, such as repeating triplets or quadruplets, might carry information that is not available from knowing about spike count and rate modulation. A model using the spike count distribution, the low-pass filtered PSTH (bandwidth below 30 Hz), and, to a small degree, the interspike interval distribution predicts the numbers and types of exactly-timed triplets and quadruplets that are indistinguishable from those found in the data. From this it can be concluded that the coarse (<30 Hz) sequential correlation structure over time gives rise to the exactly timed patterns present in the recorded spike trains. Because the coarse temporal structure predicts the fine temporal structure, the information carried by the fine temporal structure must be completely redundant with that carried by the coarse structure. Thus, the existence of precisely timed spike patterns carrying stimulus-related information does not imply control of spike timing at precise time scales.

INTRODUCTION

Interpreting the information encoded in single neuronal responses requires knowing which response features carry information. Despite a great deal of study, which response features are important is still not completely certain. One approach is to identify the smallest number of parameters that are needed to represent all aspects of neuronal responses encoding information. In formal language, we would like to find the minimum description length for representing neuronal responses. In one sense, to describe a neuronal response, we must specify the arrival time of each spike. We are interested less in the spike train itself, however, than in its role in transmitting information. Therefore, only those aspects of the response that carry unique information need be included.

We have used two approaches. First, we measured the information carried by different response representations with the goal of identifying those carrying information. Second, we sought to develop a statistical model that uses only a few experimental measurements to produce simulated spike trains that are indistinguishable from those recorded during our experiments. Below, we first review the parameters that are known to describe spike trains. We then use those parameters to describe a model that produces simulated spike trains that are indistinguishable from the data.

OVERVIEW OF METHODS

We recorded 32 LGN and 19 V1 single neurons in awake, fixating monkeys. Stimuli were presented, one at a time, for 280 to 350 ms

[†]Corresponding author
fax: +1 (301) 402 0046

(depending on the exact experiment) centering on the classical receptive field and covering some of the near surround. The spike times were recorded with 1 ms resolution. The stimulus sets included simple bright and dark bars on a gray background, patches of oriented sine wave gratings from 1 to 6 cycles per degree, and Walsh patterns. In the LGN, the stimuli included bright and dark bars, and center-surround stimuli of both positive and negative contrasts, with centers ranging from smaller than to larger than the neuron's receptive field center. The gratings, Walsh patterns, and center-surround stimuli were contrast-modulated around the background luminance.

Data analysis included statistical descriptions and information theoretic methods. Information theory quantifies how well an output signal, here a neuron's response, can be used to identify the input, here a visual stimulus (Shannon & Weaver, 1949, Cover & Thomas, 1991). Before knowing a neuron's response, we are uncertain about which member of a stimulus set was presented. Information is the reduction in uncertainty about which stimulus was presented, provided by knowing the neuron's response. Intuitively, information quantifies how distinct the responses to different stimuli are. The less the responses to different stimuli overlap, the smaller the chance that more than one stimulus elicited a particular response, and the higher the information. The more the responses to different stimuli overlap, the more difficult it is to determine which stimulus elicited a particular response, and the lower the information. Thus, our minimal description of neuronal responses must include both the average responses to the stimuli and their variabilities, namely, we must know the distribution of responses to each stimulus. Estimating information from data can be difficult and should be undertaken carefully. The problems that are associated with estimating information from small data sets are beyond the scope of this discussion, but are discussed extensively elsewhere (Kjaer et al., 1994; Golomb et al., 1997; Panzeri, Treves, 1996).

For the work below, our use of information theory is largely comparative. We would like to know which representations of neural responses

allow the classification of stimuli with the greatest certainty, namely, which representations carry more information. Often we want to know whether adding an element to the representation of the response (that is, increasing the dimension of the code) adds information.

RESULTS

Response Strength

Since the earliest single-neuron recordings, it has been clear that the number of impulses or response strength is easily modulated by changing experimental parameters. For example, in V1 of monkey cortex, the responses to a stationary sine wave grating or bar centered on the receptive field change as the orientation is changed, giving rise to the classic tuning curve.

Despite this obvious relation between orientation and response strength, the interpretation of the responses is more complicated than first appears because the number of impulses varies widely across repeated trials. This variability is usually dealt with by averaging across (hopefully many) repetitions of the stimulus condition. Because the brain does not have the luxury of the experimentalist (observing many responses) and must decode the response from a single stimulus presentation, the brain cannot average across trials. Instead, the response must be decoded across neurons. Asking how many independent neurons are needed to decode the response to determine, for example, the orientation of a bar or grating, is natural. This question can be answered by examining how much information is needed to determine the orientation, and how much information is provided by a single neuron.

Past methods for calculating information have estimated response distributions directly from the data, although, as was pointed out above, doing so can be difficult. In the past few years, reliable methods have been developed (Kjaer et al., 1994; Golomb et al., 1997; Panzeri, Treves, 1996; Victor, Purpura, 1996). Recently we have taken another approach. Rather than estimating such

distributions directly from the data, we modeled the data using distributions with well-studied properties. For our spike counts, Gaussian distributions truncated at zero fit the data well enough to make accurate information estimates, whereas Poisson distributions fit the data substantially less well and lead to poor estimates of transmitted information (Gershon et al., 1998; Wiener & Richmond, 1998). Using the correct distribution is especially important for interpreting the origin of exactly timed spike patterns (see the section below on precise temporal patterns).

The means and variances of stimulus-elicited responses have been shown to be related (Fig. 1) (Tolhurst et al., 1981a; Tolhurst et al., 1981b; Tolhurst et al., 1983; van Kan et al., 1985; Vogels et al., 1989; Gershon et al., 1998). Gershon et al (1998) and Wiener and Richmond (1998) showed

that using this relation between mean and variance, and the truncated Gaussian model of response distributions, allows accurate calculation of the stimulus-related information that is carried in the spike count. In general, the amount of information carried in the spike counts of visual system neurons is about 0.4 to 0.5 bits. Thus, two or three independent neurons would be needed to provide the one bit of information that is necessary for dividing a stimulus set in two, and between 6 and 9 neurons to decode which of 8 (3 bits) oriented gratings appeared.

Rate Variation and Latency:

The firing rate generally changes during the response to a single stimulus presentation. Finding responses with the same average firing




Fig. 1: Relation between mean and variance across all presented stimuli for a single V1 complex cell. Each point shows the mean and variance of spike counts elicited by a single stimulus. The $\log(\text{variance})$ is linearly related to the $\log(\text{mean})$. The regression line (solid) has a slope of 1.16 and an intercept of 0.5. The data would not be well-approximated by a Poisson process (dotted line). The median variance to mean ratio for these data was 2.3 (interquartile range: 1.7-3.1). The variance to mean ratio for a Poisson distribution would be 1.0.

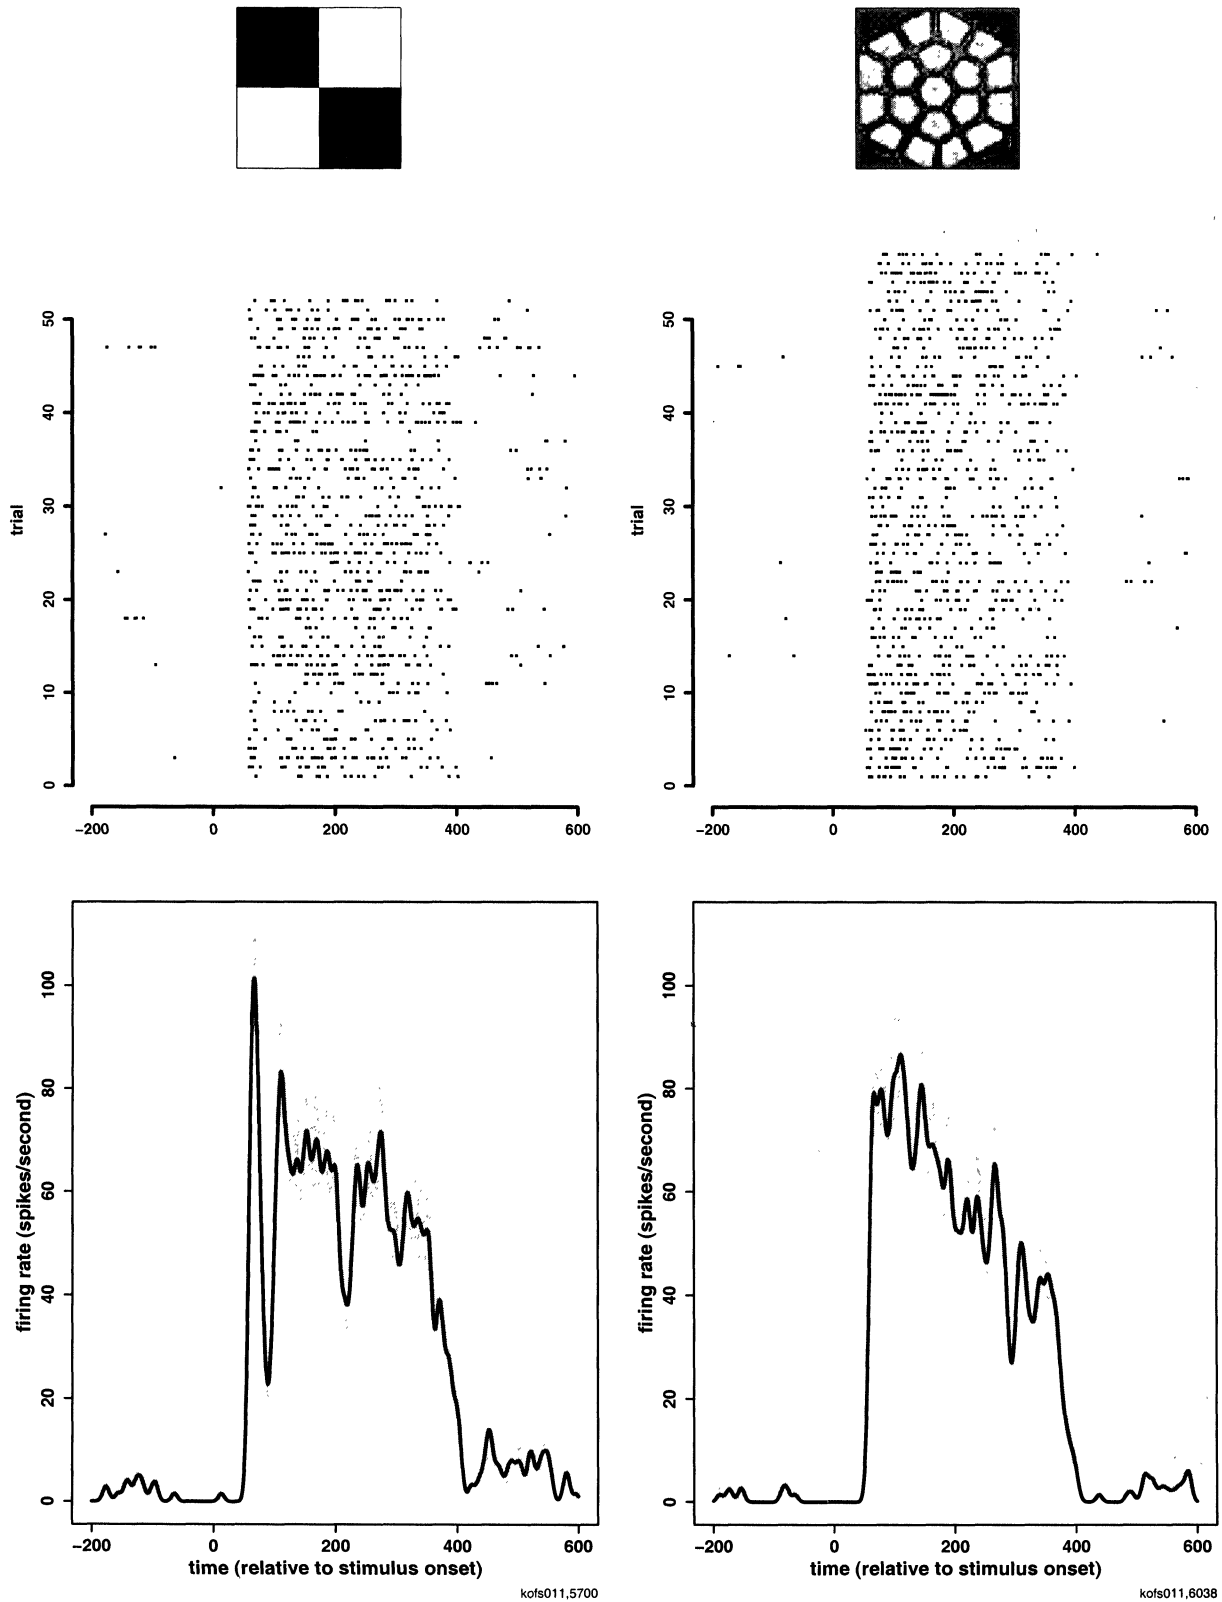


Fig. 2: Responses elicited by two stimuli from a V1 supragranular complex cell that are equal in response strength and different in temporal pattern. Two temporal features can be used to differentiate these responses. The first is the presence of the ~ 25 ms gap in the early response on the left and not on the right. The second is the progressively decreasing strength seen prominently in the response on the right.

rate, but different patterns, over time is not difficult (Fig. 2). Smoothing the responses over tens of milliseconds (low-pass filtering below 30 Hz) preserves all the stimulus-related information that is available in the unsmoothed responses (Heller et al., 1995). Including this slow variation in the response accounts for about 25% more information than that in the response strength alone (Optican, Richmond, 1987; Richmond et al., 1990; Tovee et al., 1993; Heller et al., 1995; Victor, Purpura, 1996).

The latency of a response, that is, the delay with which a change in the stimulus elicits a change in firing rate, is considered a particularly important feature of rate modulation. Gawne et al (1996) showed that the latency is strongly related to the contrast or luminance of the stimulus. Recently, we confirmed this result with other stimuli, including gratings (Fig. 3) (Wiener et al.,

1998). The minimum latency in V1 is related to the contrast across all response strengths.

The minimal description of neuronal responses must include some representation of the rate modulation. Whether the rate modulation must be represented at a precision greater than the tens of milliseconds described above continues to be a subject of study.

Exactly Timed Spike Patterns

So far we have seen that both the spike count (including its distribution) and the slowly varying pattern of rate change (including latency) are related to the stimulus. Each carries information that is unavailable from the other. It is also natural to represent a neuronal response as a series of discrete spike arrival times.

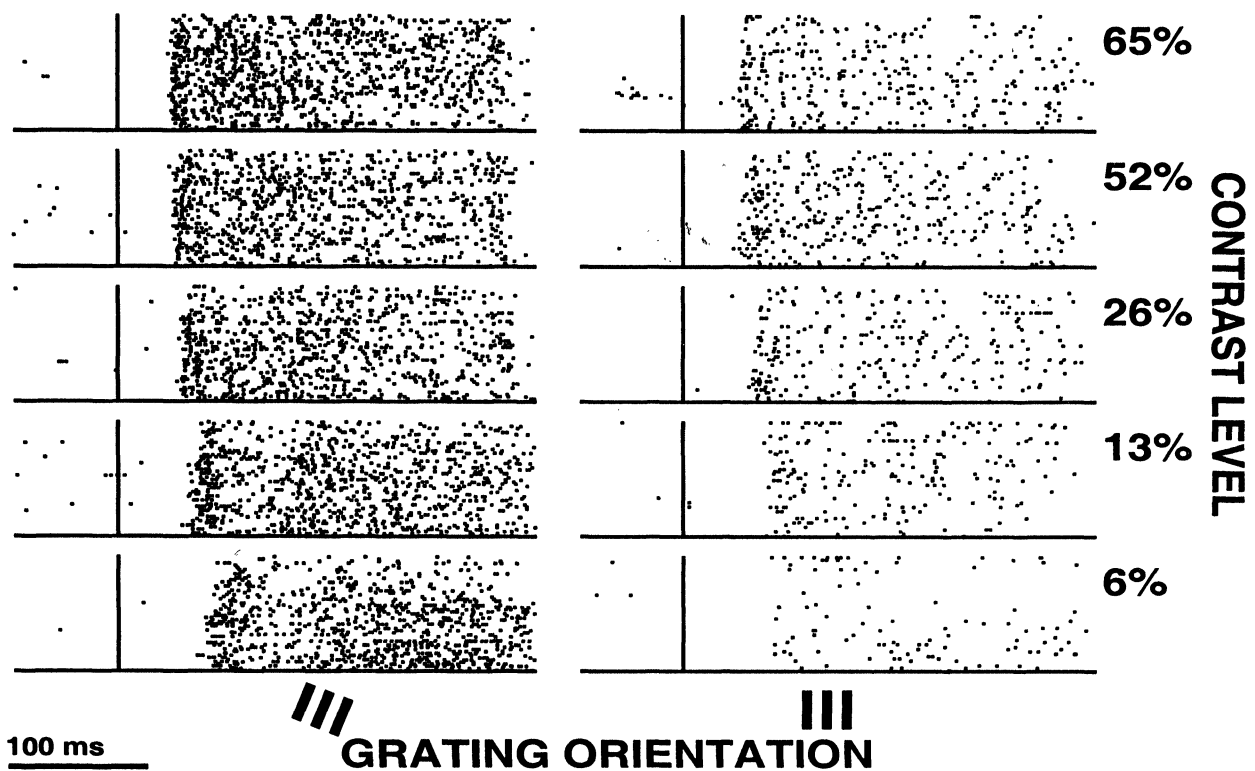


Fig. 3: Responses elicited by gratings of several contrasts at optimal and nonoptimal orientations from a V1 complex cell. There are substantial changes in response strength induced by changes in orientation, whereas the latency changes by, at most, a small amount. The latencies become substantially longer (~30 ms in this example) as the contrast decreases. The minimum latency at each contrast is the same at both orientations. The response elicited by the lowest contrast optimally oriented grating is large with a long latency (bottom left column) whereas the response elicited by another grating is weak and has a short latency (top right column). Thus, latency and response strength are independent.

If the time of every spike is important, the number of distinct response patterns is substantially larger than the number of spikes, theoretically making it possible for response patterns to encode much more information than that in spike count alone. In the simplest model, there would be $n!/k!(n-k)!$, that is, n choose k , different ways of placing k spikes in n time bins. The large number of the potentially available degrees of freedom has led both experimentalists and theoreticians to consider whether spike patterns measured with millisecond precision carry information that has not yet been identified. Detecting whether and which patterns carry information is difficult, however, if n choose k really indicates the number of degrees of freedom in a neuronal response. Although exact timing provides an extraordinary number of distinct patterns (more than 1,000,000 for 3 spikes in 200 ms), only if particular patterns are consistently present in responses elicited by particular stimuli do they transmit information.

In a provocative speculation, a proposal has been made that particular types of spike patterns across neurons might play a critical role in higher brain functions, such as the perception of objects (Abeles, 1991; Lestienne, Strehler, 1987; Lestienne, Tuckwell, 1998; von der Malsburg, Schneider, 1986). To investigate implications of this speculation, Oram et al (1999) recently carried out the information theoretical and statistical studies of the precisely timed spike patterns described below. The investigators used three types of previously studied exactly timed spike patterns (Lestienne, Strehler, 1987; Abeles et al., 1993; Abeles, Gerstein, 1988):

1. triplets for which the exact same pattern repeated one or more times (called repeating triplets below),
2. repeating quadruplets (defined in the same manner as the repeating triplets, except the patterns used four spike times), and
3. the numbers of each type of triplets occurring across all presentations of a particular stimulus.

Here we discuss the results from the repeating triplets, which are representative of those from all three types of exactly-timed spike patterns

examined.

First, the information carried in the spike count alone was compared with the information carried by the spike count plus the numbers of repeating triplets. Despite the triplets carrying some stimulus-related information, the joint code carried only the same amount of information as that carried by the spike count alone.

The information theoretical analysis shows that the number of triplets is strongly related to the spike count. Therefore, a model was sought to connect the spike count to the exactly-timed spike patterns. Several such models have been proposed. These models share the assumption that the exact times of spikes are randomly determined, but differ in the distribution from which spike times are drawn. The most common class of models uses a Poisson process (originally uniform, time-varying more recently) to determine the spike times. The time-varying Poisson process maintains the appropriate peristimulus time histogram (PSTH), but fails to match the distribution of spike counts and interspike intervals. Other models match the distribution of spike counts. Randomly reordering the interspike intervals within a train maintains the interval distribution but not the PSTH; exchanging spikes between trains maintains the PSTH, but not the interval distribution; and jittering the times of spikes in trains retains both parameters approximately but neither exactly. Thus, each model preserves some important features of the response, but none preserves all (Table 1).

The numbers of precisely timed patterns found in the spike trains that were generated using these models have failed to match the numbers in experimentally observed spike trains. This result has led researchers to reject their models and to tentatively conclude that at least some precisely timed spike patterns are determined by the stimulus condition (Abeles et al., 1993; Aertsen et al., 1991; Zhang et al., 1997; Lestienne, Strehler, 1987; Lestienne, Tuckwell, 1998; Prut et al., 1998). It is possible, however, to retain the assumption of stochasticity and to conclude instead that the models of spike time distributions are not adequate.

TABLE 1

Models for connecting the spike count to the exactly-timed spike patterns

Model	Method	Interval distribution	PSTH matched	Spike count variance	Match precise timing
Homogeneous Poisson	Uniform Poisson process	—	—	—	—
Interspike interval matched	Shuffle interspike intervals	+	—	+	—
PSTH matched Poisson (non homogeneous)	Time varying Poisson process	—	+	—	—
Spike count matched	Same as above with trials selected to match spike count	+	+	+	+

An Accurate Model of Spike Timing

To account for the results of the information theoretical analyses summarized above, Oram et al (1999) created and tested a new statistical model, the spike count matched model described below. The spike count matched model generates single spike trains in a manner similar to the nonhomogeneous Poisson process that was used by Lestienne et al (1986). Instead of assuming that the process generating spikes was Poisson, however, which would lead to a Poisson spike count distribution, the spike counts were forced to match those observed in the experimental data because, as pointed out above, the Poisson distribution fits the data poorly. If, for example, the data had six spikes in a trial, an artificial train with six spikes was generated (Fig. 4).

Recently Berry and Meister (1998) have shown that the interspike interval distribution affects how well the response can be modeled. Therefore, in the spike count matched and the nonhomogeneous Poisson models, the interspike interval distribution, generated initially by the model, was forced to match that from the data by adjusting the probabilities of the first two intervals. When generated creating a 1 or 2 ms interval, a spike was discarded randomly so that the resulting interval distribution matched the distribution in the data (see Oram et al, 1999). This adjustment of the interspike intervals had a

small but significant effect on the numbers of repeating triplets. For example, in V1 the number of repeating triplets averaged 0.57; the predictions were 0.45 for the NHPP model, 0.55 for the spike count matched model, and 0.57 for the spike count matched model with the ISI's adjusted. For the data presented here, only the first two intervals must be adjusted.

Figure 5A shows that the homogeneous Poisson model, the interval shuffling model, and the nonhomogeneous Poisson model all underestimate the number of triplets found in the data. In contrast, the spike count matched model predicts numbers of triplets that are indistinguishable from the numbers found in the data. Because the number of triplets is a stochastic consequence of the spike count and interval distributions and the PSTH, it can carry only information already available from those measures.

Investigating whether particular types of triplets, namely, those defined by particular pairs of intervals, carry additional stimulus-related information is still necessary. Because the number of triplets being counted is large (here the 625 different triplet types with both intervals ≤ 25 ms), a problem arises with multiple comparisons. Even if the spike count matched model is correct, $0.05 \times 625 = 31.25$ tests are expected to be significant at the $p < 0.05$ level. In 10,000 simulated runs of the spike count matched model,




Fig. 4: Spike count matched model. Spike trains are generated so that there is one simulated spike train to match each train in the data, with the spike count in the simulated train being forced to be equal to the corresponding train in the data. The top panel shows the average spike density plot for the responses of one V1 supragranular complex cell aligned at the time of stimulus onset. This is integrated over the whole period to give the curve in the lower panel. This integrated spike probability density function is then used to generate simulated spike trains using a uniform random number generator. The number of spikes needed is chosen from the number of spikes in the recorded data. In this example, six spikes were needed, so six random numbers would be chosen, placed on the ordinate and mapped through the cumulative spike density function as shown to generate the simulated spike train indicated by the dots at the bottom. For illustration here, six numbers separated by equal intervals were used. After transformation through the cumulative probability function the intervals are no longer equal. Using this procedure places the spikes stochastically due to the random numbers. However, the probability density function averaged across many examples will be indistinguishable from that in the upper panel.

each containing the same number of trials as our recorded data, we found, on average, exactly the abovementioned number of significant tests. This result is consistent with the hypothesis that triplets of each type arise by chance.

Figure 6 illustrates the dangers of multiple comparisons. The top panel shows the number of triplets of each type in the experimental data. The four panels below show the number of triplets of each type in four different sets of data that were simulated using the spike count matched model. Three model runs have peaks that are as large as the largest peak found in the data. Thus, although to think that the large peak from the data must be

significant is tempting, we must exercise caution because a stochastic model leads to equally large peaks. If we accept the high peaks in the data as significant, we must also accept the high peaks in the simulations as significant, yet we know that the latter were generated by a stochastic process that is directly related to the spike count.

The results show that matching the spike count distribution is critical for matching the numbers of precisely-timed patterns in the data. The reason for this can be seen by examining the relation between the firing rate and the number of repeating triplets. When the firing rate is high, the number of triplets increases very rapidly (Fig. 5C).

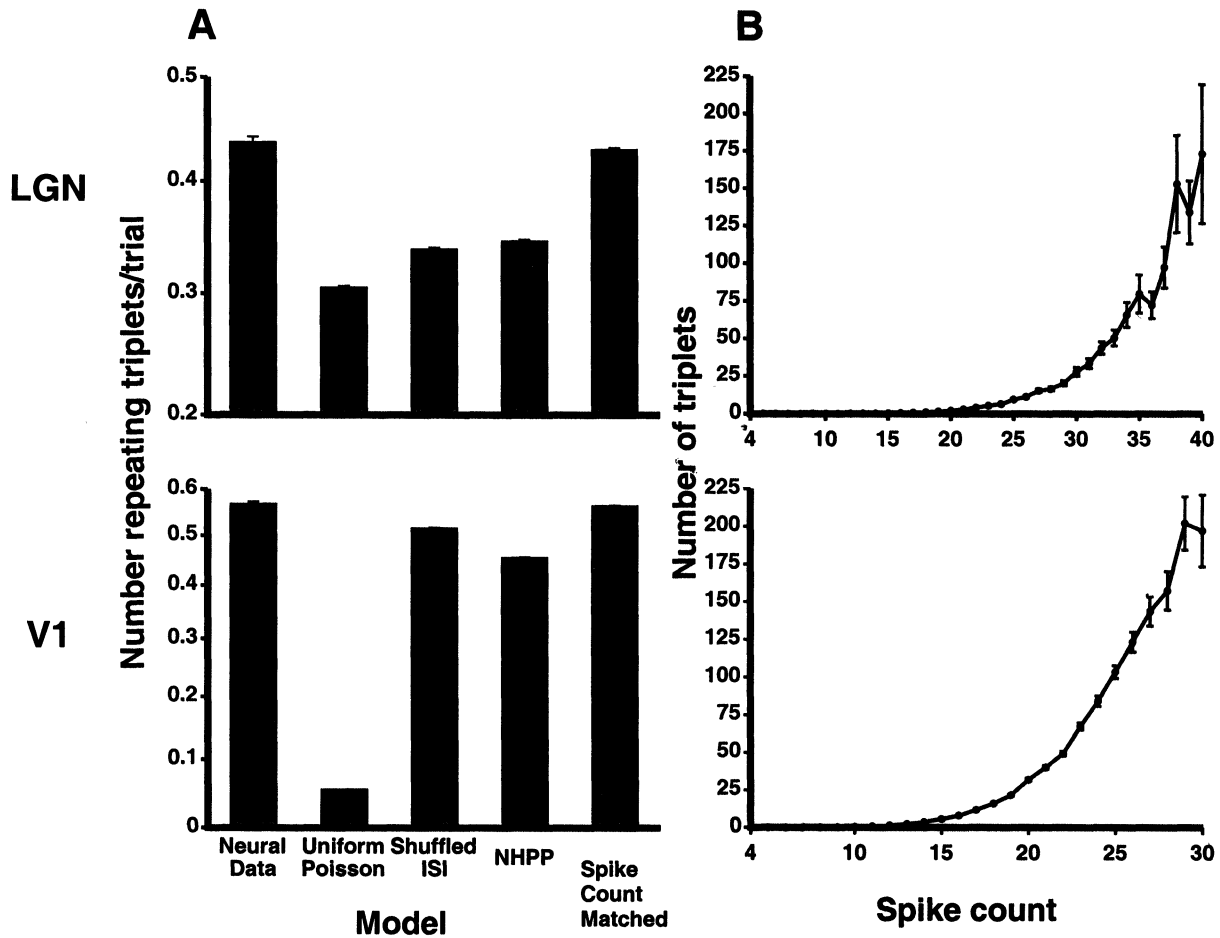
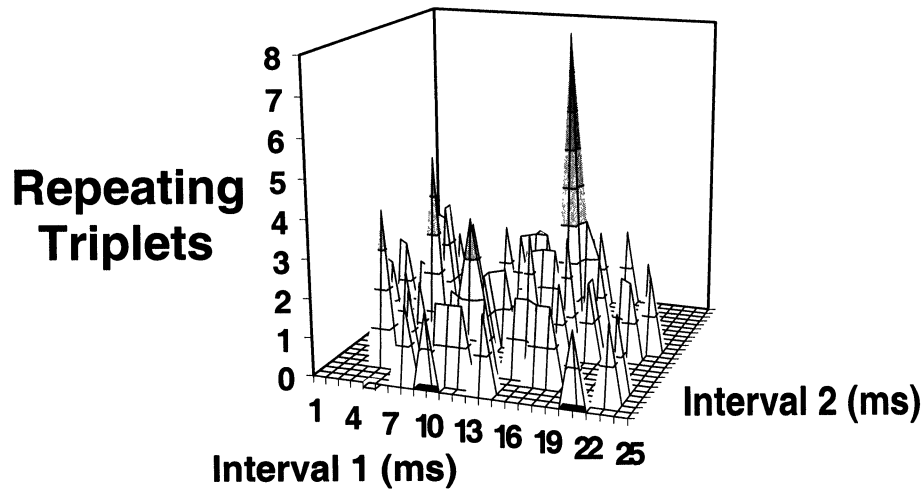


Fig. 5: A. The number of repeating triplets seen in the data (left bar) and four models: the uniform Poisson model, the interval shuffling model, the nonhomogeneous Poisson model, and the spike count matched model. The numbers of repeating triplets generated by the spike count matched model are indistinguishable from the number found in the data. The numbers generated by the other 3 models are significantly ($p < 0.001$) different than those found in the data. B. Relation of the number of triplets to spike count. The numbers of triplets increases nonlinearly at all spike counts. If the variability of the responses is underestimated the number of triplets predicted will be underestimated because the number of triplets added by a high spike count will not be completely compensated by the number lost at the corresponding low spike count.

PRIMARY VISUAL CORTEX



SPIKE COUNT MATCHED MODEL

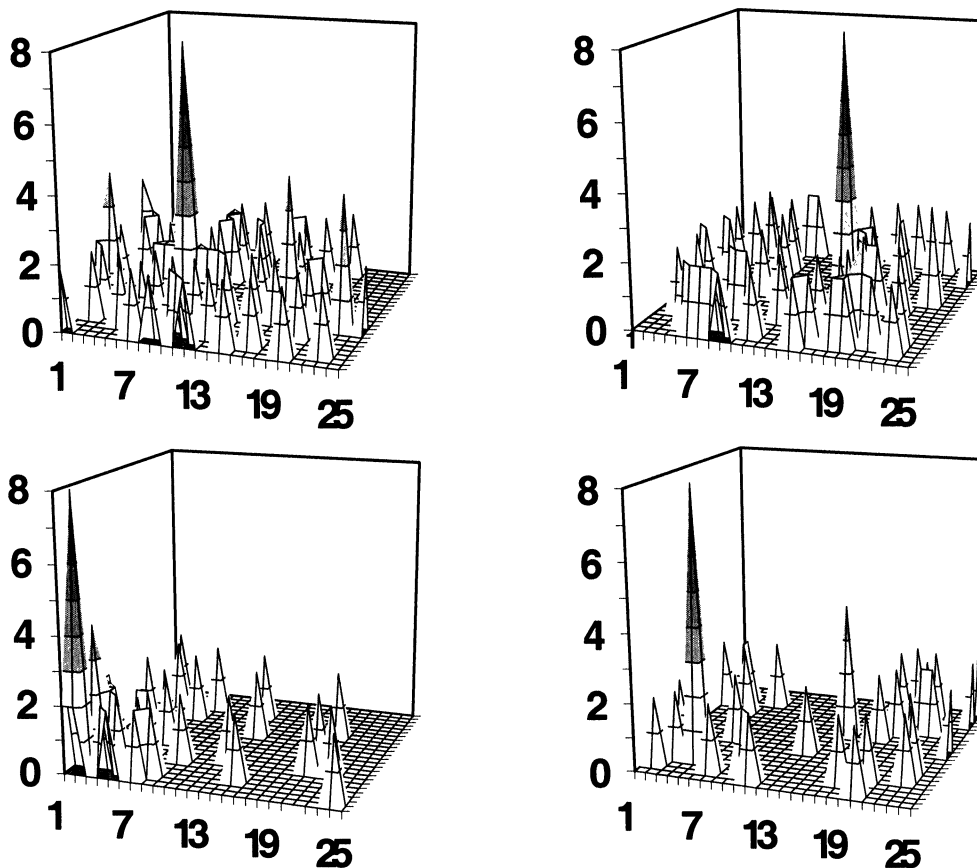


Fig. 6: Numbers of repeating triplets of all types with each interval less than or equal to 25 ms. The top panel shows the numbers of repeating triplets of each type for one stimulus in a V1 complex cell. The other 4 panels show 4 different simulation runs. These figures indicate the difficulty with multiple comparisons. Although the triplets occurred most often are different in the data and simulations, they are different across the different simulations, indicating that the numbers of each triplet type, including the ones occurring most frequently, can arise from a stochastic model.

If a few trials with large numbers of spikes are overlooked (as would happen if, for example, the assumed variance was smaller than the true variance), then the number of triplets predicted will be too low. The results are not as sensitive to changes in the PSTH. Low-pass filtering the PSTH does not affect the predicted triplet structure as long as frequencies below about 30 Hz are not attenuated. This corresponds to smoothing the PSTH with a Gaussian having a standard deviation of ~ 5 ms, which is equivalent to grouping spikes in 30 ms wide time bins.

DISCUSSION

A spike train has many degrees of freedom, and the idea that many of them might carry information is appealing. However, the results reviewed above show that accurately describing spike trains requires specifying only (a) the spike count distribution (which is approximately truncated Gaussian), (b) the PSTH (with a bandwidth of less than 30 Hz, the equivalent of measuring spike counts in 30 ms wide bins), and (c) the interval histogram. If they completely describe single neuronal responses, then these features contain all of the information that is available from those responses, no matter what representation is chosen. The accuracy of the spike count matched model shows that the fine temporal structure of responses (here triplets and quadruplets repeating within and across trials) depends stochastically only on these coarser measures. Thus, the fine temporal structures should be viewed as a transformed representation of the response.

Given these findings, one can wonder whether these precisely timed structures could be used to enhance information transmission. That possibility seems unlikely to us. The information carried by these patterns is already available from the spike count (and from the slowly varying temporal variation seen in the PSTH). Furthermore, the amount of information carried by the triplets is much less than that carried by the spike count, because per trial, the number of triplets is much more variable than the number of

spikes. There are large numbers of trials that have few or no repeating triplets and a few or a single trial with a large number of repeating triplets, sometimes hundreds. This variability is even greater for any particular triplet type. Such a large variance makes these triplet patterns a code that is less reliable than spike count. Thus, the triplets (and other precisely-timed patterns) are poor candidates for carrying critical information.

This work clearly demonstrates that the fine temporal structure of a spike train is sensitive to the distribution of spike counts (Fig 5C). Poisson distributions have often been assumed to be effective models of the response. If we consider sufficiently narrow time bins and assume that spikes appear in each independently of all the others, a Poisson process arises naturally. This appealing derivation and the mathematical tractability of the Poisson process have contributed to the widespread use of Poisson models of spike trains. These models predict a Poisson distribution of spike counts in sufficiently long time windows. Experimental data, however, including that presented here, have generally been inconsistent with the Poisson hypothesis (Baddeley et al., 1997; Berry, Meister, 1998; Bradley et al., 1987; Britten et al., 1993; Buracas et al., 1998; Gershon et al., 1998; Lee et al., 1998; Levine, Troy, 1986; Reich et al., 1997; Snowden et al., 1992; Tolhurst et al., 1983; Victor, Purpura, 1996; Vogels et al., 1989). A correlation between time bins, such as that observed in Heller et al (1995), shows that the Poisson assumption cannot be correct.

We conclude that the spike trains are consistent with a stochastic process generating all the spikes, and that serial correlation on a broad time scale (spike count distribution and PSTH) can give rise to the fine temporal structures seen in the data. In another context, Brody (1998) has shown that slow correlations in spike counts between pairs of neurons can give rise to narrow cross-correlogram peaks (a type of precise correlation). Thus, the existence of structure at fine time scales does not imply control at fine time scales. Accounting for the influence of correlations over long periods on precisely timed patterns of any type found in spike trains is always

necessary. When the fine temporal structure is predicted from the coarse temporal structure, it can carry no unique information and then, for assessing the information carried, only the coarse structure need be described. The results may be useful for interpreting the significance of precisely timed patterns of spikes across neurons for information processing.

REFERENCES

- Abeles M. *Corticonics*. Cambridge, England: Cambridge University Press, 1991.
- Abeles M, Bergman H, Margalit E, Vaadia E. Spatio-temporal firing patterns in the frontal cortex of behaving monkeys. *J Neurophysiol* 1993; 70: 1629–1638.
- Abeles M, Gerstein GL. Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J Neurophysiol* 1988; 60: 909–924.
- Aertsen A, Vaadia E, Abeles M, Ahissar E, Bergman H, Karmon B, Lavner Y, Margalit E, Nelken I, Rotter S. Neural interactions in the frontal cortex of a behaving monkey: signs of dependence on stimulus context and behavioral state. *J Hirnforsch* 1991; 32: 735–743.
- Baddeley R, Abbott LF, Booth MC, Sengpiel F, Freeman T, Wakeman EA, Rolls ET. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc R Soc Lond B Biol Sci* 1997; 264: 1775–1783.
- Berry MJ, Meister M. Refractoriness and neural precision. *J Neurosci* 1998; 18: 2200–2211.
- Bradley A, Skottun BC, Ohzawa I, Sclar G, Freeman RD. Visual orientation and spatial frequency discrimination: a comparison of single neurons and behavior. *J Neurophysiol* 1987; 57: 755–772.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA. Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 1993; 10: 1157–1169.
- Brody CD. Slow covariations in neuronal resting potentials can lead to artefactually fast cross-correlations in their spike trains. *J Neurophysiol* 1993; 80: 3345–3351.
- Buracas GT, Zador AM, DeWeese MR, Albright TD. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 1998; 20: 959–969.
- Cover TM, Thomas JA. *Elements of Information Theory*. New York: John Wiley & Sons, 1991.
- Gawne TJ, Kjaer TW, Richmond BJ. Latency: another potential code for feature binding in striate cortex. *J Neurophysiol* 1996; 76: 1356–1360.
- Gershon ED, Wiener MC, Latham PE, Richmond BJ. Coding strategies in monkey V1 and inferior temporal cortices. *J Neurophysiol* 1998; 79: 1135–1144.
- Golomb D, Hertz J, Panzeri S, Treves A, Richmond B. How well can we estimate the information carried in neuronal responses from limited samples? *Neural Comput* 1997; 9: 649–665.
- Heller J, Hertz JA, Kjaer TW, Richmond BJ. Information flow and temporal coding in primate pattern vision. *J Comput Neurosci* 1995; 2: 175–193.
- Kjaer TW, Hertz JA, Richmond BJ. Decoding cortical neuronal signals: network models, information estimation and spatial tuning. *J Comput Neurosci* 1994; 1: 109–139.
- Lee D, Port NL, Kruse W, Georgopoulos AP. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci* 1998; 18: 1161–1170.
- Lestienne R, Strehler BL. Time structure and stimulus dependence of precisely replicating patterns present in monkey cortical neuronal spike trains. *Brain Res* 1987; 437: 214–238.
- Lestienne R, Tuckwell HC. The significance of precisely replicating patterns in mammalian CNS spike trains. *Neuroscience* 1998; 82: 315–336.
- Lestienne R, Tuckwell H, Chalansonnet M, Chaput M. Repeating patterns and oscillations in mitral cell discharges of freely breathing rats. *Eur J Neurosci* (in press).
- Levine MW, Troy JB. The variability of the maintained discharge of cat dorsal lateral geniculate cells. *J Physiol (Lond)* 1986; 375: 339–359.
- Optican LM, Richmond BJ. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J Neurophysiol* 1987; 57: 162–178.
- Oram MW, Wiener MC, Lestienne R, Richmond BJ. Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J Neurophysiol* 1999; 81: 3021–3033.
- Panzeri S, Treves A. Analytical estimates of limited sampling biases in different information measures. *Network* 1996; 7: 1–21.
- Prut Y, Vaadia E, Bergman H, Haalman I, Slovin H, Abeles M. Spatiotemporal structure of cortical activity: properties and behavioral relevance. *J Neurophysiol* 1998; 79: 2857–2874.
- Reich DS, Victor JD, Knight BW, Ozaki T, Kaplan E. Response variability and timing precision of neuronal spike trains in vivo. *J Neurophysiol* 1997; 77: 2836–2841.
- Richmond BJ, Optican LM, Spitzer H. Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. I. Stimulus-response relations. *J Neurophysiol* 1990; 64: 351–369.

- Shannon CE, Weaver WW. The mathematical theory of communication. Urbana, IL; University of Illinois Press, 1949.
- Snowden RJ, Treue S, Andersen RA. The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Exp Brain Res* 1992; 88: 389–400.
- Tolhurst DJ, Dean AF, Thompson ID. Preferred direction of movement as an element in the organization of cat visual cortex. *Exp Brain Res* 1981a; 44: 340–342.
- Tolhurst DJ, Movshon JA, Thompson ID. The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Exp Brain Res* 1981b; 41: 414–419.
- Tolhurst DJ, Movshon JA, Dean AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 1983; 23: 775–785.
- Tovee MJ, Rolls ET, Treves A, Bellis RP. Information encoding and the responses of single neurons in the primate temporal visual cortex. *J Neurophysiol* 1993; 70: 640–654.
- van Kan PL, Scobey RP, Gabor AJ. Response covariance in cat visual cortex. *Exp Brain Res* 1985; 60: 559–63.
- Victor JD, Purpura KP. Nature and precision of temporal coding in visual cortex: a metric- space analysis. *J Neurophysiol* 1996; 76: 1310–1326.
- Vogels R, Spileers W, Orban GA. The response variability of striate cortical neurons in the behaving monkey. *Exp Brain Res* 1989; 77: 432–436.
- von der Malsburg C, Schneider W. A neural cocktail-party processor. *Biol Cybern* 1986; 54: 29–40.
- Wiener MC, Oram MW, Richmond BJ. Response latency is related to stimulus contrast, but not to response strength. *Society for Neuroscience Abstracts* 1998; 24: 1258 (Abstract).
- Wiener MC, Richmond BJ. Using response models to study coding strategies in monkey visual cortex. *Biosystems* 1998; 48: 279–286.
- Zhang J, Riehle A, Requin J, Kornblum S. Dynamics of single neuron activity in monkey primary motor cortex related to sensorimotor transformation. *J Neurosci* 1997; 17: 2227–2246.