

Power Spectrum Analysis of Bursting Cells in Area MT in the Behaving Monkey

Wyeth Bair,¹ Christof Koch,¹ William Newsome,² and Kenneth Britten²

¹Computation and Neural Systems Program, California Institute of Technology, Pasadena, California 91125 and

²Department of Neurobiology, Stanford University School of Medicine, Stanford, California 94305

It is widely held that visual cortical neurons encode information primarily in their mean firing rates. Some proposals, however, emphasize the information potentially available in the temporal structure of spike trains (Optican and Richmond, 1987; Bialek et al., 1991), in particular with respect to stimulus-related synchronized oscillations in the 30–70 Hz range (Eckhorn et al., 1988; Gray et al., 1989; Kreiter and Singer, 1992) as well as via bursting cells (Cattaneo et al., 1981a; Bonds, 1992). We investigate the temporal fine structure of spike trains recorded in extrastriate area MT of the trained macaque monkey, a region that plays a major role in processing motion information. The data were recorded while the monkey performed a near-threshold direction discrimination task so that both physiological and psychophysical data could be obtained on the same set of trials (Britten et al., 1992). We identify bursting cells and quantify their properties, in particular in relation to the behavior of the animal.

We compute the power spectrum and the distribution of interspike intervals (ISIs) associated with individual spike trains from 212 cells, averaging these quantities across similar trials. (1) About 33% of the cells have a relatively flat power spectrum with a dip at low temporal frequencies. We analytically derive the power spectrum of a Poisson process with refractory period and show that it matches the observed spectrum of these cells. (2) About 62% of the cells have a peak in the 20–60 Hz frequency band. In about 10% of all cells, this peak is at least twice the height of its base. The presence of such a peak *strongly* correlates with a tendency of the cell to respond in bursts, that is, two to four spikes within 2–8 msec. For 93% of cells, the shape of the power spectrum did not change dramatically with stimulus conditions. (3) Both the ISI distribution and the power spectrum of the vast majority of bursting cells are compatible with the notion that these cells fire Poisson-distributed bursts, with a burst-related refractory period. Thus, for our stimulus con-

ditions, no explicitly oscillating neuronal process is required to yield a peak in the power spectrum. (4) We found no statistically significant relationship between the peak in the power spectrum and psychophysical measures of the monkeys' performance on the direction discrimination task. (5) For cells firing bursts, ROC (receiver operating characteristic) analysis shows that the "event" rate, where an event is either a single burst of spikes or an isolated spike, is on average a more sensitive measure of visual stimulus direction than the total number of spikes, used previously (Britten et al., 1992), implying that the number of spikes in a burst is less stimulus dependent than the overall firing rate or the rate of bursts.

[Key words: behaving monkey, extrastriate cortex (MT), oscillations, power spectrum, interspike interval analysis, bursting]

What neural code is used by the brain to decipher sensory events and translate them into a percept of the visual scene? Because the mean firing frequency in response to a sensory stimulus is reproducible under identical stimulus conditions and varies predictably and smoothly with such stimulus parameters as velocity, contrast, orientation, and so on, it is widely held to be the primary variable relating neuronal response to sensory experience (Lettvin et al., 1959, or the 5. dogma in Barlow, 1972). This belief is supported by the existence of a quantitative relationship between the firing rates of single cortical neurons and psychophysical judgements made by behaving monkeys (Werner and Mountcastle, 1963; Barlow et al., 1987; Newsome et al., 1989a; Vogels and Orban, 1990; Zohary et al., 1990; Britten et al., 1992). Some electrophysiologists have focused on the idea that the detailed dynamics of the neuronal response may carry significant information (e.g., Poggio and Viernstein, 1964; Chung et al., 1970; Strehler and Lestienne, 1986; Optican and Richmond, 1987; Abeles, 1990; Eskandar et al., 1992; Zipser et al., 1993; see also Bialek et al., 1991). A great deal of attention has recently been given to the reports of stimulus-induced semisynchronous neuronal oscillations in the 30–70 Hz range in the visual cortex of the anesthetized cat (Eckhorn et al., 1988; Gray and Singer, 1989; Gray et al., 1989; Ghose and Freeman, 1992) and the awake monkey (Kreiter and Singer, 1992). Moreover, Gray et al. (1989) report that oscillating neurons up to 10 mm apart can be phase-locked with a phase-shift close to zero; that is, these neurons usually fire within ± 3 msec of each other (for a review, see Singer, 1994).

Only a few reports have focused on the possible significance of bursting for neuronal coding. Cattaneo et al. (1981a,b) report that complex (but *not* simple) cells in area 17 of anesthetized

Received June 21, 1993; revised Oct. 20, 1993; accepted Oct. 26, 1993.

We thank a number of individuals with whom we discussed these findings over the last two years, in particular Francis Crick, Joel Franklin, Ehud Zohary, and William Softky. This work was supported by the Office of Naval Research, the Air Force Office of Scientific Research, the National Science Foundation, the National Eye Institute, and the McDonnell-Pew Program in Cognitive Neuroscience. W.B. received support from a National Science Foundation Graduate Fellowship and is now supported by the L. A. Hanson Foundation. K.B. received support from an NIMH training grant to the Department of Neurobiology at Stanford University.

Correspondence should be addressed to Wyeth Bair, Computation and Neural Systems Program, 139-74, Division of Biology, Pasadena, CA 91125.

Copyright © 1994 Society for Neuroscience 0270-6474/94/142870-23\$05.00/0

and alert cats frequently respond with bursts. They show that the frequency of bursts (or "grouped spikes") varies strongly with the spatial frequency and orientation of sinusoidal drifting gratings, while the frequency of "isolated spikes" only weakly depends on these parameters, encoding rather the contrast of the stimulus. Bonds (1992) found in his analysis of cat striate neurons that the structure of spike trains—in his case the frequency and duration of bursts—can vary substantially on the basis of how the firing rate was generated. For instance, presentation of stimuli at nonoptimal orientations at high contrasts yields bursts that are shorter than those generated by lower-contrast stimuli at optimal orientations. Legéndy and Salzman (1985) hypothesized functional significance for burst-firing patterns in spontaneously active striate neurons of alert cats, but by their definition, bursts included mostly long periods (0.5–2.0 sec) of significant elevation in firing rate. Finally, Crick (1984) postulated that the neuronal expression of selective visual attention is the production of bursting in a subset of thalamic neurons. This bursting, in combination with a short-term and transient alteration in the synaptic strength (as proposed by von der Malsburg, 1981), could lead to the short-term formation of transient cell assemblies at the level of cortex (see also Crick and Koch, 1990, 1992).

Motivated by these findings and proposals, we investigated the temporal properties of cortical cells in the awake and behaving monkey. The data were obtained from an ongoing series of experiments linking the responses of neurons in extrastriate area MT (or V5) to the psychophysical performance of trained monkeys (Newsome et al., 1989a,b; Britten et al., 1992). MT is specialized for the analysis of visual motion, as indicated by its preponderance of directionally selective neurons (Zeki, 1974; Maunsell and Van Essen, 1983) and by the motion-specific effects of lesions and electrical microstimulation in MT (Newsome and Pare, 1988; Salzman et al., 1992). The monkeys were trained to discriminate the direction of motion in a stochastic visual display while the responses of MT neurons were simultaneously recorded. As reported previously, the responses of individual MT neurons, considered simply to be the total number of spikes occurring during the period of visual stimulation, can provide a remarkably accurate account of a monkey's performance on the near-threshold direction discrimination task. Single-neuron responses, analyzed by a method based in signal detection theory, provide a good description of psychophysical threshold and of the shape of the psychometric function relating performance to the strength of the motion signal (Newsome et al., 1989a; Britten et al., 1992). In addition, small trial-to-trial variations in the strength of a neuron's response to identical, weak motion stimuli can be significantly correlated with the directional judgments made by the monkey on the same trials (Newsome et al., 1989b). In all prior analyses of these data, however, the temporal characteristics of the spike train were ignored.

We now analyze this single-cell database with an eye toward describing the temporal structure of MT responses and uncovering any relationship between the temporal structure and the psychophysical performance of the animal. We only consider data from well-isolated single neurons, leaving the analysis of other data, including well-isolated pairs of simultaneously recorded neurons (Zohary et al., 1992) and multiunit data, to a future report. We find that the temporal structure of MT responses is characteristic for each neuron and does not change in a stimulus-specific manner. Spikes are distributed almost randomly in time for some neurons, but are highly nonrandom

for other neurons, being characterized by occasional "bursts" in which spikes are tightly clustered in time. Both firing patterns are well described by simple models that require no intrinsically oscillatory process. A signal detection analysis indicates that bursting neurons convey more information about the direction of motion in the stimulus if bursts are considered to be individual signaling events.

Some of these results are briefly described elsewhere (Bair et al., 1992, 1993).

Methods

We first summarize the methods used to obtain the spike trains and then describe the data analysis techniques. Given the importance of applying the underlying mathematical transformation correctly, we justify our analysis in detail.

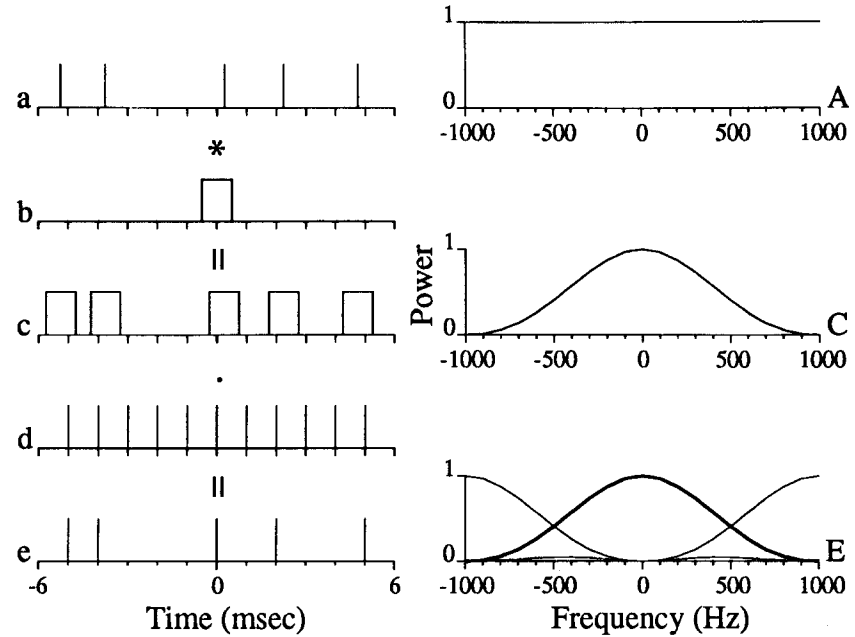
Data collection

Experimental methods for the collection of the original data are described in detail by Britten et al. (1992). Three adult macaque monkeys were trained to report the direction of motion in a dynamic random dot display in which a certain fraction, c , of the dots moved coherently at a common speed in one direction, while the remaining dots moved in random directions and at random speeds. For $c = 0$ all dots moved randomly, for $c = 1$ all dots moved in a common direction, the neuron's preferred direction, and for $c = -1$ all dots moved in the opposite direction. For a given block of trials, the random dot stimuli were placed within the receptive field of a single MT neuron, and the coherent motion signal was randomly presented in the preferred direction of the MT neuron ($c \geq 0$) or in the direction 180° opposite (null direction, $c \leq 0$) and the fraction of dots, $|c|$, carrying the coherent signal was varied randomly from trial to trial to cover a prespecified range of coherence values near psychophysical threshold. For a typical experiment, at least 210 trials were performed: 15 trials at each of six preferred and six null direction motion coherence levels plus 30 trials at $c = 0$ (random noise). Far more data were obtained for some experiments since additional blocks of trials were run as long as the cell remained well isolated.

An individual trial began with the onset of a fixation point presented on an oscilloscope 57 cm distant from the animal. After the monkey directed its gaze toward the fixation point, the random pattern appeared within the receptive field of the MT neuron for 2 sec. The monkey attended to the random dot display and judged the direction of the coherent motion signal while maintaining its gaze on the fixation point. At the end of the viewing interval, the fixation point and the random dot stimulus were extinguished, and two light-emitting diodes (LEDs) appeared corresponding to the two possible directions of the coherent motion signal. The monkey indicated its decision regarding the direction of motion by making a saccadic eye movement to the appropriate LED; correct choices were rewarded with water or juice. The monkey's eye movements were monitored continuously throughout the experiment using a scleral search coil system (Robinson, 1963).

Action potentials were recorded extracellularly from 216 MT neurons while monkeys performed the direction discrimination task. Thus, physiological data and psychophysical data were obtained on the same trials. On each trial the physiological data consisted of a spike train recorded continuously during the fixation interval, the 2 sec visual stimulation interval, and the intertrial interval. The time of occurrence of each action potential was recorded with a resolution of 1 msec. In the previous analysis (Britten et al., 1992), psychophysical data were compiled into psychometric functions relating percentage of correct choices to the strength of the motion signal. A psychophysical threshold, c_{system} , was measured for each experiment, where threshold was considered to be the motion coherence level that supported 82% correct performance. This psychophysical threshold characterizes the perceptual sensitivity of the monkey to the motion signals under the specific conditions of each individual experiment. Neuronal sensitivity was measured from the responses to preferred and null direction motion obtained over a range of coherence levels; responses were considered to be the total number of spikes that occurred during the 2 sec visual stimulation interval. Using a method based on signal detection theory, a "neuro-metric function" was computed that expressed the theoretical performance of an ideal observer who judges the direction of motion in the visual stimulus based only on the responses of the MT neuron being

Figure 1. Aliasing in the power spectrum $S(f)$ of a spike train. The ideal spike train (*a*) is a sequence of δ -functions at arbitrary time instants, but in recorded spike train data (*e*), spikes are assigned to sampling points. To go from *a* to *e* we convolve with a 1-msec-wide binning function (*b*) to yield *c* and then multiply by a 1 kHz sampling function (*d*) to yield (*e*). To the right are power spectra for the special case where *a* is random (Poisson). Ignoring the DC component, *A* is flat. Convolution with *b* multiplies the power spectrum by the square of the *sinc* function, yielding *C* as the spectrum of *c*. Sampling in this special case corresponds to adding up infinitely many copies of the *sinc*² at 1 kHz intervals (*E*), which yields a flat power spectrum (not shown) like *A*. The thick curve in *E* demonstrates the contribution to the final power spectrum from the true spectrum, while the thin curves (shifted *sinc*² functions) show the contribution from aliasing. There is relatively little aliasing for $|f| < 200$ Hz.



analyzed. Computed performance of the ideal observer was plotted as percentage correct choices as a function of motion coherence, and “neurometric” thresholds, c_{cell} , were extracted in the same manner described for psychometric thresholds. In general, MT neurons were remarkably sensitive to the motion signals in the stochastic display. In roughly half of the experiments, the neurometric function was statistically indistinguishable from the psychometric function measured on the same set of trials. Across the entire set of experiments, information based on counting the total number of spikes correlated well with the monkeys’ behavior, and the geometric mean ratio of neuronal threshold to psychophysical threshold was 1.19 (Britten et al., 1992). One goal of the present study is to determine whether some temporal characteristics of MT spike trains correlate with either the monkeys’ perceptual performance or our prior measurements of neuronal sensitivity based on integrated spike counts.

From the initial database of 216 cells, four are not considered here. Three of these were not recorded for the full 2 sec (*w044*, *w045*, *w046*), and the fourth cell showed an abrupt change in firing rate during the experiment and was presumed to be damaged (*j036*).

Data analysis

Computation of power spectra. For a real and continuous function $F(t)$ observed between $t = -T$ and $t = +T$, the associated continuous Fourier transform (FFT) \tilde{F} at the frequency f is given by

$$\tilde{F}(f) = \int_{-T}^{+T} e^{-i2\pi ft} F_T(t) dt, \quad (1)$$

where $F_T(t) = F(t)$ for $|t| \leq T$, and 0 outside the observation interval. The autocorrelation function associated with F_T is given by

$$R(t) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^{+T} F_T(t_1) F_T(t_1 + t) dt_1. \quad (2)$$

Following the Wiener-Khinchine theorem, the Fourier transform of the autocorrelation is equal to the power spectrum, that is,

$$S(f) = \int_{-\infty}^{+\infty} e^{-i2\pi ft} R_T(t) dt, \quad (3)$$

or

$$S(f) = \lim_{T \rightarrow \infty} \frac{1}{2T} \tilde{F}(f) \tilde{F}^*(f), \quad (4)$$

where $*$ denotes the complex conjugate. It can be seen that the power spectrum is always real and symmetric.

However, given the discrete sampled nature of our spike trains, we require the use of the discrete Fourier transform with the two associated

problems of (1) the variance inherent in the estimate of the discrete power spectrum and (2) aliasing due to a finite sampling interval. We perform Fourier transforms on spike trains using the standard FFT algorithm and compute one-sided estimates of the power spectral density using overlapping data segments and windowing (Press et al., 1988). To emphasize the difference between the true power spectrum $S(f)$ and one the we compute on the basis of the sampled data, we denote the latter by $S'(f)$. The data submitted to the FFT algorithm is a sequence of 1s and 0s, where each 1 represents an action potential in a spike train sampled at 1 kHz. Transforming a 2-sec-long spike train yields a one-sided spectrum with a frequency resolution of about 0.5 Hz from 0 up to the Nyquist frequency of 500 Hz.

Because we are only interested in studying broad trends in the data over a relatively wide band of frequencies, we do not require such high-frequency resolution. Furthermore, the variance associated with the estimation of the power spectrum can be reduced by using larger-frequency bins, that is, by sacrificing frequency resolution. Thus, we break the 2-sec-long trial into smaller segments, typically using 12 overlapping data segments of 256 msec duration, thereby utilizing 1664 msec of the 2000 msec spike train. We always begin the first segment at 336 msec to eliminate from analysis the transient response to the onset of the stimulus; however, we find essentially no difference in the results when the analysis is performed using the entire 2 sec spike train (not reported here). For each 256 msec segment, a two-sided power spectrum was computed. Given the fact that the power spectrum is always symmetric, we normalize the two-sided spectrum to a one-sided spectrum with 128 entries lying at equally spaced intervals between 0 and 500 Hz (with $\Delta f \approx 4$ Hz).

To improve the spectral estimate further, a triangular Parzen window was applied to each segment to reduce spectral leakage arising due to the finite duration of spike trains (Harris, 1978). Since windowing would otherwise throw away data at the ends of each segment, it becomes important to use overlapping segments. In addition to the averaging due to data segmentation and symmetry for a single 2-sec-long trial, we usually—unless stated otherwise—average the power spectra over all trials for a particular *c* for an individual cell. Since we average in the frequency domain, our method is not sensitive to the exact phase relationship of the response with respect to stimulus onset. When computer-generated data are shown, averages are over the equivalent of 1000 trials of 2000 msec duration each.

While we can reduce the variance in the spectra by averaging, we cannot avoid aliasing due to temporally sampled data. To what extent does aliasing play a dominant role in shaping our spectra? A continuous abstraction of a spike train is a set of occurrence times for action potentials that are idealized as Dirac $\delta(t)$ impulse functions. In recording these occurrence times, the continuous function is not sampled in the usual manner; rather, the action potentials are shifted to nearby sam-

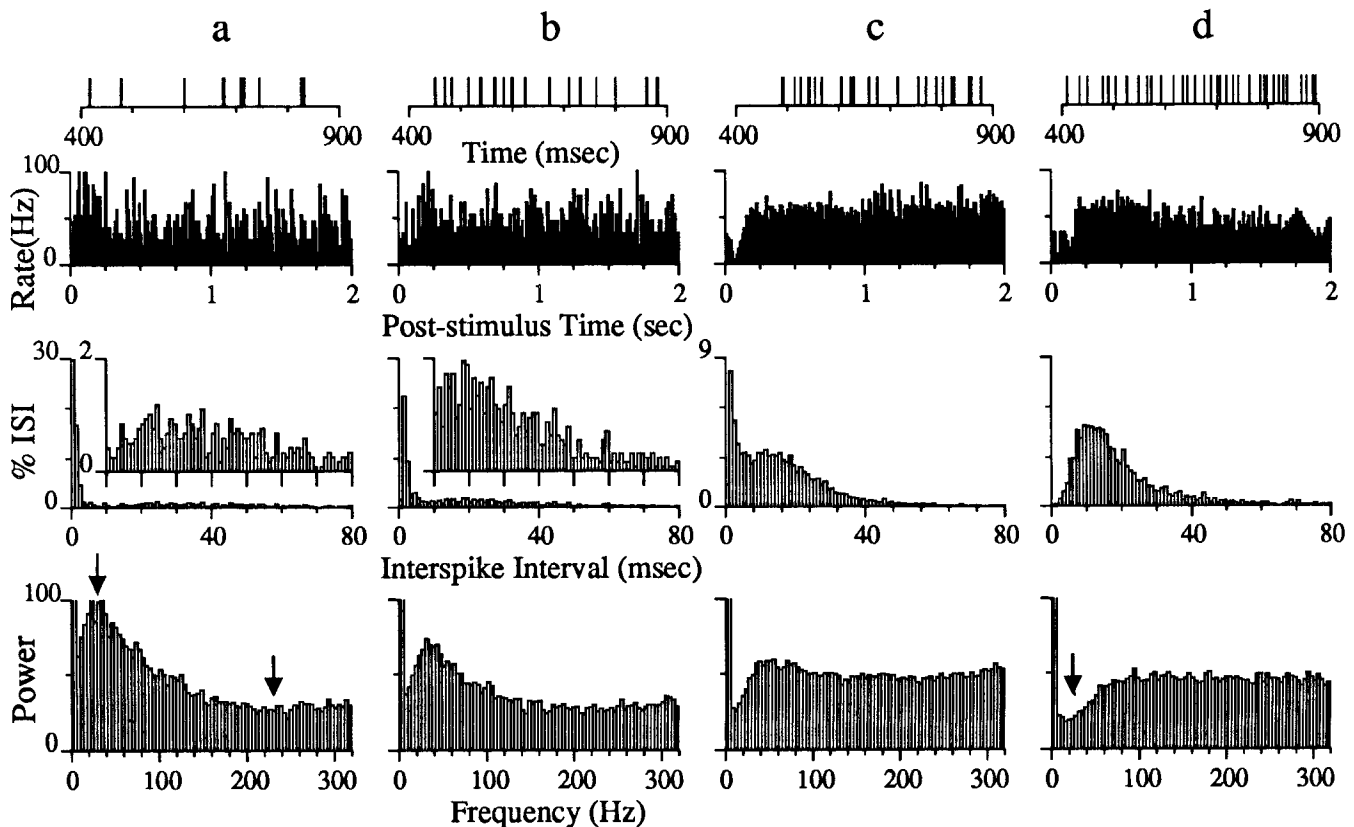


Figure 2. Spike train statistics for four area MT neurons responding to $c = 0.256$ motion. Each column (a–d) corresponds to a different cell, and the cells are arranged from most bursting (a) to least bursting (d). The *top row* shows 500 msec segments of spike occurrence times. The distributions of intervals between spikes are shown by the ISI histograms in *third row*. Over 50% of intervals fall in the 1, 2, or 3 msec bins in ISI (a) while less than 2% do so in ISI (d) (note scale change for ISI in c and d). For a and b ISI insets (horizontal axes begin at 10 msec and remain aligned with main histogram) expand the vertical axis to show the peak near 20–30 msec, which corresponds to intervals between bursts. Power spectral densities, $S'(f)$ (*bottom row*), have peaks in the 20–60 Hz range for the cells (a–c), which have peaks in the ISI below 4 msec. The dip in $S'(f)$ at low frequencies (arrow in d) indicates the presence of a refractory period for spikes or bursts. All four cells fire at roughly similar average rates (40–60 spikes/sec; see PSTHs, *second row*) and the response is relatively maintained following the transients during the first 200 msec of stimulation. Plots are based on averages from 15 trials (a, b), 60 trials (c), and 30 trials (d). Frequent burst firing causes PSTHs (a and b) to be excessively noisy.

pling points. If this were not the case, then most action potentials would be missed altogether. This type of sampling is a form of data binning in which each bin has width equal to the sampling interval and the assumption is made that at most one event occurs per bin. Given a sampling interval of 1 msec, this is a reasonable assumption for most neurons. We present a modification of Poststimulus time histograms (PSTHs) to understand what effect this has on the spectrum.

Figure 1 summarizes this analysis for a special case that has a simple and revealing solution. The left column illustrates how the binning process that converts the continuous spike train (in a) into sampled data (in e) is described as a convolution with a boxcar function (in b) followed by multiplication with a comb function (in d). We assume the special case where the continuous spike train is totally random (Poisson) and therefore has a flat power spectrum (in A) with a δ -function at the origin (which we ignore here; see Stochastic models, in Results). Convolution with the boxcar in the time domain corresponds to multiplying the power spectrum by the square of a *sinc* function [$\text{sinc}(x) = \sin(x)/x$; in Fig. 1C]. Multiplying by the comb function in time corresponds to convolving the power spectrum with a comb with inverse spacing, which replicates and sums copies of the *sinc*² function at 1 kHz intervals (Fig. 1C). The sum of the original *sinc*² plus the infinite number of *sinc*² functions shifted by 1 kHz is a flat spectrum, which is what we started with; however, only the copy of the *sinc*² centered at zero is contributing to the true power spectrum, while the other shifted copies are contributing aliased frequencies (Fig. 1E). In this case, frequencies below about 200 Hz are relatively uncorrupted and this is the frequency range that concerns us. As discussed below, the power spectrum of most cells in the database is relatively flat, especially at higher frequencies. The worst

cases of aliasing occur for neurons that have very regular spike trains (Schild and Schultens, 1986), at least an order of magnitude more regular than we observe here.

Quantification of spectrum shape. In Experimental results (below), we utilize a scheme for classifying cells based on the shape of their power spectra (for a related approach, see Ghose and Freeman, 1992). The classification of a cell depends upon the classification of its set of power spectra, one at each c value. Here we describe in detail how to compute the set of spectra for a cell, and how to classify the cell based on that set.

The estimated power spectral density for a cell at a given c , $S'(f)$, is computed as the average of the spectra for all spike trains recorded at that c value normalized by their average spike rate. For this normalization, the computation of average spike rate must take into account the effect of the overlapping data windows and the shape of the Parzen window. These corrections are particularly important for spike trains with very few spikes. If the individual spike trains were Poisson, then the expected value of each normalized spectrum would be 1.0 at all nonzero frequencies. This normalization reduces the variance in S' due to fluctuations in spike rate from trial to trial while preserving relationships between peaks and dips within each spectrum.

As discussed in Experimental results, many cells have peaks in their spectra in the region below 60 Hz (see Fig. 2, cells a and b, bottom), and we wish to measure the height of that peak with respect to the associated dip or lowest baseline level at higher frequencies [see arrows on $S'(f)$, Fig. 2, cell a, bottom left] and at the same time determine if the peak is present for all c . The location of the peak (if one exists) is computed by sliding a fixed-width window along the set of spectra while looking for the frequency f_p (at the center of the window) that maximizes

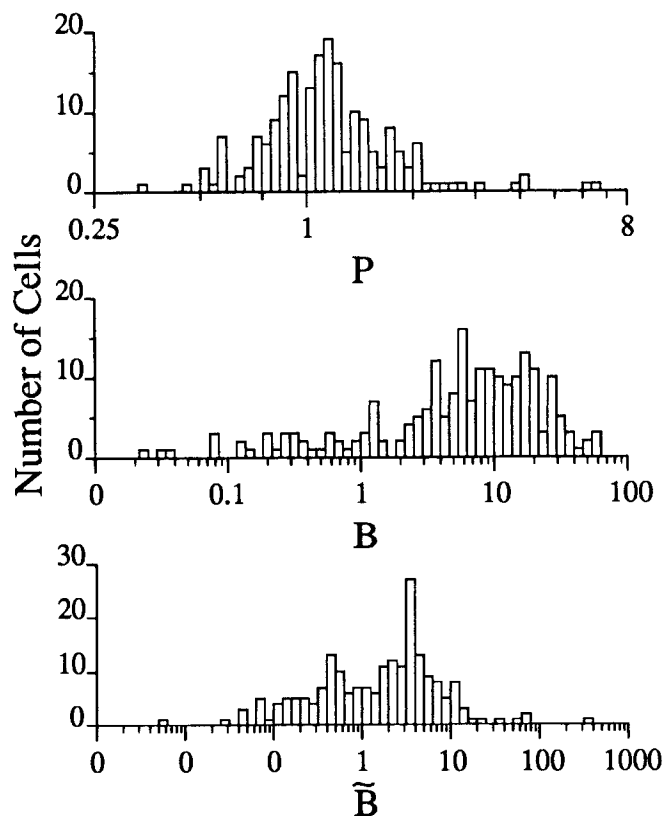


Figure 3. Frequency histograms for our database of 212 MT cells for P , the measure of the shape of $S'(f)$, and two measures of burstiness, B and \tilde{B} , averaged over trials for all c . The *top two plots* show that the distribution of cells with respect to the statistics P and B (the fraction of the ISI in the 1, 2, and 3 msec bins) is primarily a continuum. The dip near 1.0 in the distribution of P is an artifact of the classification of cells as burst or nonburst (see Methods). All burst cells have $p > 1.0$; nonburst cells, $p < 1.0$; and mixed cells that are neither one or the other have P near 1.0. The distribution of B has a long left tail due to cells that rarely fire a second spike within 3 msec. The *bottom plot* shows the distribution for a second measure of burstiness, \tilde{B} , the ratio of the 2 msec bin to the 5 msec bin in the ISI histogram. B and \tilde{B} are highly correlated for our database, but we use B since it is less sensitive to noise than the ratio measure.

the integral within the window summed over all S'_c . The constraint $20 \text{ Hz} < f_p < 60 \text{ Hz}$ is used to keep the window away from the peak at $f = 0$ (i.e., the DC component) and to avoid scanning higher frequencies where peaks are absent. The average value of S'_c within the window will be called the peak level, P_p . A second sliding window is used in a similar, but minimizing, procedure to find the lowest point, or baseline level, at higher frequencies. The location of the lowest point is quantified by f_b , the center of the window, and the average value within the window, P_b , will be called the baseline level. We constrain f_b to be greater than f_p and less than the cutoff frequency, 500 Hz. All windows are seven spectral bins wide (3.9 Hz/bin) and therefore sacrifice accuracy of localization for noise immunity.

Once the peak frequency f_p and the baseline frequency f_b are determined, a cell is classified as a *burst cell* (below we explain the link between peaks and burst firing) if for at least 90% of all c values $P_p > P_b$, that is, if at nearly all coherence levels, the spectrum has a peak in the 20–60 Hz range. If a cell fails to be classified as a burst cell, then we attempt to localize a dip in the power spectrum [see arrow on $S'(f)$, Fig. 2, cell *d*, bottom right, which has a dip near 20 Hz] using another sliding window. We take f_d to be the center frequency of the seven-bin window that minimizes the integral within the window summed over all c , where $20 \text{ Hz} < f_d < 500 \text{ Hz}$. We take the average value of the spectra within a window centered at f_d to be P_d . We classify a cell as *nonburst* if for at least 90% of all c values $P_d < 1.0$, that is, if at nearly

all coherence levels the spectrum has a dip below the expected baseline level (which manifests itself at high frequencies) for a Poisson-like spike train, which is 1.0 due to our spike rate normalization. This definition would result in classifying a cell with Poisson-distributed spikes (that therefore has a flat power spectrum) as neither burst nor nonburst, but since all cells studied here show evidence of refractory periods, this case does not occur in practice. Note that for a pacemaker cell “oscillating” in the 20–60 Hz band, P can become arbitrarily large as the oscillation becomes increasingly regular.

If a cell fails to be classified as either burst or nonburst, it is classified as *mixed* since at some c values it lacks a significant peak, while at others it lacks a significant dip, in the 20–60 Hz range. To avoid classifying a cell based on too little data, trials with less than six spikes in the 336–2000 msec time window are discarded, and c values with less than eight valid trials are not represented in the set of spectra. Finally, a classification is made only when there are at least three different c values with valid S'_c . Typically, there are six S'_c for $c < 0$, six for $c > 0$, and one for $c = 0$.

A measure of the shape of the power spectrum, P , is associated with each classified neuron. For burst cells, the ratio of the peak to the baseline is used: $P = P_{\text{burst}} = P_p/P_b$. For nonburst cells, the ratio of the dip level to the ideal baseline, 1.0, is used: $P = P_{\text{nonburst}} = P_d$. The value for mixed cells depends on the subcategorization; that is, P follows the definition for burst if S'_c had a peak for the majority of c values but follows the definition for nonburst otherwise. When discussing the shape of power spectrum, we will simply refer to P when the particular definition is understood from context.

Other methods. The PSTHs are computed from the single-trial data by averaging over all trials with identical stimulus conditions, using a bin width of 10 msec. They are normalized to show spike rate rather than counts per bin. Interspike interval (ISI) distributions are computed with 1 msec bin width. Power spectra, $S'(f)$, are usually normalized to match continuous spectra under the assumption that spikes can be described as Dirac δ -functions. In this case, the vertical offset is roughly proportional to the spike rate, and for nonburst cells, the flat level at higher frequencies is usually an accurate reflection of the spike rate, as in Figure 12. Under the second spike-rate normalization (discussed above), spectra are divided by the average spike rate so that all are nearly the same height to allow comparison of shapes, as in Figure 5.

Vertical truncation of histogram plots is indicated by open histogram bars near the top of the graph (Figs. 2, bottom; 5, upper left; 15, both ISI plots and lower right; 18, lower left and right).

Results

Experimental results

We begin by describing the population of MT cells with respect to two statistical measures of the temporal fine structure of spike trains: the ISI distribution and the power spectrum $S'(f)$. The first measure is an order-independent statistic since it contains no information about the temporal order in which the intervals occur. For instance, all short intervals could have occurred at the beginning of each trial and all long intervals at the end, or each short interval could have been followed by a long one. The estimate $S'(f)$ is order dependent since it depends on temporal relationships between events at scales beyond single intervals. Although many different $S'(f)$ may be associated with a particular ISI distribution, we find for this database that the shape of the ISI distribution predicts the shape of the power spectrum quite well, and that the tendency of a cell to fire bursts of action potentials is the basis for the prediction.

Estimates of the ISI and $S'(f)$ are shown in Figure 2 for four cells from the database. Segments of typical spike trains from each cell are shown at the top, and below them, the PSTHs show that the average firing rate is relatively constant throughout the period over which we compute the ISIs and power spectra, from 336 to 2000 msec. Although not shown here, including the initial transients had little effect on the shape of $S'(f)$. The autocorrelation functions, $R'(t)$ (not shown), for these cells do not show ringing, even when the associated spectrum $S'(f)$ has a prom-

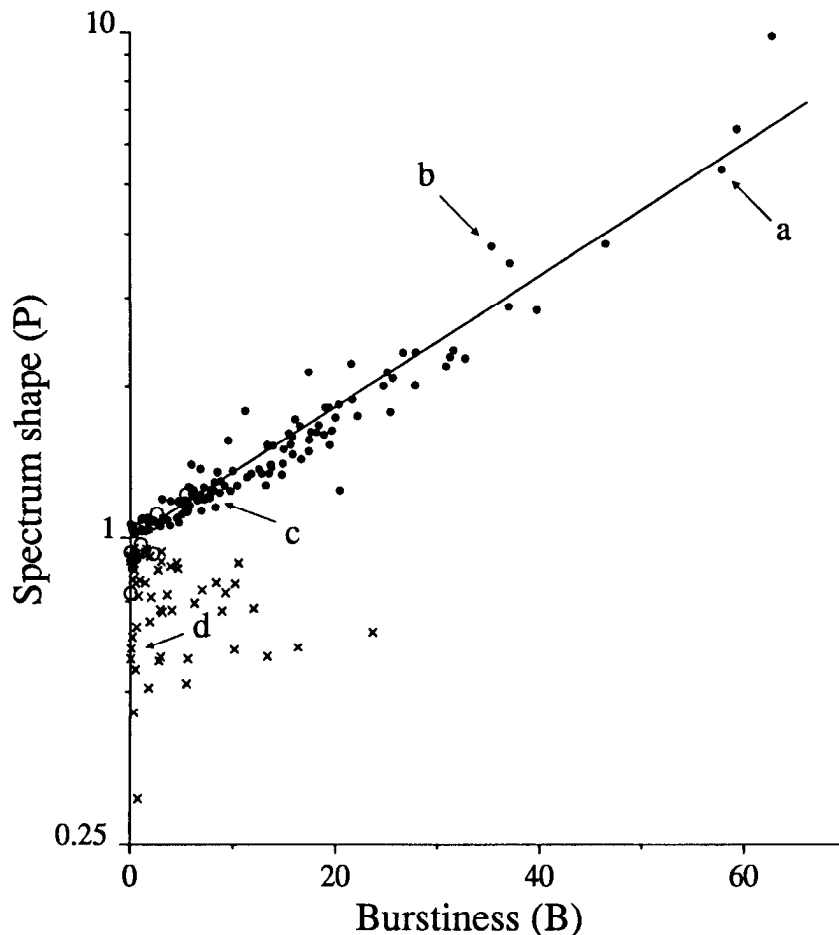


Figure 4. $\log(P)$ is plotted against B for burst (solid circles), nonburst (crosses), and mixed (open circles) cells. Among burst cells, the correlation is strong between $\log(P)$ and B , as shown by the fit from linear regression (line). Hypothetical cells firing Poisson spike trains would fall along the horizontal line $p = 1$, separating the burst and nonburst cells. Mixed cells straddle this line and are neither strongly burst nor nonburst. Since B often has a small but systematic variation with spike rate, values shown here are computed by averaging only over trials with responses that are statistically indistinguishable from those at $c = 0$. Results are very similar when values are averaged over trials for all c . The arrows show the points corresponding to the four cells of Figure 2.

inent peak of the type seen in Figure 2 (cell *a*, bottom). Only a single cell showed strong ringing in $R'(t)$, and this was for $c = 1$.

Bursting cells

We were intrigued by the persistent tendency of certain cells to burst frequently, that is, to discharge a group of two to four tightly clustered action potentials with ISIs of no more than 3 msec (see, in particular, the ISI histogram in Fig. 2, cells *a* and *b*). Cells that respond frequently in this manner to the visual stimulus show a bimodal ISI histogram with a pronounced peak at short, that is, 1–3 msec, intervals and a second, much lower and broader, peak at longer intervals (see ISI insets in Fig. 2, cells *a* and *b*). The first peak is caused by the interval distribution within a burst, while the second peak is partly due to intervals between bursts (interburst intervals, IBIs) and partly due to intervals between isolated spikes.

Figure 2 illustrates the trend in burstiness observed in our data, from strongly bursting (cell *a*) to weakly bursting (cell *c*) and nonbursting cells (cell *d*). Figure 2 also reveals a second and correlated trend, that some cells have a peak in their power spectrum between 20 and 60 Hz (cells *a* and *b*, somewhat in cell *c*), while others (cell *d*) have a spectrum that is flat with a dip at low frequencies (except at the origin $f = 0$). The mean center frequency of the peak in the estimated spectrum S' for all cells with such a peak is 41 Hz (with an SD of 9 Hz), so in the remainder of this article, we will refer to this frequency range as the 40 Hz band, with the understanding that the location of the peak varies from cell to cell. Of 212 cells analyzed, 71 had

relatively flat averaged spectra, $S'(f)$, with a dip at low frequencies, and 131 had peaks in the 40 Hz range of their averaged spectra. The remaining 10 cells had too little data to judge accurately the shape of the spectrum.

As described in Methods, we classified the cells as either *burst* or *nonburst* based on the shape of their power spectra using a scheme that takes into account the possibility that the spectra might change as the stimulus motion coherence c varies. The criteria for this classification were designed to be strict so that cells that did not always, regardless of stimulus direction and coherence, behave as burst or nonburst would be classified as *mixed*. Of 212 cells, 10 were eliminated from classification because they did not meet the minimum standards for number of spikes, trials, and coherence levels. Of the remaining 202 cells, 125 were classified as burst, 61 as nonburst, and 13 as mixed (in three cases, visual inspection disagreed with the classification algorithm, so these cells were dropped from consideration). It is striking that 93% of the cells were classified as either burst or nonburst because it indicates that this rough categorization is stimulus invariant.

As also discussed in Methods, we define a measure P of the shape of the power spectrum. For burst cells, P is defined as the ratio of the height of the peak in the 40 Hz range to the level of the baseline dip at higher frequencies (see arrows on $S'(f)$, Fig. 2, cell *a*). For nonburst cells, P is defined as the ratio of the level of the dip at low frequencies (see arrow on $S'(f)$, cell *d*) to the flat level at higher frequencies. For an ideal Poisson cell, $P = 1$ regardless of the classification as burst or nonburst; for a

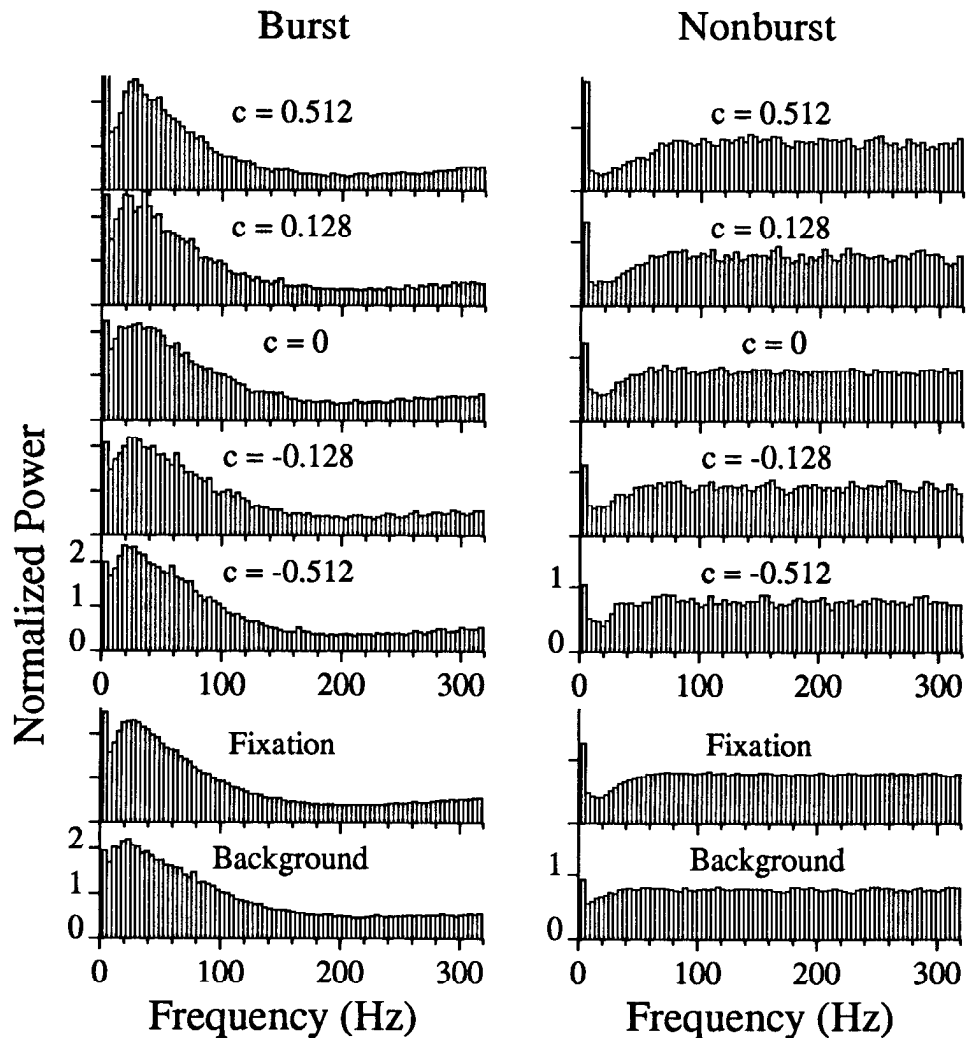


Figure 5. Power spectra, $S'(f)$, are shown for motion coherence values ranging from highly coherent preferred direction motion ($c = 0.512$) to highly coherent null direction motion ($c = -0.512$) for the burst cell (a) and the nonburst (d) cells from Figure 2. Spectra are also shown for background, that is, spontaneous activity, and fixation conditions. The spectra vary little, except that the dip below 20 Hz becomes more prominent when spike rate increases. This can be explained by the potentially greater effect of a refractory period at higher spike rates. These spectra are normalized by spike rate.

pacemaker cell “oscillating” in the 20–60 Hz band, P becomes arbitrarily large as the oscillation becomes increasingly regular.

With P as a measure of the shape of S' , we developed a measure of burstiness based on the ISI. Given the well-known distinction between bursting and nonbursting cells based on intracellular current injections in rodent slice preparations (McCormick et al., 1985; Connors and Gutnick, 1990), we attempted to find a metric that would classify all of our cells into two (or more) segregated groups according to the degree of burstiness. For this purpose, we introduce the measure B as the percentage of the ISI histogram in the 1, 2, and 3 msec bins.

The variable B is similar to other proposed measures of burstiness that are based on the proportion of the ISI distribution below a cutoff value (Cattaneo et al., 1981b; Abeles, 1982). We also considered another measure of burstiness, \tilde{B} , based on the ratio of the number of intervals in the 2 msec ISI bin to the 5 msec ISI bin. This variable has the potential advantage that it is able to distinguish between a bursting cell with a bimodal ISI histogram and a very fast firing cell that has a unimodal ISI histogram concentrated below about 10 msec. However, \tilde{B} is quite sensitive to fluctuations in the trough between the peaks of a bimodal histogram, and its value is less stable. We will use B as the measure of burstiness here but point out that B and \tilde{B} tend to be highly correlated, at least for our database.

Figure 3 shows the frequency distribution for these three statistical measures for all cells averaged over all stimulus conditions. We interpret the histogram for P to represent a unimodal distribution. The dip at unity is an artifact of our classification system because P is based on regions of the power spectrum that are chosen for maximizing the peak-to-trough ratio or minimizing the trough-to-baseline ratio for burst and nonburst cells, respectively. The long left tail of the distribution for B shows that many cells have less than 1% of their intervals shorter than or equal to 3 msec, such as cell d of Figure 2. The distribution for \tilde{B} is spread over many orders of magnitude and shows a hint of bimodality. Overall, however, it is difficult to segment the data into two classes based on these histograms, since many burst and nonburst cells fall in overlapping regions in the histograms for B and \tilde{B} . We stress, therefore, that the burst and nonburst classifications are primarily tools for defining two ends of what appears to be a continuum.

For burst cells, P changes relatively little with stimulus condition and appears to reflect primarily an intrinsic property of these cells in an alert and trained monkey. As we show next, in such cells P is highly correlated with B . For nonburst cells, B often changes systematically with spike rate and is therefore not as revealing about intrinsic properties.

The close connection between bursting and the shape of the

power spectrum is illustrated in Figure 4. Here the value of B for individual cells is plotted against the associated P . The values shown here are averaged over trials at $c = 0$ and trials at other low c values for which the monkey's and the neuron's responses were statistically indistinguishable from responses at $c = 0$. The three groups, burst, nonburst, and mixed, are plotted together in Figure 4. For the burst cells (solid circles), there is an obvious strong correlation between $\log(P)$ and B , with the line indicating an empirical, exponential fit: $P = e^{0.03B}$. Thus, the more a cell tends to fire action potentials in tight bursts rather than as isolated spikes, the higher the peak in its power spectrum in the 40 Hz band. We will explore the reasons for this behavior further in Stochastic models, below. Cells with purely Poisson-distributed spikes would fall along the horizontal line $P = 1.0$, with low firing rates near the origin, and higher firing rates to the right. The strong correlation evident here between B and P justifies classifying cells as either "burst" or "nonburst" based on the shape of the power spectrum.

Relation of the 40 Hz peak to prior measures of neuronal and psychophysical performance

Previous analyses of this data set have identified several interesting parallels between the psychophysical performance of the monkeys and the responses of single MT neurons. In all of these analyses, the response of a neuron was considered to be the total number of spikes occurring during the period of visual stimulation (Newsome et al., 1989a,b; Britten et al., 1992). We now consider to what extent temporal structure, here the presence and amplitude of the 40 Hz peak in the power spectrum, reflects any aspect of the visual stimulus or the monkey's behavior. Specifically, we ask the following questions. (1) Does the prominence of the peak vary with the strength of the motion signal, c ? (2) Is the peak affected by the behavioral state of the animal? (3) Does the spectral peak develop or change with time during the course of a 2-sec-long trial? (4) How is the peak correlated to prior measurements of cell sensitivity based on total spike counts? (5) Is the size of the peak correlated with the monkey's decisions concerning direction of motion for a particular stimulus condition?

We suggested in a previous section that the burstiness of a cell, quantified by the ratio P , is largely independent of the strength of the motion signal in the visual stimulus c . Qualitatively, this point is supported by the spectra illustrated in Figure 5, computed for one burst and one nonburst cell, which appear fairly constant in shape for varying values of c , and by Figure 6, which shows the mean and SD for P as a function of c for a burst and nonburst cell.

To analyze the relationship of P and c quantitatively, we first conducted a one-way ANOVA for each neuron to determine whether P varied significantly within the range of c tested; 118 of 202 neurons (58%) failed to show heterogeneity by this test ($p > 0.05$), and we conclude that P is completely independent of c for these cells. For the remaining neurons we performed a multiple regression analysis to determine whether c influenced P in a systematic manner. For the great majority of our neurons, the mean response increased with c , causing changes in the shape of the power spectrum that are related to the presence of the refractory period (see Stochastic models, below). We therefore included mean firing rate as a coregressor in our multiple regression analysis to disentangle the effects on P of c and mean firing rate.

Only 20 of 202 MT neurons (10%) showed a significant re-

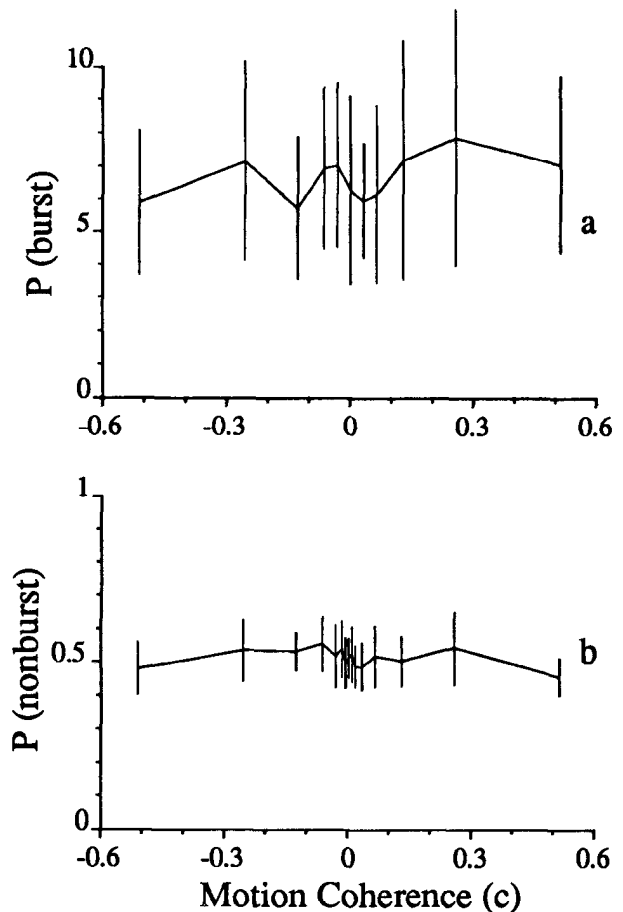


Figure 6. For most cells, P does not change substantially over the range from $c = 0$ to values well above the threshold level. The average value of P over all trials at each coherence level is plotted for a strongly bursting cell (top) and a strongly nonburst cell (bottom). Error bars show SDs.

lationship of P and c (multiple regression, $p < 0.05$); the slope of this relationship was negative for 13 cells and positive for the remaining seven cells. For all neurons, furthermore, the slope of the regression line relating P to c was sufficiently small that the classification of a cell as burst (mean $P > 1$) or nonburst ($P < 1$) was unambiguous. It appears, therefore, that c has no strong or systematic impact on P for our population of MT neurons considered as a whole. In most of our subsequent analyses, therefore, P is averaged across c to obtain a single index of burstiness for each cell.

Is the spectral peak influenced by the behavioral state of the animal? To answer this question, we computed the index of burstiness, P , for spike trains obtained under three different behavioral conditions. In the "choice" condition, the animals attended to the random dot stimuli with the intent of making a discrimination. The same visual stimuli were presented in the "fixation condition," but the animals were only rewarded for maintaining fixation on a visual target; no discrimination was required. Finally, the "background" refers to spontaneous neuronal activity that was acquired during the interval between trials.

The index of burstiness, P , did not vary between the "choice" and "fixation" conditions for the 82 cells for which "fixation" data was available (paired t test, $p > 0.05$), suggesting that the

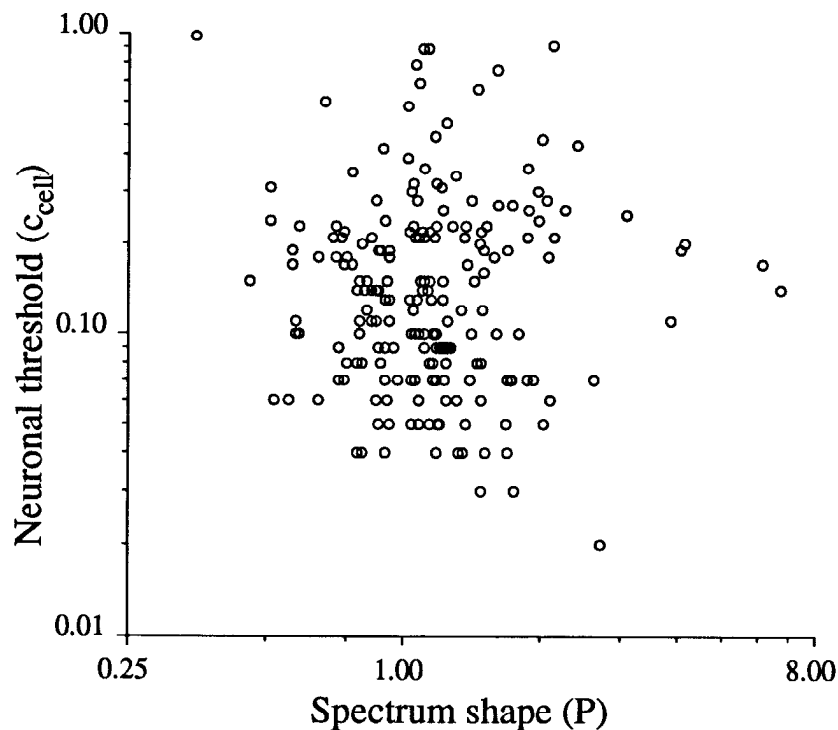


Figure 7. A comparison of neuronal threshold to the shape of the power spectrum. There is no significant correlation ($r = 0.045$, $p = 0.61$) between $\log(c_{\text{cell}})$ and $\log(P)$. Neuronal threshold, c_{cell} , is the coherence level that supports 82% correct decisions by an ideal observer counting total number of spikes. P is the ratio of the height of the peak (dip) in the 40 Hz band of $S'(f)$ to the baseline level for burst (nonburst) cells. We obtain similar results when correlating P against measures of the animal's psychophysical performance on motion discrimination (not shown).

monkey's intent to make a discrimination had no effect on the spectral peak. There was a significant change in P between the choice and background conditions (paired t test, $p < 0.05$), but the effect was quite small: P decreased by an average of 3% for burst cells ($n = 122$) and increased by an average of 8% for nonburst cells ($n = 59$). We therefore infer that P is substantially independent of behavioral state, a conclusion that is supported by visual inspection of power spectra like those illustrated at the bottom of Figure 5.

We next inquired whether the spectral peak developed or changed with time during the course of a 2 sec trial. Analyzing the evolution of the shape of the power spectrum is difficult due to the small amount of data that most spike trains contain in a period as short as a few hundred milliseconds. The average spike rate over our entire database is 19 Hz (with an SD of 18 Hz), so the exact placement of any one spike will have a large contribution toward the overall shape of the power spectrum for short windows. Because of this, we address a special case of this question that allows averaging over trials. This method is therefore limited to detecting changes that are locked to the stimulus onset.

We divided each trial (starting 336 msec after the onset of the stimulus to eliminate initial transients) into six equal time windows that overlapped by one-third of their width. The average value of P was computed from individual spectra for all windows of similar time lags that fulfilled a minimum spike criterion of five spikes per window. Only 10% of burst cells and 21% of nonburst cells showed a significant correlation between P and time (Spearman rank-order correlation coefficient, $p < 0.05$). Of those cells, P increased by an average of 11% and 14% for burst and nonburst cells, respectively. Because P shows no correlation with time during the trial for 86% of cells and changes little for the other cells, we compute only one spectrum per trial in other analyses.

In a prior analysis of this data set, signal detection theory was

used to compute a neuronal "threshold" that expressed the sensitivity of each neuron to motion signals in the display (Britten, et al., 1992; see Methods). Threshold was defined to be the coherence value at which the neuron signaled the direction of motion with a criterion level of reliability. Thresholds varied widely among neurons in the data set, and we therefore tested for the hypothesis that burstiness as measured in the present analysis could be systematically related to the measure of sensitivity computed in the prior study. Figure 7 shows cell threshold plotted against the index of burstiness, P , for all neurons that yielded a reliable estimate of P . The scatterplot contains no structure signifying a relationship between the two measures, an impression that is confirmed by calculation of a correlation coefficient ($r = 0.045$, $p = 0.61$).

Psychophysical threshold also varied across these experiments since the testing conditions were changed to match the preferences of each cell (Britten et al., 1992). For some purposes, therefore, it is useful to express the sensitivity of each neuron relative to psychophysical sensitivity by calculating for each experiment the ratio of neuronal to psychophysical threshold. To determine whether P is related to cell sensitivity expressed in this manner, we calculated a correlation coefficient between the log of the "threshold ratio" and $\log(P)$, but again we found no relationship ($r = 0.097$, $p = 0.28$). Thus, the prominence of a peak in the 40 Hz region of the power spectrum does not correlate with prior measures of cell sensitivity.

Finally, we asked whether the prominence of the peak varied in a systematic way with the decision made by the monkey on successive presentations of a given motion condition. Previous analyses have demonstrated a trial-to-trial covariation between neuronal response and psychophysical decision when the response is considered to be the integrated spike count (Newsome et al., 1989b; K. H. Britten, W. T. Newsome, M. N. Shadlen, S. Celebrini, J. A. Movshon, unpublished observations). We therefore wondered whether a measure of temporal structure in

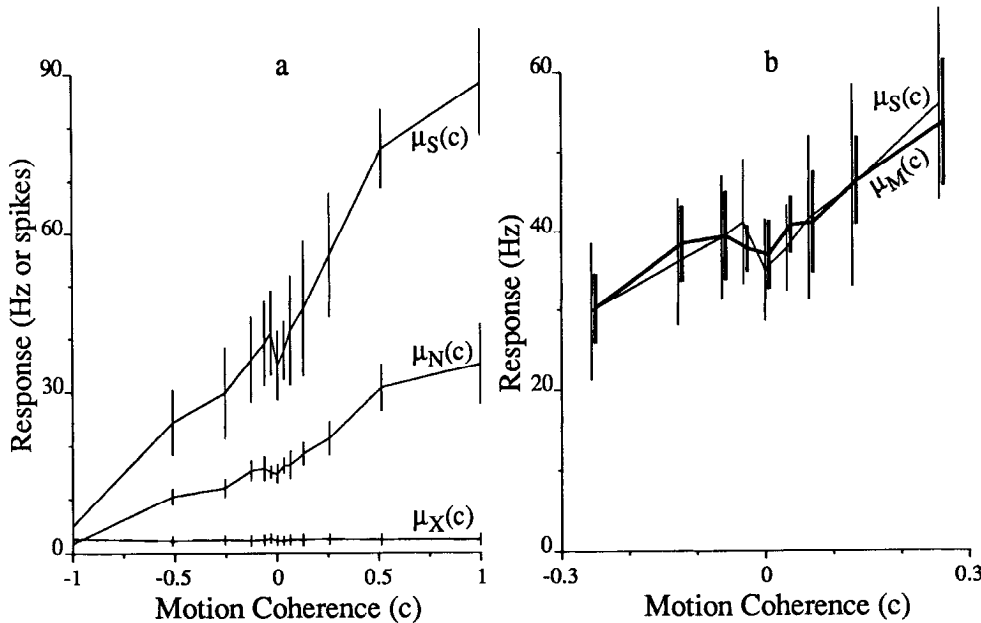


Figure 8. One example of how counting events rather than spikes can yield better direction discrimination. Events are defined to be either isolated spikes or bursts (groups of spikes with all intervals ≤ 3 msec). *a*, The tuning of mean spike rate [$\mu_S(c)$], mean event rate [$\mu_N(c)$], and the number of spikes per event [$\mu_X(c)$] is plotted relative to motion coherence for the strongly bursting cell j001. $\mu_S(c)$ and $\mu_N(c)$ are tuned with c , but $\mu_X(c)$ is flat at about 2.5 and is treated as a constant, μ_X , in the text. *b*, Focusing on the region around $c = 0$, $\mu_S(c)$ (thin line) is plotted against $\mu_M(c)$ (thick line, offset), computed by multiplying $\mu_N(c)$ from *a* by $\mu_X = 2.5$ spikes/event. It is apparent by the smaller SDs and similar slope that $\mu_M(c)$, and thus event rate, is a better basis for predicting c than is raw spike rate. Error bars show SDs.

the spike train, P , would be similarly correlated with performance.

As in the prior analysis, we eliminated the influence of the visual stimulus itself on the monkey's decisions by carrying out the analysis only for $c = 0$ (completely random motion) and other small coherence levels for which neither the monkey nor the neuron discriminated the direction of motion at levels greater than chance (effectively random motion). On these trials, the monkey "guessed" the correct direction since the visual stimulus itself contained no effective information about the correct choice. For each cell we computed the difference between the average value of P for trials on which the monkey decided in favor of the neuron's preferred direction and the average value of P when the monkey decided in favor of the null direction. We found no statistically significant difference in P between these two conditions for either burst or nonburst cells (paired t test, $p > 0.05$), and we therefore conclude that the prominence of the spectral peak in our data set and for our stimulus conditions is not related to the monkey's behavioral choice.

We found a similar result when testing whether P was correlated with correct versus incorrect decisions by the monkey at the coherence level closest to the monkey's psychophysical threshold, c_{system} (typically at $c = 0.128$). At this level, there are a significant number of incorrect response trials, and yet the monkey is not simply guessing. Again, we found no statistically significant difference (paired t test, $p > 0.05$) for either burst or nonburst cells, and therefore conclude that the prominence of the spectral peak is not related to correct and incorrect responses by the monkey.

Since it is a widely held belief that changes in temporal structure (such as an increase in burstiness) can result from cell damage caused by the electrode, we tested for a change in P from trial to trial over the course of the experiment. We found that 13% of burst cells showed a significant increase in P , 67% on average, during the experiment, while 12% showed a significant but small (only 8% on average) decrease in P . Among nonburst cells, 20% showed a significant increase, while 18% showed a significant decrease in P . The magnitudes of the in-

crease and decrease among the nonburst cells were both 10% on average.

Treating bursts as events

We previously appealed to the neuronal threshold, c_{cell} , as a measure of an ideal observer's ability to decide the direction of motion of the stimulus based on the output of the neuron (Newsome et al., 1989a; Britten et al., 1992), assuming that the relevant neuronal output is the number of spikes fired during the stimulus period without considering whether those spikes occurred in bursts or as isolated action potentials. What happens if we quantify the neuron's output by the number of "events," where an event is either a burst or an isolated spike, and recompute an associated neuronal threshold, \tilde{c}_{cell} ? One could well argue on biophysical grounds that a burst of spikes could be more powerful in evoking a postsynaptic response than the same number of isolated spikes.

Consistent with our definition of B , events are defined as the longest sequences of spikes with all ISIs less than or equal to 3 msec (values between 3 and 8 msec give very similar results). With this definition, single isolated spikes as well as bursts are counted as individual events. A nonburst cell will have nearly the same number of events as spikes, while a burst cell will have many fewer events than spikes. Figure 8*a* compares the tuning of a strongly bursting cell's response measured in spikes/second (upper curve), events/second (middle curve), and spikes/event (lower curve). Typical of our database, spikes/event is not tuned with c ; therefore, the curve for events is similar to that for spikes, but scaled down by the average number of spikes per event, here 2.5. In Figure 8*b*, the thick line shows the events/second curve scaled up by 2.5 spikes/event so that it overlays the spikes/second curve. From the relatively smaller SDs for normalized events/second, it is clear that for this cell events/second is a more useful neuronal signal for predicting the direction of coherent motion.

Because the neuronal code that carries motion information in cortex is not known, and since likewise we do not know whether neurons postsynaptic to the one recorded differentiate

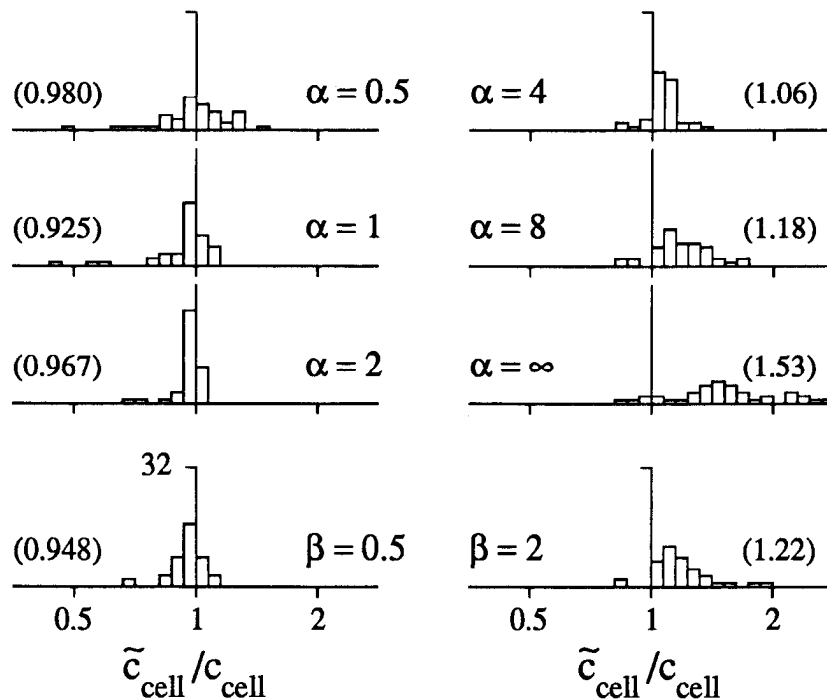


Figure 9. Comparing neuronal thresholds based on spike rate to those based on weighted event counts. For 41 burst cells ($p \geq 1.5$), frequency histograms show the threshold ratio of \tilde{c}_{cell} to c_{cell} . In the *upper six histograms*, \tilde{c}_{cell} is computed from ROC analysis based on the number of single spikes plus α times the number of bursts. The greatest leftward shift in the distribution (*numbers in parentheses show means*), representing the largest average improvement in neuronal threshold, is achieved for $\alpha = 1$, which corresponds to using $\mu_N(c)$, event count, as the neuronal signal. (The counts near 0.5 indicate cells for which the neuronal thresholds were roughly halved by this procedure.) For $2 \leq \alpha \leq 3$, this procedure is very similar to counting individual spikes, since bursts are composed of typically two or three spikes. Histograms for $\alpha = 0.75, 1.5$, and 3.0 (not shown) have means 0.944, 0.940, and 1.02, respectively. As $\alpha \rightarrow \infty$, single spikes are ignored and only bursts are counted. The *bottom two histograms* show results from two schemes that weight events based on the number of spikes per event raised to the power β . The square root yields an improvement in threshold since it reduces the effect of variance in the number of spikes per event, while squaring emphasizes the variance, and worsens the thresholds. Overall, these plots indicate that an ideal observer with knowledge of the arrangement of spikes in bursts will be better able to predict the direction of motion, particularly at near-zero coherence levels, than an observer knowing only the total number of spikes.

between bursts and isolated spikes, we tried various schemes for weighting the contribution of events to the output signal based on a function of the number of spikes per event. First, we weighted isolated spikes, that is, single-spike events, as 1 and bursts, events of multiple spikes, as α , with α varying between 0.5 and 8. We also used a different weighting scheme, where each event, irrespective of whether burst or isolated spike, is weighted according to its number of spikes raised to a power, β . Note that $\beta = 0$ corresponds to the first weighting scheme with $\alpha = 1$, and $\beta = 1$ corresponds to our original scheme, which does not differentiate between bursts and isolated spikes. In addition, we consider $\beta = 1/2$ and $\beta = 2$.

To assess the advantage of these schemes, we recomputed neuronal thresholds based on the modified output signals for the 41 burst cells where the peak in the power spectrum was at least 50% above the baseline ($p \geq 1.5$). More weakly bursting cells are ignored because we expect no effect when isolated spikes greatly outnumber bursts. Figure 9 shows the frequency histogram of $\tilde{c}_{\text{cell}}/c_{\text{cell}}$, where \tilde{c}_{cell} is the neuronal threshold based on the modified signal. The shifts of the distributions are significant ($p < 0.05$) for all histograms shown except for $\alpha = 0.5$. Leftward shifts indicate that the thresholds improved (became lower) when the modified signal was used in place of spike count. The greatest improvement occurred for $\alpha = 1$ (i.e., $\beta = 0$) and corresponds to a 7.5% decrease in threshold. For three cells, thresholds were roughly cut in half. In other words, allowing an ideal observer

to count bursts as single events enhances his ability to predict the direction of motion of the stimulus, on average.

Weighting bursts more heavily ($\alpha > 1$) or less heavily ($\alpha = 0.5$) than single spikes did not improve thresholds. Squaring the number of spikes within the burst also led to higher (worse) thresholds, while taking the square root yielded an improvement.

Based on these results, and on the relative variance-to-mean ratios for event count and spike count seen in Figure 8, we believe that the improvement, particularly for $\alpha = 1$, is due to a reduction in relative variance that occurs by ignoring the number of spikes within events. This effect is easily demonstrated by a simple stochastic model. Consider the model that a bursty spike train is governed by two distributions, that of the number of events N and that of the number of spikes per event X . Assume that N is Poisson distributed with rate parameter $\mu_N(c)$, a function of stimulus coherence, and that X is distributed with mean μ_X and variance σ_X^2 . Using basic results from the theory of branching processes (Feller, 1968), the mean of S , the number of spikes per trial, is then

$$\mu_S(c) = \mu_N(c)\mu_X, \quad (5)$$

and the variance (see Appendix for proof) is

$$\sigma_S^2(c) = \mu_N(c)(\mu_X^2 + \sigma_X^2), \quad (6)$$

where we use the fact that $\mu_N(c)$ is both the mean and variance

of the Poisson distribution for N . Rather than comparing the neuronal output signal S , based on spikes, directly to the event count N , we consider the random variable $M = \mu_N N$, which has the same mean as S , that is,

$$\mu_M(c) = \mu_N(c)\mu_X. \quad (7)$$

This corresponds to multiplying the event/second curve in Figure 8a by 2.5 spikes/event so that it lies directly on top of the spikes/second curve and allows direct comparison of the sizes of the SDs (i.e., Fig. 8b). The variance of the scaled-up event count M is given by

$$\sigma_M^2(c) = \mu_N(c)\mu_X^2, \quad (8)$$

which no longer has the term from Equation 6 involving the variance of the number of spikes within an event. Therefore, counting events rather than spikes should allow an observer to better predict the direction of motion of the stimulus under the conditions of this model. The critical condition here is that the number of spikes per event is independent of stimulus condition, and this appears to be the case, as demonstrated in Figure 8, for most of the bursting cells, particularly at near-threshold coherence levels. This is further supported by our observation that P changes little with coherence level.

Stochastic models

What stochastic models of neuronal firing can give rise to the observed power spectral densities and ISI distributions? To answer this question completely would require a detailed understanding of the biophysics of individual cortical cells as well as the dynamics and connectivity of the network in which the MT cells from which we recorded are embedded. Instead, we focus on the statistical properties of the discharge frequency of individual cells in a qualitative manner, bypassing the need for detailed single-cell or network models. We believe that this method is justified by our results; we can model the power spectra and ISI distributions using very simple two- or four-parameter distribution functions.

Poisson-distributed action potentials

Our starting point is the fundamental result that the power spectrum of a random, Poisson train of action potentials is flat at all frequencies except for a δ -peak at the origin. This agrees with our intuition that all spectral components should be equally represented in a completely random spike train. To be more precise, we will model a spike train as a shot-noise process $f(t)$, where the function $h(t)$ describes the shape of a single shot, here at first a single action potential, and $S_h(f)$ is the associated energy spectrum. A train of infinitely many action potentials is given by

$$g(t) = \sum_{i=-\infty}^{+\infty} h(t - t_i). \quad (9)$$

We assume here that the spike train is generated by a stationary and ergodic process. If the spikes occur at random, that is, without any memory of the previous spike but with an average firing rate λ , the power spectrum of such a random spike train is (Champeney, 1973)

$$S(f) = \lambda S_h(f) + 2\pi \langle |g(t)|^2 \delta(f), \quad (10)$$

where the average value of g is related to $h(t)$ by $\langle g(t) \rangle = \lambda \int h(t)dt$. In the familiar case where we assume that an individual

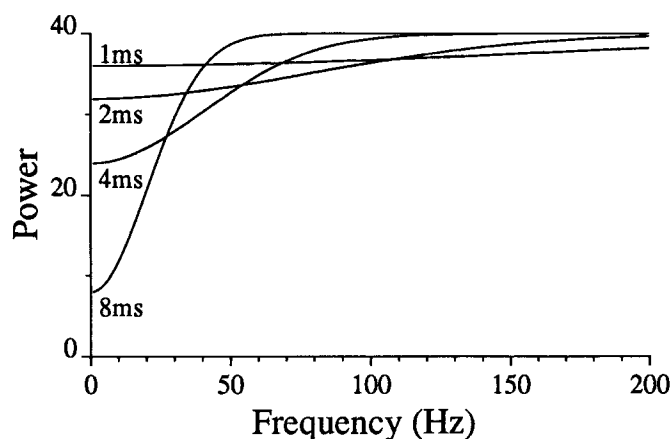


Figure 10. Derived power spectra S_{Poisson} for a random (Poisson) spike train with refractory period. The power spectrum S_{Poisson} for a Poisson process with a refractory period modeled as a Gaussian-shaped depression in the renewal density (see Eq. 15) is plotted for $\sigma = 1, 2, 4$, and 8 msec at $\lambda = 40$ Hz. As the length of the refractory period increases, the trough becomes deeper at lower frequencies. This model for $S(f)$ only holds if $\lambda \leq 1/(\sqrt{2\pi}\sigma)$. Outside of this range (i.e., if the inverse of the mean spike frequency is on the order of the refractory period), the spectrum develops a peak at that inverse of the mean refractory period, and the renewal density can no longer be modeled as a constant minus a Gaussian.

action potential is adequately described by a δ -impulse function, that is, $h(t) = \delta(t)$, $S_h(f) = 1$, the above equation reduces to

$$S(f) = \lambda + 2\pi\lambda^2\delta(f). \quad (11)$$

The interspike interval (ISI) probability density function for this case is given by

$$\text{ISI}(\Delta) = \lambda e^{-\lambda\Delta} \quad \Delta t > 0. \quad (12)$$

Thus, if we observe a Poisson spike train for a sufficiently long time, its ISIs should be distributed according to a single decaying exponential.

We will also make use of a more general expression for the power spectrum that includes cases where the occurrence of an action potential is dependent on the last time an action potential was initiated. We embody this dependency in the renewal density function $p(t)$. Assume that a spike was generated at time t_1 . The probability for the next spike to occur between $t_1 + t$ and $t_1 + t + dt$ is given by $p(t)dt$ (for details, see Perkel et al., 1967). The power spectrum of shot-noise with this dependency among the "shots" is

$$S(f) = \lambda S_h(f) \left[1 + \int_{-\infty}^{\infty} e^{-i2\pi ft} p(t) dt \right], \quad (13)$$

for all values $f \neq 0$ (see Appendix L in Champeney, 1973). For a Poisson process, the probability of spiking per unit interval is always constant and is characterized by the mean rate; thus, $p(t) = \lambda$.

Neurons, however, do not fire totally without memory, because for a variable time following the generation of an action potential the spiking threshold is elevated, making it more difficult to discharge the cell (absolute and relative refractory periods).

The effect of a refractory period can be modeled analytically with the help of the renewal density function. The shape of $p(t)$ can, indeed, be measured directly by computing the probability

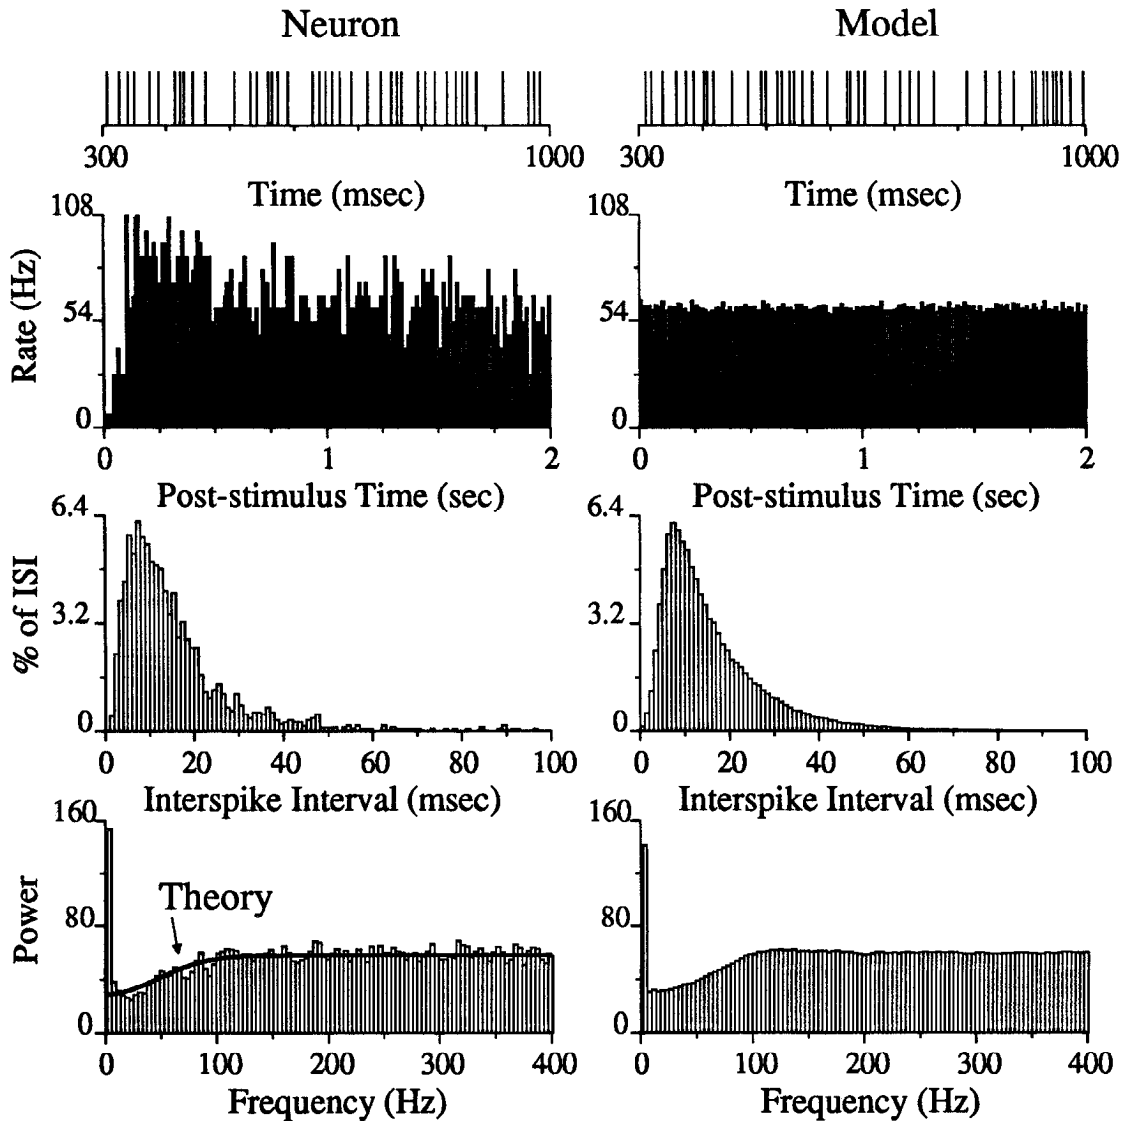


Figure 11. Comparison of spike train statistics for nonburst cell *e047* to those for a simple numerical model. The statistics for the neuron (*left column*) were computed by averaging over 15 trials at $c = 0.128$. The spikes (*top trace*) are less clustered than random, as demonstrated by the absence of short intervals in the ISI and the dip at low frequencies in $S'(f)$. The corresponding numerical model (*right column*) consists of a computer-generated Poisson process (mean rate, 86 Hz) superimposed with a Gaussian-distributed refractory period (mean, 5 msec; SD, 2.0 msec; truncated at 0 and renormalized). The model does not account for initial transients in the data and averages over the equivalent of 1000 2-sec-long trials, so the PSTH is flat and all plots are less noisy for the model. The levels of PSTH and power spectrum (above 200 Hz) demonstrate that the resulting process has an overall mean rate matching that of the neuron. The absence of short intervals in the ISIs and the dips at low frequencies in the power spectra are in close agreement between the neuron and the model. This model is *not* intended to be a best fit for the data, but rather a demonstration that the location and size of the dip are qualitatively accounted for by a random process with a stochastic refractory period of appropriate duration. The *solid curve* superimposed on the neuron's spectrum (*bottom left*) corresponds to the analytical power spectrum for a Poisson process with a refractory period (Eq. 15) with $\lambda = 58$ Hz and $\sigma = 3.5$ msec. Again, this qualitatively accounts for the dip.

for the observed cell to fire an action potential in the short time interval $t_i + t$ and $t_i + t + dt$, assuming that it had fired at time t_i . For the binary data we have here (per sampling interval of $\Delta t = 1$ msec, either zero or one spike can occur), $p(t)$ is directly proportional to the autocorrelation function $R(t)$. For our non-bursting cells (e.g., Fig. 2, cell *d*), $R(t)$ (not shown) is well fitted by a constant minus a small Gaussian around the origin, indicating a reduced probability of firing around $t = 0$. We therefore assume for the renewal density

$$p(t) = \lambda - \lambda e^{-(t^2/2\sigma^2)}. \quad (14)$$

Replacing $p(t)$ into Equation 13 yields the power spectrum of

an infinite train of Poisson-distributed δ -impulses with refractory period

$$S_{\text{Poisson}}(f) = \lambda(1 - \sqrt{2\pi}\lambda\sigma e^{-2(\pi f\sigma)^2}), \quad (15)$$

for $f \neq 0$. In order to ensure that S_{Poisson} is always positive, the maximum firing rate must be limited: $\lambda \leq 1/(\sqrt{2\pi}\sigma)$. This spectrum, parameterized by two parameters, the mean rate λ and the width of the refractory period σ , is constant for large values of f but dips toward its minimum at $f = 0$. Figure 10 shows $S_{\text{Poisson}}(f)$ for $\lambda = 40$ Hz and for $\sigma = 1, 2, 4$, and 8 msec. A longer refractory period causes a deeper trough at low frequencies. Note that this result appears at odds with intuition,

Cell Name	λ (Hz)	σ (msec)
w014	80.3	2.0
j233	56.0	1.5
j200	46.8	1.0
e085	39.3	7.5
w213	30.6	6.5
j129	10.8	13.0

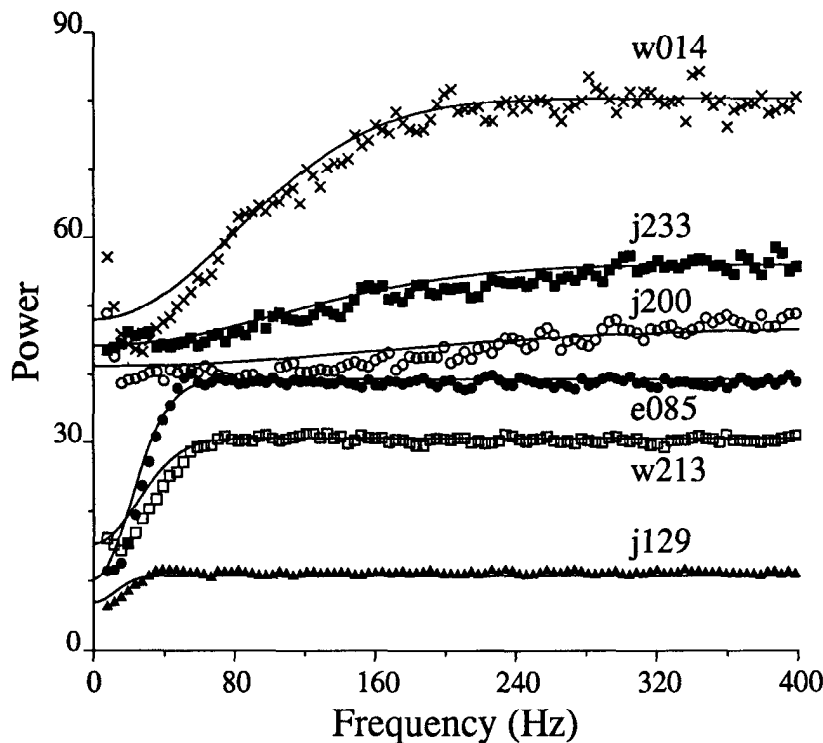


Figure 12. A comparison of power spectra, $S'(f)$ (data points) for six non-bursting MT cells and their corresponding analytical curves, $S_{\text{Poisson}}(f)$ (solid lines), based on the expression for a process with randomly (Poisson) distributed spikes and a refractory period (Eq. 15). These examples illustrate the ability of the simple analytical model to account qualitatively for the location and size of the dip in the power spectrum over a broad range of firing rates. The particular shape of the dip is often fit poorly since the form of the renewal density may not match the Gaussian shape imposed by analytical model (see Eq. 14). The different levels of the various power functions reflect the different mean spikes rates. The parameters of the analytical expressions are shown in the table.

since a refractory period seems to demand a peak in the neighborhood of the inverse of the smallest ISI. However, this is only true if the firing rate λ is so high that the mean time between spikes approaches the refractory period. In that case, Equation 15 no longer describes the resulting spectrum. Clearly, in the case of a fixed absolute refractory period, the mean spike rate must be no greater than the inverse of the refractory period. The additional factor of $1/\sqrt{2\pi}$ arises because of the Gaussian depression in the renewal density used to model a stochastic refractory period.

We also numerically simulated this situation using Poisson-generated shot-noise (with a mean spiking rate $\lambda = 58$ Hz) and a Gaussian-distributed absolute refractory period (of 5 msec mean, 2 msec SD, and truncated below zero). In other words, each time a spike was generated, the Gaussian distribution specified the refractory period associated with that particular spike. After this refractory period, the probability for the next spike to occur is a constant $\lambda \Delta t$. We compare in Figure 11 the spectrum and ISI distribution from an MT cell with a relatively flat spectrum and a dip at low frequency against this simple model. Both the synthetic and the experimentally recorded spike trains (see Fig. 11, top) are subject to the same analysis. The associated PSTH is flat for the computer-generated process since our model

does not account for the transient component of the neuron's response. It is obvious that the ISI distribution and power spectrum for the synthetic process are very similar to those for the MT cell. In particular, both spectra show a dip at low frequencies.

We superimposed the analytical expression $S_{\text{Poisson}}(f)$ (with $\lambda = 58$ msec, $\sigma = 3.5$ msec) onto the power spectrum for the neuron (Fig. 11); it appears to provide an excellent fit to the computer-generated and the measured spectra. We performed this fitting procedure of $S'(f)$ against the two-parameter function $S_{\text{Poisson}}(f)$ of Equation 15 for 61 nonburst MT cells. Six examples of the fits are shown in Figure 12 for neurons with various firing rates, and the values of λ and σ for all nonburst cells are shown in Figure 13. As expected, the refractory period shortens as firing frequency increases. Altogether, we find it remarkable that such a simple stochastic model of cell firing accounts for the shapes of the power spectra of many cells in our database.

Poisson-distributed bursts

Can we use a similar model to account for the 41 MT cells with a peak in their power spectrum whose amplitude is at least 50% higher than the baseline? To answer this question, we consider

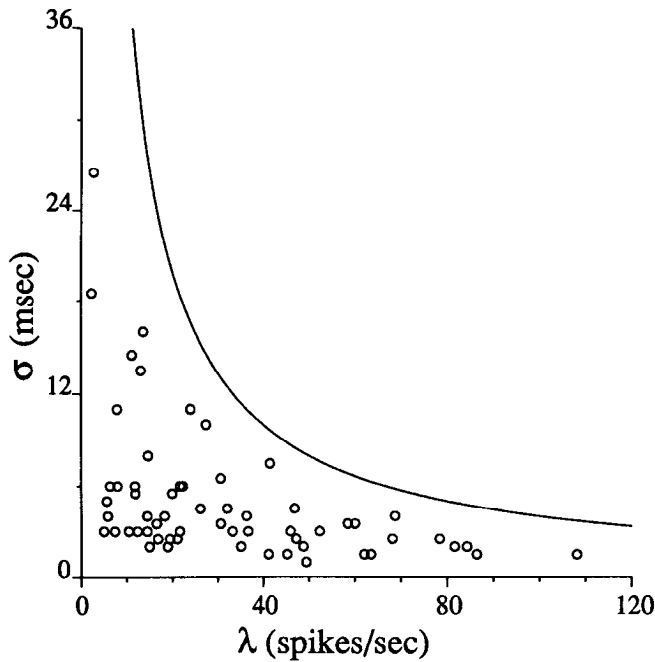


Figure 13. The distribution of parameters for fits of nonburst cells' spectra to the analytical model of the power spectrum of a Poisson process with a refractory period (see Eq. 15). The refractory period parameter σ is plotted against the mean firing rate λ for 61 nonburst cells (data points). The solid line shows the boundary outside of that the model no longer holds, that is, for which $\lambda > 1/(\sqrt{2\pi\sigma})$. To the upper right of this line, the firing rate becomes too high to support the corresponding refractory period under our model of the renewal density.

the IBI (interburst interval) distribution; if bursts occur at random but with a fixed absolute refractory period, their distribution should correspond to a shifted exponential, that is, $\text{IBI}(\Delta t) = \gamma e^{-\gamma(\Delta t - t_0)}$ for $t \geq t_0$ and 0 elsewhere, where t_0 is the duration of the absolute refractory period and γ is the mean rate for bursts. If, on the other hand, bursting cells are pacemakers, that is, if they regularly fire in bursts at a fixed interval, the IBI should be sharply peaked around γ . Figure 14 shows the average normalized IBI distribution for 37 cells (those with $p \geq 1.5$ and more than 200 IBIs). The logarithm of the distribution appears linear in the normalized (see Fig. 14 caption) time range of 40–160 msec and falls off at shorter intervals, consistent with a numerical model (thick curves) of Poisson-distributed bursts with a burst-related refractory period that we develop in the next paragraph in terms of a single neuron.

To emulate the data for a single neuron shown in the left column of Figure 15, we synthesize the following point process. Similar to the previous section, we generate “events” using a Poisson process (with $\lambda = 32$ Hz) with a Gaussian-distributed refractory period (of mean 16 msec and 7 msec SD; this distribution was truncated below zero and renormalized). Each event was then replaced with a burst of action potentials, that is, δ -functions, where the length of the burst in milliseconds was approximately Gaussian distributed (mean, 5.2 msec; SD, 1.1 msec) and the spikes within the burst were chosen with approximately Gaussian spacing (mean, 1.8 msec; SD, 0.5 msec). The mean rate λ and the Gaussian refractory distribution were chosen to fit the measured IBI distribution. The parameters of the Gaussian distribution for the length of the burst and the density within the burst were also chosen to fit the neuronal data. If this model is simplified by assuming that the spikes

within the bursts are generated by a Poisson process (similar to a model proposed by Smith and Smith, 1965), then the power spectrum would remain flat above 200 Hz, rather than gradually rising as seen at the bottom of Figure 15.

The right column of Figure 15 shows the resultant ISI and power spectrum, which are matched against similar functions for a bursting MT cell (Fig. 15, left column). What is surprising is that the synthetic data shows a peak in the power spectrum at about 31 Hz, *without* any underlying oscillations. How can this occur? A simple analytical model proves to be insightful.

We again appeal to the power function of an infinite train of shot-noise [where each individual shot is described by $h(t)$], with refractory period modeled with the renewal density $p(t)$ (Eq. 13). While before we assumed that individual spikes can best be described using a $\delta(t)$ function, we now model a burst by a boxcar of amplitude A and half-width L centered around the origin. We set L to the half-width of the typical burst and A to normalize the area of the boxcar to account for the number of spikes within the typical burst. The energy spectrum associated with such an event is given by the square of a *sinc* function, that is, by

$$S_h(\omega) = \frac{A^2 \sin^2(2\pi Lf)}{\pi^2 f^2} \quad (16)$$

The power spectrum of such Poisson events with a refractory period is

$$S_{\text{burst}}(f) = \lambda \frac{A^2 \sin^2(2\pi Lf)}{\pi^2 f^2} (1 - \sqrt{2\pi\lambda\sigma} e^{-2(\pi f\sigma)^2}), \quad (17)$$

for $f \neq 0$.

We superimposed $S_{\text{burst}}(f)$ onto the neuron's spectrum in Figure 15 (lower left) and found that both functions show a peak at the same frequency. The reason for the peak is the fact that S_{burst} is the product of $\text{sinc}(f)^2$, a decreasing function of f around the origin, and an monotonically increasing function, $1 - e^{-f^2}$. Figure 16 shows the estimated power spectrum $S'(f)$ as well as the associated best fit on the basis of Equation 17 for five burst cells. The analytical model does not account for variations in the burst width and occasional isolated spikes. Also, due to the use of the boxcar function to mimic bursts, we have no control over the fine structure of the spikes within the burst and therefore $S_{\text{burst}}(f)$ does not match well at high frequencies. What is important in this model is that the spectrum of this point process shows a peak, in the absence of any underlying oscillator model. A similar result may be obtained in this case by using IBIs drawn from a broad Gaussian distribution with a mean value close to 25 msec. More neuronal data would be required to distinguish between the appropriateness of these models.

To emphasize the fact that the presence of bursts—in combination with a refractory period—can lead to a peak in the power spectrum, we used all 210 trials at different values of c for one particular cell, *j001*, and replaced every occurrence of a burst by a single spike, located at the center of the burst (Fig. 17). The associated power spectrum changes dramatically in character, from the usual peaked one to a flat spectrum with a dip at low frequency, compatible with the notion that once bursting has been accounted for, what remains are Poisson-distributed events modulated by the presence of a refractory period. If bursts would tend to occur every 25 msec or so, then this procedure should have led to a spectrum with a large peak around 40 Hz. For our data, bursts account satisfactorily for the peaks in the power spectrum. This is also witnessed by the rate of burst occurrence λ , shown in the table in Figure 16, where

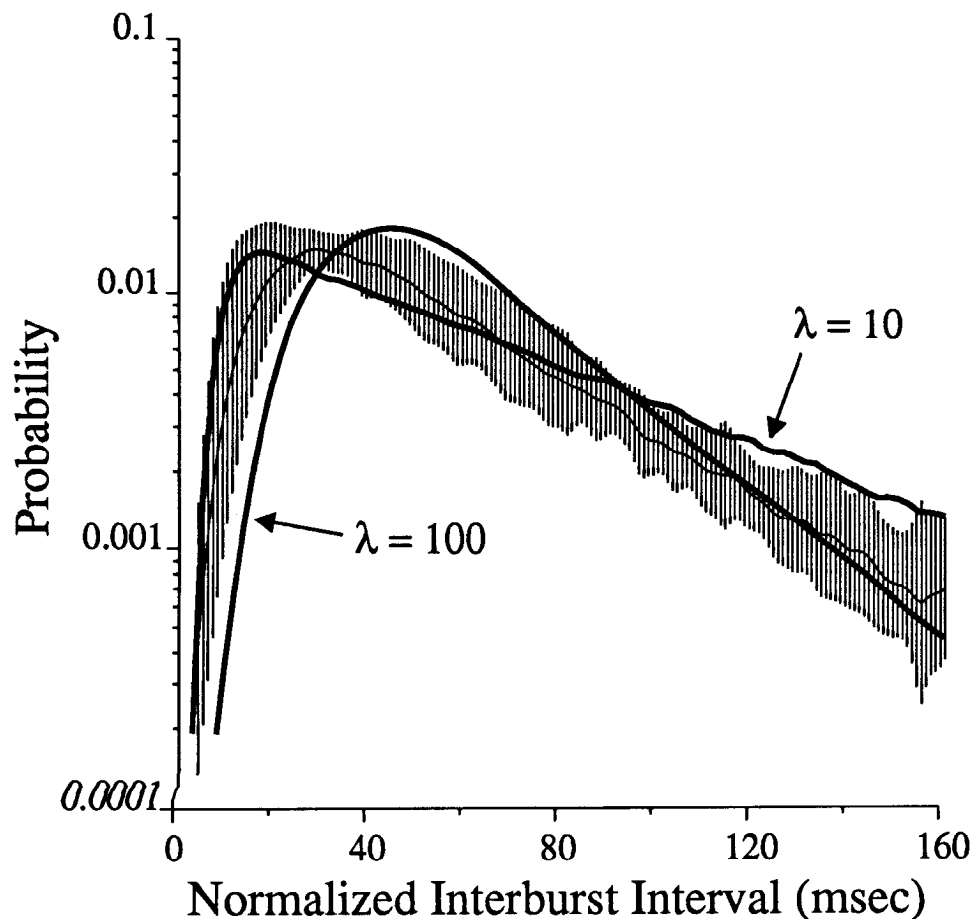


Figure 14. The average IBI distribution for 37 burst cells ($p \geq 1.5$; four cells with too little data were discarded) and a numerical model. Individual IBIs were computed from all trials at $c = 0.512$ ($c = 0.256$ for four cells) using the 8 msec criterion (see Results), normalized to have mean IBI equal to 60 msec (the population mean) by compressing or stretching the abscissa while preserving area, and smoothed with a Gaussian ($\sigma = 3.0$ msec) to eliminate isolated zero values. The average of logarithm of the IBI (thin line; error bars show SD) is roughly linear from 40–160 msec. Since the mean IBI varied from 20 to 100 msec over the 37 cells, the linear range of 40–160 msec in the plot corresponds to ranges from 13–53 msec to 65–265 msec in actual time. Two thick curves show IBI distributions from a numerical model of a Poisson burst-generating process with Gaussian-distributed burst-related refractory period ($\mu = 17$ msec, $\sigma = 8$ msec, as in Fig. 15; see Results). Model curves are shown for $\lambda = 10$ and 100 bursts/sec. By varying λ , the model accounts for the slope and approximate shape of the normalized IBIs while holding constant the parameters of the stochastic refractory period. The variance of the data is smallest near the intersections of the model curves, consistent with the notion that varying the model parameters induces little change in the IBI in these regions. The model somewhat overestimates the fraction of intervals in the linear range of 40–160 msec. Beyond 160 msec, values become undefined due to frequent zeros in the individual IBI histograms.

λ is distributed between 10 and 20 Hz, below the range where the peaks appear in the corresponding spectra.

Discussion

The intent of the research reported here is to study some aspects of the time structure of spike trains recorded in cortex of the behaving monkey on the basis of the power spectral density, an order-dependent measure, and the ISI histogram, an order-independent measure. Furthermore, we would like to assess whether knowing anything about the time course of the cell's discharge can lead to more accurate predictions concerning the stimulus or the monkey's response than simply counting spikes. In this investigation, we have confined our analysis to data from well-isolated single units.

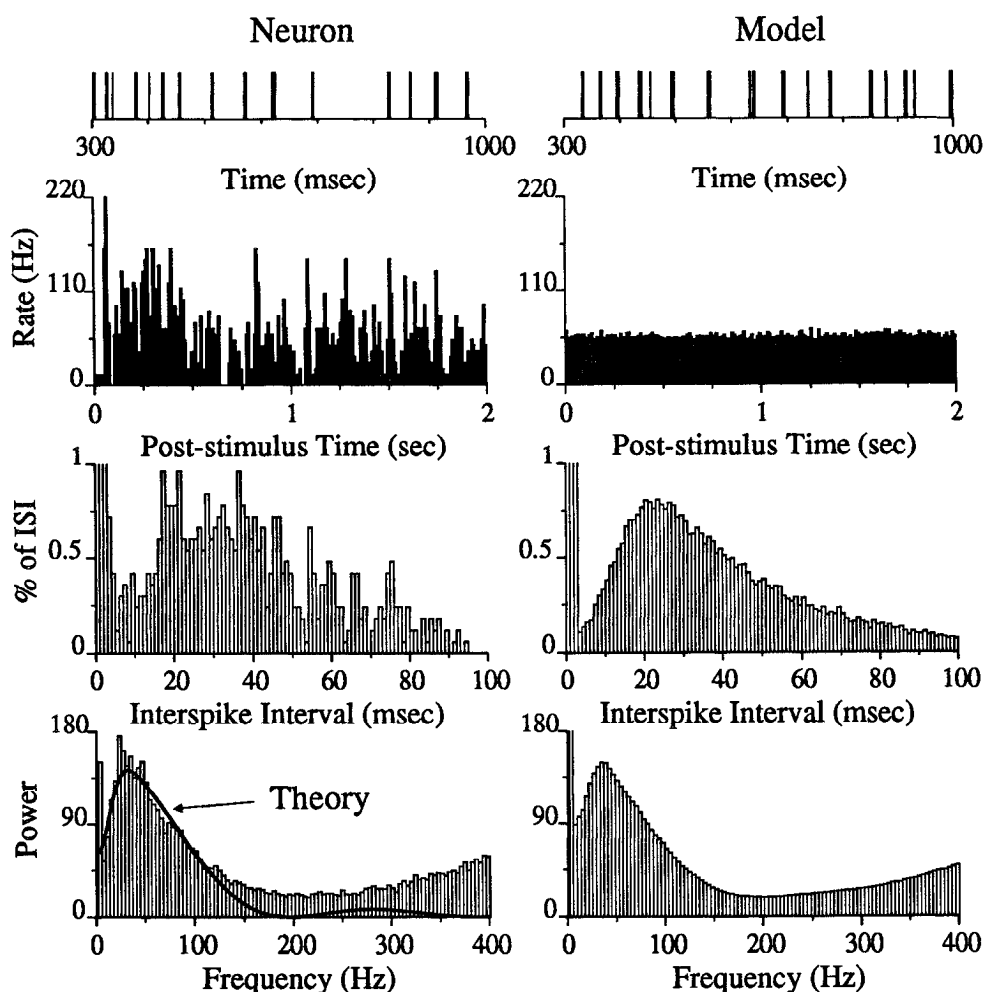
Random, nonbursting cells

We found that about one-third (71) of our MT cells can be adequately described by a Poisson process of mean spiking rate λ with a refractory period (modeled here by a Gaussian-distributed refractory period), in the sense that the experimentally determined power spectra $S'(f)$ and ISI distributions match the

analytically (and numerically) determined ones (Eq. 15, Fig. 11). In particular, the power spectra are flat, with a dip toward low temporal frequencies. Spikes from these cells do not occur in bursts; that is, they are almost always at least 4 msec (and usually much more) apart. We wish to point out that a Poisson process with refractory period is almost the simplest statistical description possible, with only two degrees of freedom, yet it appears to describe the measured discharge patterns relatively well. To our knowledge this is the first time that such a process is identified by its characteristic power spectrum.

In a related study (Softky and Koch, 1992, 1993), we computed the coefficient of variation C_v associated with the spike trains from the same data set of nonbursting MT cells (normalized for their nonstationary firing rates) as well as from V1 cells in the behaving monkey responding to bars and other textured stimuli (Knierim and Van Essen, 1992) and found values of $C_v \approx 1$, consistent with a Poisson process. We also analyzed the number of spikes occurring in different trials in response to the same stimulus and found that the variance in the number of spikes scales approximately as the average number of spikes to the 5/4 power (Softky and Koch, 1993). In a pure Poisson

Figure 15. Comparison of spike train statistics for burst cell *j001* to those from a simple numerical model. The statistics for the neuron (left column) were computed by averaging over 15 trials at $c = 0.256$. The spikes (top trace) are more clustered than random, as demonstrated by the excess (62%, truncated on plot) of intervals in the 1, 2, and 3 msec bins of the ISI. The PSTH is particularly noisy because spikes occur in bursts. The corresponding numerical model (right column) is an extension of that used for nonburst cells (see Fig. 11) in which each spike generated from an underlying Poisson process with refractory period is now replaced by a burst of spikes where the burst length and the temporal structure of spikes within the burst are chosen to match the data (see Results). Similar to the nonburst model, this is not intended to be a best fit to the data, but a demonstration that a process firing bursts randomly with a burst-related refractory period can account qualitatively for the location, size, and shape of the peak near 33 Hz and the dips near the origin and at higher frequencies. The solid curve superimposed on the neuron's spectrum (bottom left) corresponds to the analytical power spectrum given by Equation 17 with $\lambda = 20.4$ events/sec, $\sigma = 12.5$ msec, $A = 550$ spikes/sec, and $L = 2.55$ msec. This cell is one of the 10% of cells whose peak in the 40 Hz range is twice as tall as the dip at higher frequencies.



process, the variance should be equal to the mean, while for a fractal point process the variance can be larger than the mean (Teich, 1992; Usher et al., 1994). A number of studies have used this measure as indicative that the firing of cortical cells in striate and extrastriate monkey cortex are consistent with a description of spiking as a Poisson process (Tolhurst et al., 1983; Parker and Hawken, 1985; Vogels et al., 1989; Zohary et al., 1990; Snowden et al., 1992). Thus, at least for long spike trains in the trained monkey, the associated ISI and the power spectrum are compatible with the notion that the underlying point process can be described by an almost memoryless Poisson process with refractory period.

Bursting cells

More complex temporal dynamics are shown by the large fraction of cells (131 of 212) that frequently discharge in bursts, that is, two to four spikes within 2–8 msec or less (see the raster plots in the top row of Fig. 2). The fraction of the total number of spikes in a train that are less than 3.5 msec apart (our measure of "burstiness," B) ranges from an extreme value above 60% to 0. Unfortunately, we were not able to separate our 212 cells into two clearly segregated subpopulations using this or a related measure, since the distribution of cells varies continuously from strongly bursting to nonbursting (Fig. 3). Thus, any grouping of cells into "bursting" and "nonbursting" will be arbitrary to some extent. However, the amount of burstiness associated with in-

dividual cells remains relatively constant for all visual stimuli tested and, in particular, is independent of motion coherence (Figs. 5, 6, and our definition of P in Data analysis in Methods). Furthermore, we found no systematic relationship between burstiness and the onset or the duration of the experiment for the majority of cells. Finally, it should be remembered that due to the perceptually demanding nature of the experiment, the monkey had to be highly alert throughout each trial. These observations argue against the possibility that the bursting is related to the onset of drowsiness or sleep.

Intracellular current injection into cells in rodent slices of sensory-motor cortex has revealed three distinct types of neurons (McCormick et al., 1985; Connors and Gutnick, 1990; Agmon and Connors, 1992). The majority of these *in vitro* cells respond to the sustained current by a train of action potentials, which adapt within 50–100 msec to a more moderate discharge rate ("regular-spiking" cells). A second class of neurons, only infrequently recorded from, is capable of high discharge rates with little or no adaptation ("fast-spiking" cells). A third set of neurons respond to the depolarization by generating a short burst of two to four spikes, followed by a long hyperpolarization. This cycle of burst and hyperpolarization persists for as long as the current stimulus persists ("intrinsically bursting"). In slice tissue, the regular-spiking cells correspond to pyramidal neurons, fast-spiking cells to GABAergic nonspiny stellate cells, and the intrinsically bursting cells to layer V pyramidal cells (Agmon

Cell Name	λ (Hz)	σ (msec)	A (Hz)	L (msec)
w150	13.0	22.4	558	2.3
w143	14.5	16.5	526	2.1
w190	17.7	17.5	421	2.0
w200	13.4	13.5	426	2.0
w180	15.4	11.8	374	1.8

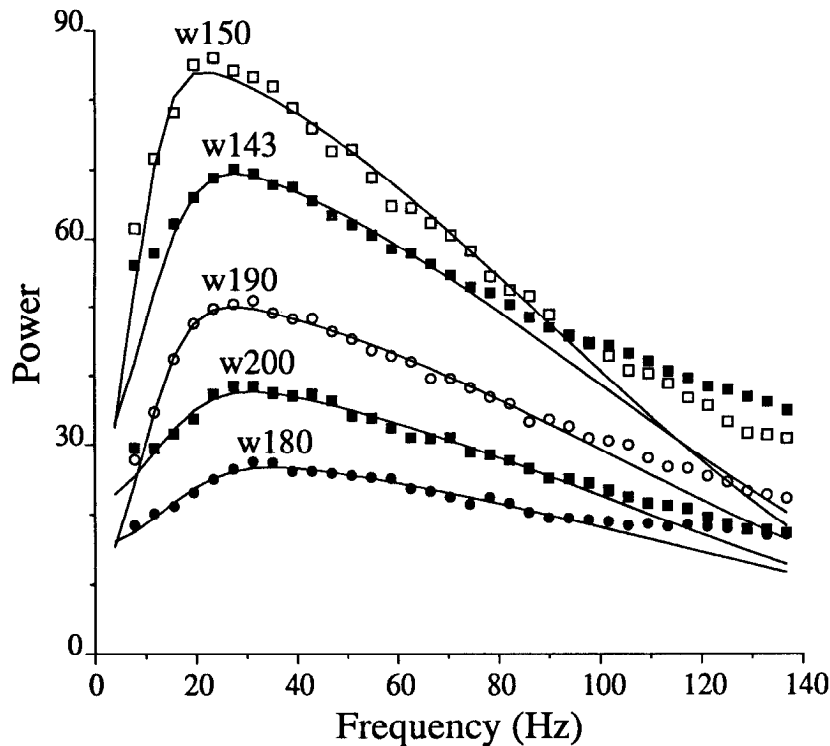


Figure 16. A comparison of power spectra, $S'(f)$ (data points), for five strongly bursting MT cells and their corresponding analytical curves, $S_{\text{burst}}(f)$ (solid lines), based on the expression for a shot-noise process with randomly (Poisson) distributed bursts (modeled as boxcar functions) with a burst-related refractory period (Eq. 17). The parameters of the fits are shown in the table: λ is the mean rate of the shot-noise, σ is the refractory period parameter, A is the height of the boxcar function (i.e., the spike rate within bursts), and L is the half-width of the boxcar function (i.e., the half-width of the burst). The analytical curves do not follow the spectrum at higher frequencies because they do not model individual actions potentials (see Results).

and Connors, 1992). Because only little is known about the distribution of these cell types in the monkey, we can at present only speculate to what extent our “bursting” cells in MT correspond to these layer V intrinsically bursting cells characterized in slice preparations. However, the evidence presented above certainly suggests that the propensity of our MT cells to respond with bursts is not dependent on the nature of the visual stimulus, but rather appears to be an intrinsic property of certain cells.

What is the statistical distribution of bursts? We converted spike trains of some bursting neurons into “bursting trains” by the simple rule that a burst was defined as the longest subtrain with no ISI greater than 8 msec (using values as low as 3 msec here made only very little difference in the result). The resulting IBI distributions (see Fig. 14 for average) are not readily compatible with a neuronal process that generates bursts at any fixed temporal interval. In most cases, the IBI distributions can be fitted assuming a Poisson distribution of bursts combined with a Gaussian-distributed burst refractory period (e.g., cell j001; see Figs. 14, 15). In some cases, the IBI has a long tail, arguing against a simple exponential decay. We suspect that occasional excesses of long IBIs is partly the reason that the Poisson IBI model (thick curves, Fig. 14) somewhat overestimates the fraction of IBIs in the range of 40–160 msec. The relatively short duration of the trials considered here does not allow a conclusive study of intervals that fall beyond the 160 msec (normalized) value in Figure 14. We never observed narrowly peaked IBIs,

suggesting that individual MT cells are not acting as pacemakers. This is partly, however, a question of linguistic convention, since any cell having an IBI distribution with a single peak (such as that shown in Fig. 14), might *in principle* be considered to “oscillate” at the inverse of this peak. We do not, however, find this to be a very helpful definition.

Cells with a peaked spectrum

About two-thirds of all our MT cells (131 of 212) have a single peak in their power spectra in the 40 Hz range (mean, 41 Hz; SD, 9 Hz; Figs. 2, top row; 15). Such a peak in the power spectrum is *not*, however, associated with ringing or oscillatory behavior in the Fourier transform of $S'(f)$, that is, the autocorrelation function $R'(t)$. Only in a single cell (j001) during very high levels of motion coherence did we ever observe an oscillatory response in $R'(t)$. This appears to be quite different from the study of Kreiter and Singer (1992), who report that a large fraction of MT cells in the awake monkey show oscillations.

What simple statistical model of neuronal firing can give rise to a peak in the power spectrum? We show that the power spectrum of a Poisson process with a Gaussian-distributed refractory period is monotonic increasing, leveling off toward a constant value at higher frequencies (Eq. 15). The spectrum associated with a boxcar-like burst is $[\sin(f)/f]^2$, a monotonic decreasing function around the origin. The power spectrum of

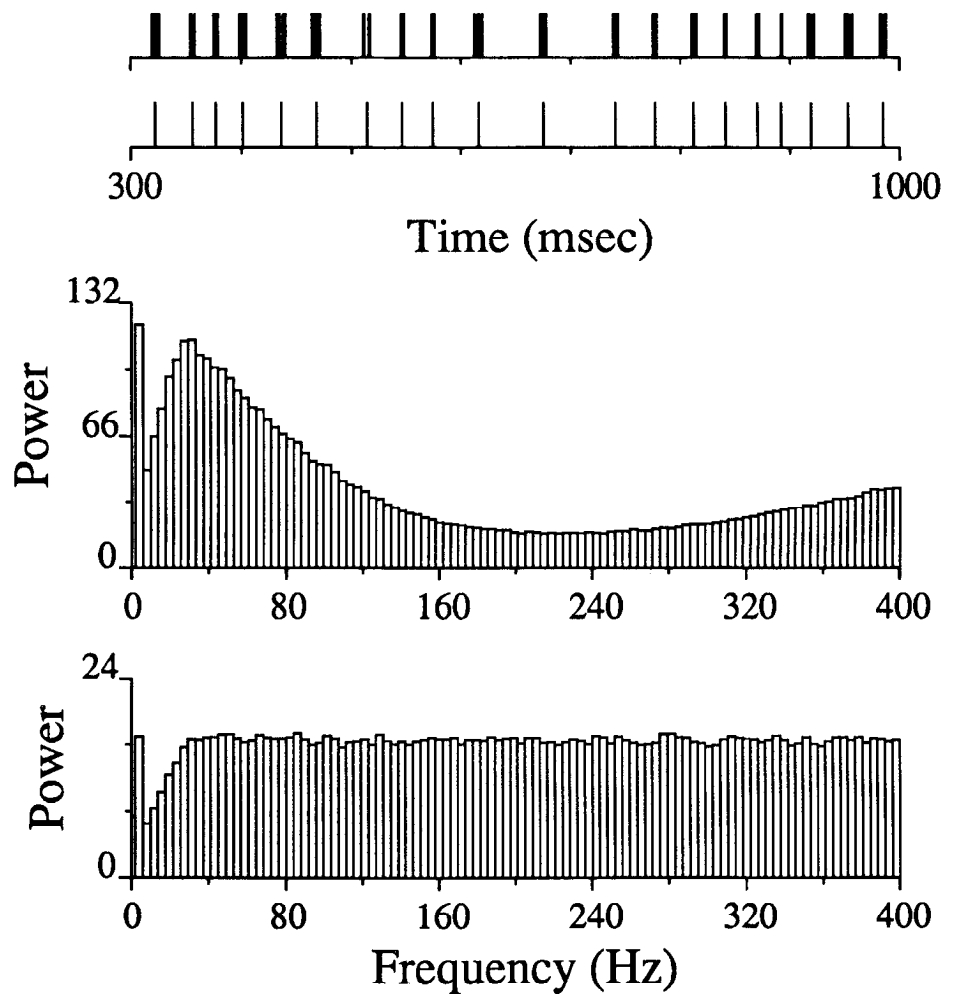


Figure 17. A demonstration that a peak in the power spectrum may be due to the presence of bursts, rather than regularity in their temporal alignment. When bursts from a neuron (*upper impulse plot*) are replaced by single spikes (*lower impulse plot*), the peak in the original power spectrum (*upper spectrum*) disappears (*lower spectrum*). A burst is defined as the longest subtrain of consecutive action potentials with no ISI greater than 8 msec (using 3 msec gives a very similar result). Each burst is considered to be an event at the mean occurrence time of all action potentials within that burst. This supports our notion that the bursts themselves are randomly placed (with a burst-related refractory period) and are *not* locked to a regular oscillatory pattern.

a process that randomly fires bursts followed by a refractory period is then given by the product of these two equations (Eq. 17). Given that one function is increasing with f and the second one decreasing, the product of the two will have a local maximum. In our case, if bursts are treated as boxcar functions that are about 4 msec wide, occur at a frequency of between 10 and 20 Hz, and are followed by a refractory period of between 10 and 25 msec (see table, Fig. 16), the peak in $S'(f)$ lies in the 20–50 Hz range (Fig. 15). These values were obtained from the distribution of the bursts themselves and can also be justified on biophysical grounds (Connors and Gutnick, 1990). We find it surprising to what extent simple analytical models can account for the observed ISI distributions and power spectra of bursting cells at frequencies less than 100 Hz. Our computer simulation of such a firing process, which differs from the analytical model by resolving the boxcar bursts into individual impulses and the renewal density into a stochastic refractory period, gives a better match of $S'(f)$ at higher frequencies.

If the occurrence of every burst in a spike train is replaced by a single action potential throughout the entire spike train (and isolated action potentials remain single spikes), the power spectrum $S'(f)$ totally changes its character (Fig. 17), from a spectrum with a peak to a flat spectrum with a dip at low frequencies, compatible with our notion that bursts themselves are distributed according to a Poisson distribution with a burst-related refractory period. If, for instance, the peak in the spec-

trum is due to periodically occurring bursts, our procedure should have revealed a spectrum with a clear peak, rather than the flat spectrum with a dip. We believe that our method of replacing bursts with “events” is a useful diagnostic tool for removing the confounding influence of bursts on the power spectrum, revealing the underlying dynamics.

Another way in which a neuronal “oscillator” model differs from our “random burst” model is in the distribution of IBIs; the former gives rise to an IBI distribution tightly clustered around the inverse of the oscillation period, while the latter model is associated with a decaying exponential IBI modified by a refractory period.

As discussed in the previous section (Bursting cells), our data are consistent with the random burst model (Fig. 14); however, the two models are difficult to distinguish when the oscillator model becomes less regular. We can show using computer-generated data what is expected in the case of the oscillator model. Figure 18 demonstrates the appearance of the ISI and spectrum $S'(f)$ in the case of a neuron which fires isolated spikes (top) or bursts (bottom) with an approximately Gaussian ISI or IBI. In the case of isolated spikes, the power spectrum remains flat with a dip related to the apparent refractory period induced by the Gaussian ISI for distributions with a broad range of SDs ($\sigma \geq 12$ msec). Once the SD becomes smaller (Fig. 18, top; $\sigma = 7$ msec), a prominent peak arises in the spectrum (see Fig. 18, upper right arrow) related to the inverse of the mean of the

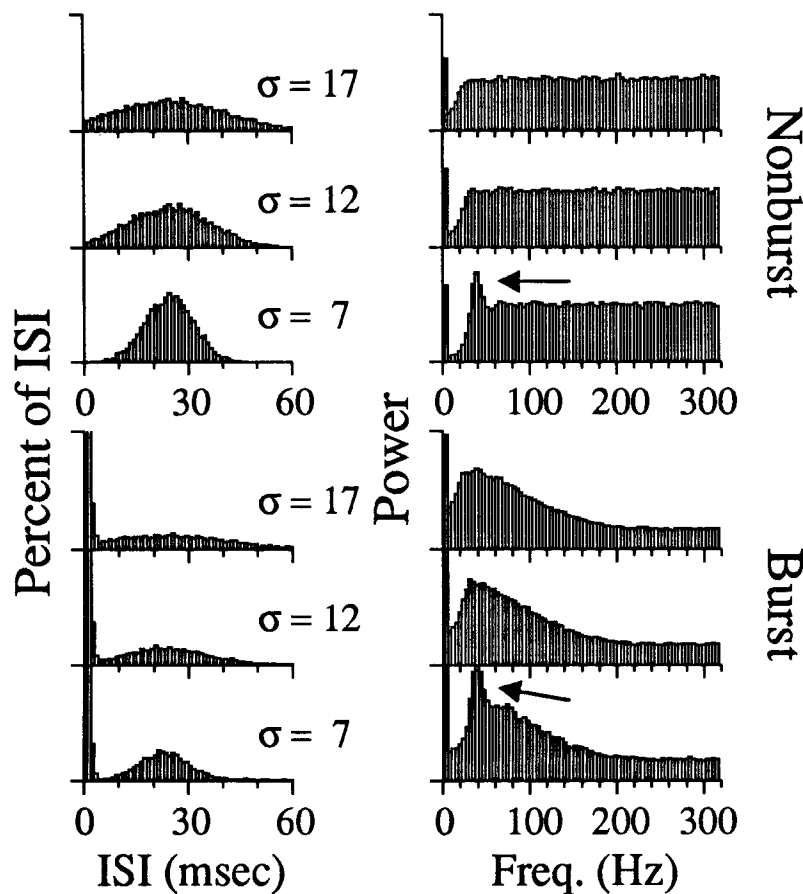


Figure 18. The difference between peaks in the power spectrum due to regular oscillation and peaks due to bursts. The *top six plots* show computer-generated data for the Gaussian ISI model with a mean of 25 msec. The *left column* shows a series of ISI histograms in which the SD, σ , of the Gaussian distribution is varied. For $\sigma > 12$ msec, there is no observable peak in the power spectra (*right column*), only a dip due to the apparent refractory effect imposed by the Gaussian. At $\sigma = 12$ msec, a peak is just beginning to form at 40 Hz (the inverse of the mean of the ISI), and as σ decreases further, the peak (shown by arrow for $\sigma = 7$ msec) rapidly increases in size. A similar progression occurs for the Gaussian IBI model (*bottom set of six plots*) in which spikes from the Gaussian ISI model are replaced by bursts. However, in the Gaussian IBI model, the spectrum has a broad bulge that sweeps up to a peak near 40 Hz due to the structure of the bursts (see Fig. 17), and this peak is present for all values of σ . A narrower peak appears above the broad peak due to bursts only when the variance of the Gaussian IBI becomes small. We therefore make a distinction between a peak in the power spectrum that is due to the local temporal structure of bursts and a peak that is due to a very narrow ISI or IBI distribution, which is a sign of regularity, that is, oscillation, in the timing of spikes or bursts. It is common to see sharp peaks of the type pointed to by the *arrows* when neurons respond to artifacts in a stimulus, such as the 60 Hz refresh rate of a cathode-ray tube.

Gaussian ISI. The bottom portion of Figure 18 shows a similar result, except that the isolated spikes are replaced by bursts (see the bimodal ISI distributions) and the peak due to small values of σ is superimposed on the peak due to bursts, explained in Poisson-distributed bursts in Results. The narrow, that is, more localized, peak should be associated with oscillations since it arises due to the regularity of the ISI and *not* due to the interaction of bursting with the refractory period. Localized peaks in the power spectrum, although not observed in our data, are associated with ringing in the cross-correlogram that has been reported in data from other laboratories.

Treating bursts as signaling events

Because of the possible special relevance of bursts to signaling in the brain (e.g., Koch and Crick, 1994), we evaluated to what extent bursts convey a different message from that conveyed by a collection of individual spikes. Following Cattaneo et al. (1981a,b) and Bonds (1992), we plotted the tuning curve as a function of motion coherence for three different measures of cell response (Fig. 8a): spikes/second, events/second—where an event is either a burst of spikes or an isolated spike—and spikes/event. Different from Bonds (1992), the average spikes/event (which includes individual spikes) does *not* vary with motion coherence c . Furthermore, if the event/second response is scaled up by the mean number of spikes per event, it closely follows

the spike rate tuning curve (Fig. 8b). Thus, events or bursts per se have the same overall dependency on c as does the spike rate. However, due to elimination of the variability in the number of spikes per event when using event count, the variance of the scaled-up curve, $\mu_M(c)$, is less than the variance associated with the mean spike rate, $\mu_S(c)$, in particular around low values of motion coherence.

Because of the reduced variability of this measure, we expect it to be a more reliable indicator of the direction of motion of the stimulus. This intuition is confirmed by our analysis. We repeated the original ROC analysis (based on signal detection theory) of Newsome et al. (1989a) but allowed more flexibility in defining the signal on which the ideal observer would operate. They assumed that an ideal observer (referred to as the “Stanford” observer) counts each action potential in the 2-sec-long spike train from an MT cell and uses this number as the basis for his analysis. We added a twist to this by weighting bursts differently from single spikes. While each isolated action potential contributed 1 toward the final sum, all spikes in a burst make a total contribution of α . Setting $\alpha = 1$ implies that the entire burst only contributes as much as a single, isolated action potential toward the final decision.

The result, as shown in Figure 9, is unambiguous. Setting $\alpha = 1$, corresponding to using the event rate as the neuronal signal, improves the neuronal threshold for most of the 41 strong burst

cells we analyzed here. The mean improvement was 7.5%, and in three cells the thresholds dropped by roughly a factor of 2. In other words, the “Caltech” ideal observer who distinguishes spikes from bursts can—for these cells—determine the correct direction of motion (using a statistical criterion) at a lower level of coherence c than the “Stanford” observer. $\alpha = 0.5$ and 2 gave smaller improvements (as did $\alpha = 0.75$ and 1.5; not shown). As α is made larger ($\alpha \rightarrow \infty$), bursts are more and more emphasized at the expense of single isolated spikes and the thresholds increase by 53%. The threshold also increases when bursts are weighted by the square of the number of spikes in each burst ($\beta = 2$). Weighting events by the square root of the number of spikes per event ($\beta = 0.5$), on the other hand, decreases thresholds (Fig. 9), since it decreases the variability contributed to the final signal.

From the point of view of our fictitious pair of ideal observers, the Caltech observer does better than his Stanford counterpart by replacing each occurrence of a burst of spikes by a single spike. In this sense it can be said that a crude measure of the temporal organization of spike trains does better in terms of signal detection theory than a simple spike count. This is not to say, however, that more sophisticated measures of temporal organization, possibly taking account of the simultaneous activity of many neurons, cannot do better yet (e.g., Aertsen et al., 1989; Richmond and Optican, 1992; Singer, 1994).

We do not know at this point the code that neurons postsynaptic to MT cells use to decide whether the stimuli move in one or the other direction. The fact that the neuronal threshold of many cells is frequently lower than the psychophysical threshold of the entire animal (Newsome et al., 1989a) requires an explanation as to why the animal does not do better than it does (invoking correlated activity among cells and population coding; Britten et al., 1992). Our results point to an additional explanation: if bursts are substantially more efficient in elevating postsynaptic firing rate than isolated spikes, corresponding to $\alpha \gg 1$, thresholds would increase and averaging over many cells would be required in order to mimic the psychophysical thresholds. It is important at some point that this question be resolved experimentally, possibly using a combination of *in vivo* slice techniques with behavioral studies.

Functional considerations

What is the function of bursts? Why should cortex have two types of long-range projection cells, one signaling isolated spikes and the other responding frequently with bursts of spikes? It has been argued (Koch and Crick, 1994) that bursting neurons are much more efficient at accumulating calcium in their axonal terminals than cells that fire isolated spikes (i.e., four spikes within a 10 msec interval cause a much larger increase in intracellular calcium at the end of the last spike than four spikes within a 40 msec interval). Because intracellular calcium accumulation in the presynaptic terminal is thought to be mainly responsible for various forms of short-term potentiation (in particular, facilitation and augmentation; Magleby, 1987), it may well be that the primary function of bursting neurons is to induce this non-Hebbian (i.e., nonassociative) type of synaptic plasticity at its postsynaptic targets outside of the cortical system. In essence, the burst of spikes could turn on short-term memory, which would then decay over several seconds (see also Crick, 1984). One might then expect there to be a relationship between bursting and short-term learning.

It is important to know whether our “bursting” cells corre-

spond to the “intrinsically bursting” cells identified by intracellular current injection. The latter cells appear to be confined (at least in rat and guinea pig slice) to layer V (Agmon and Connors, 1992). In rat area 17, these cells have been shown to project outside cortex, in this case to the ipsilateral superior colliculus, while the remaining pyramidal cells in layer V project to the contralateral striate cortex (Kasper et al., 1991). Recent *in vivo* recordings in awake cat motor cortex have revealed that cells at or below a depth of 800 μm (corresponding to layer 5) show strong bursting activity (Baranyi et al., 1993). It is not known to what extent such cells exist or are localized to particular layers in primate cortex.

Correlation to behavior

Finally, we return to a question that provided primary motivation for this study. Is the animal’s perception of the stimulus, as evidenced by performance, influenced by temporal structure in the spike trains, particularly with respect to the peak in the power spectrum near 40 Hz? We correlated the presence and strength of the peak in the spectrum to both the stimulus and the behavior of the monkey. As seen in Figures 5 and 6, we found no significant correlation between the fraction of dots moving in one or the other direction and P , the measure of the peak in the power spectrum, for most cells. We repeated this measurement using the integral of power in the 40 Hz band of $S'(f)$ with similar results. Furthermore, if the monkey is not forced to respond to the stimulus or even in the absence of the motion stimulus, the basic propensity of a cell to show this peak remains. This is related, of course, to our earlier result that bursting in these cells does not depend on stimulus conditions.

We find a similar lack of correlation between P and the various measures of behavior used in a previous comparison of neuronal responses and psychophysical performance (Newsome et al., 1989a; Britten et al., 1992). For instance, we tested for significant changes in the distribution of P when the monkey made correct versus incorrect decisions at near-threshold coherence levels and when the monkey made preferred versus null guesses for zero coherence motion, but we found no significant (paired t test, $p > 0.05$) correlations.

Figure 7 shows another attempt at studying the relationship between the peak in the spectrum and the behavior of the monkey. As is evident, no correlation exists between P and the level of the neuronal threshold, c_{cell} , that is, the fraction of dots moving in the cells preferred direction at which the cell can “decide” the correct direction of motion (using an ROC criterion; Newsome et al., 1989a; Britten et al., 1992). A similar lack of correlation exists between P and the ratio of single-cell thresholds to the threshold of the animal ($c_{\text{cell}}/c_{\text{system}}$) and between P and the decision related probability of each neuron (not shown). Thus, for our stimulus conditions, the presence or strength of a peak in the power spectrum of well-isolated units does not tell us anything about the behavior of these animals.

When analyzing our data set, it should be kept in mind that the three monkeys from which the MT cells were recorded were extensively trained using operant conditioning techniques. In order to perform correctly the motion discrimination task at threshold levels, up to 6 months of training was required (Britten et al., 1992). We analyzed in a preliminary manner data from MT cells from one naive monkey who was only trained to fixate (E. Zohary and W. Bair, personal communication), and we found no significant difference in the distribution of burst and nonburst cells or in the shape of the associated power spectra. However,

it is possible that more subtle differences in the temporal fine-structure exists between cells in the naive and in the trained animal.

Although we cannot say how well our results will generalize to other cortical areas or other behavioral tasks, we believe that the approach taken here—correlating temporal structure in spike trains to the simultaneous behavior of an awake animal—is a necessary step in establishing the role played by temporal firing patterns in the animal's perception of visual stimulation.

Appendix

Here we assume that a cell fires N events during the course of a fixed duration stimulus and that the i th event is composed of X_i spikes, where N and X_i are independent random variables, $1 < i < N$. All of the X_i are mutually independent and drawn from a common distribution, and we use X to refer to a variable drawn from this distribution with no reference to a particular event. The total number of spikes fired during the trial is

$$S_N = X_1 + X_2 + \dots + X_N, \quad (18)$$

and we will prove for the random variable S_N that

$$E(S_N) = E(N)E(X), \quad (19)$$

$$\text{VAR}(S_N) = E(N)\text{VAR}(X) + \text{VAR}(N)E^2(X). \quad (20)$$

We use E for expectation, VAR for variance, and P for probability.

The probability that k spikes are fired during a trial is

$$P\{S_N = k\} = \sum_{n=0}^{\infty} P\{S_N = k \mid N = n\}P\{N = n\}. \quad (21)$$

Using this and the definition of expectation,

$$E(S_N) = \sum_{k=0}^{\infty} kP\{S_N = k\} \quad (22)$$

$$= \sum_{k=0}^{\infty} k \sum_{n=0}^{\infty} P\{S_N = k \mid N = n\}P\{N = n\} \quad (23)$$

$$= \sum_{n=0}^{\infty} P\{N = n\} \sum_{k=0}^{\infty} kP\{S_N = k \mid N = n\} \quad (24)$$

$$= \sum_{n=0}^{\infty} P\{N = n\}E(S_N \mid N = n). \quad (25)$$

Evaluating the conditional expectation in the previous line, we get

$$E(S_N \mid N = n) = E\left(\sum_{i=1}^n X_i\right) = nE(X), \quad (26)$$

which, when substituted into Equation 25, completes the proof of Equation 19:

$$E(S_N) = E(X) \sum_{n=0}^{\infty} nP\{N = n\} = E(X)E(N). \quad (27)$$

To compute the variance of the number of spikes, we first compute $E(S_N^2)$. Reasoning as in Equations 22–25, but now for S_N^2 rather than S_N , we get

$$E(S_N^2) = \sum_{n=0}^{\infty} P\{N = n\}E(S_N^2 \mid N = n). \quad (28)$$

Evaluating the conditional expectation from the previous line, we get

$$\begin{aligned} E(S_N^2 \mid N = n) &= E\left(\sum_{i=1}^n X_i\right)^2 \\ &= E\left(\sum_{i=1}^n X_i^2 + \sum_{i=1}^n \sum_{j \neq i} X_i X_j\right) \\ &= nE(X^2) + (n^2 - n)E^2(X), \end{aligned}$$

which on substitution into eq. 28 yields

$$\begin{aligned} E(S_N^2) &= E(X^2) \sum_{n=0}^{\infty} nP\{N = n\} + E^2(X) \sum_{n=0}^{\infty} n^2P\{N = n\} \\ &\quad - E^2(X) \sum_{n=0}^{\infty} nP\{N = n\} \\ &= E(X^2)E(N) + E^2(X)E(N^2) - E^2(X)E(N) \\ &= E(N)\text{VAR}(X) + E^2(X)E(N^2). \end{aligned}$$

Using the expressions for $E(S_N)$ and $E(S_N^2)$, the variance of the number of spikes is

$$\begin{aligned} \text{VAR}(S_N) &= E(S_N^2) - E^2(S_N) \\ &= E(N)\text{VAR}(X) + E^2(X)E(N^2) - E^2(X)E^2(N) \\ &= E(N)\text{VAR}(X) + E^2(X)\text{VAR}(N), \end{aligned}$$

which completes the verification of Equation 20.

In Equation 6, $E(N)$ and $\text{VAR}(N)$ are replaced by the mean (which is equal to the variance for a Poisson distribution) number of events, $\mu_N(c)$, and the mean and variance of the number of spikes per event are μ_X and $\sigma^2(X)$.

References

- Abeles M (1982) Quantification, smoothing, and confidence limits for single-units' histograms. *J Neurosci Methods* 5:317–325.
- Abeles M (1990) *Corticonics*. Cambridge: Cambridge UP.
- Agmon A, Connors BW (1992) Correlation between intrinsic firing patterns and thalamocortical synaptic responses of neurons in mouse barrel cortex. *J Neurosci* 12:319–329.
- Aertsen AMHJ, Gerstein GL, Habib MK, Palm G (1989) Dynamics of neuronal firing correlation: modulation of "effective connectivity." *J Neurophysiol* 61:900–917.
- Bair W, Koch C, Newsome W, Britten K (1992) Power spectrum analysis of MT neurons from awake monkey. *Soc Neurosci Abstr* 18:12.
- Bair W, Koch C, Newsome W, Britten K (1993) Temporal structure of spike trains from MT neurons in the awake monkey. In: *Computation and neural systems*, Chap 75 (Eeckman FH, Bower JM, eds). Norwell, MA: Kluwer.
- Baranyi A, Szenté MB, Woody CD (1993) Electrophysiological characterization of different types of neurons recorded *in vivo* in the motor cortex of the cat. I. Patterns of firing activity and synaptic responses. *J Neurophysiol* 69:1850–1864.
- Barlow HB (1972) Single units and sensation: a neuron doctrine for perceptual psychology? *Perception* 1:371–394.
- Barlow HB, Kaushal TP, Hawken M, Parker AJ (1987) Human contrast discrimination and the threshold of cortical-neurons. *J Opt Soc Am* 4:2366–2371.
- Bialek W, Rieke F, Vansteveninck RRD, Warland D (1991) Reading a neural code. *Science* 252:1854–1857.
- Bonds AB (1992) Dual inhibitory mechanisms for definition of receptive field characteristics in cat striate cortex. In: *Advances in neural information processing systems*, Vol 4 (Moody JE, Hanson SJ, Lippmann RP, eds), pp 75–82. San Mateo, CA: Kaufmann.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.

- Cattaneo A, Maffei L, Morrone C (1981a) Two firing patterns in the discharge of complex cells encoding different attributes of the visual stimulus. *Exp Brain Res* 43:115–118.
- Cattaneo A, Maffei L, Morrone C (1981b) Patterns in the discharge of simple and complex visual cortical cells. *Proc R Soc Lond [Biol]* 212:279–297.
- Champeney DC (1973) Fourier transforms and their physical applications. New York: Academic.
- Chung SH, Raymond SA, Lettvin JY (1970) Multiple meaning in single visual units. *Brain Behav Evol* 3:72–101.
- Connors BW, Gutnick MJ (1990) Intrinsic firing patterns of diverse neocortical neurons. *Trends Neurosci* 13:99–104.
- Crick F (1984) Function of the thalamus reticular complex: the searchlight hypothesis. *Proc Natl Acad Sci USA* 81:4586–4590.
- Crick F, Koch C (1990) Towards a neurobiological theory of consciousness. *Semin Neurosci* 2:263–275.
- Crick F, Koch C (1992) The problem of consciousness. *Sci Am* 267:152–159.
- Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: a mechanism of feature linking in the visual cortex? *Biol Cybern* 60:121–130.
- Eskandar EN, Richmond BJ, Optican LM (1992) Role of inferior temporal neurons in visual memory: I. Temporal encoding of information about visual images, recalled images and behavioral context. *J Neurophysiol* 68:1277–1295.
- Feller W (1968) An introduction to probability theory and its applications, Vol 1, 3d ed. New York: Wiley.
- Ghose GM, Freeman RD (1992) Oscillatory discharge in the visual system: does it have a functional role? *J Neurophysiol* 68:1558–1574.
- Gray CM, Singer W (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc Natl Acad Sci USA* 86:1698–1702.
- Gray CM, König P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338:334–337.
- Harris FJ (1978) On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc IEEE* 66:51–84.
- Kasper E, Larkman A, Blakemore C, Judge S (1991) Physiology and morphology of identified projection neurons in rat visual cortex studied *in vitro*. *Soc Neurosci Abstr* 17:114.
- Knierim J, Van Essen D (1992) Neuronal responses to static textural patterns in area V1 of the alert macaque monkey. *J Neurophysiol* 67:961–980.
- Koch C, Crick F (1994) Some further ideas about the neuronal basis of visual awareness. In: *Large scale neuronal theories of the brain* (Koch C, Davis J, eds), in press. Cambridge, MA: MIT Press.
- Kreiter AK, Singer W (1992) Oscillatory neuronal responses in the visual cortex of the awake macaque monkey. *Eur J Neurosci* 4:369–375.
- Legény CR, Salzman M (1985) Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons. *J Neurophysiol* 4:926–939.
- Lettvin JP, Maturana HR, McCulloch WS, Pitts WH (1959) What the frog's eye tells the frog's brain. *Proc Inst Rad Eng* 47:1940–1951.
- Magleby KL (1987) Short-term changes in synaptic efficacy. In: *Synaptic function* (Edelman GM, Gall WE, Cowan WM, eds), pp 21–56. New York: Wiley.
- Maunsell JHR, Van Essen D (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed and orientation. *J Neurophysiol* 49:1127–1147.
- McCormick DA, Connors BW, Lighthall JA, Prince DA (1985) Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J Neurophysiol* 54:782–806.
- Newsome WT, Pare EB (1988) A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J Neurosci* 8:2201–2211.
- Newsome WT, Britten KH, Movshon JA (1989a) Neuronal correlates of a perceptual decision. *Nature* 341:52–54.
- Newsome WT, Britten KH, Movshon JA, Shadlen M (1989b) Single neurons and the perception of visual motion. In: *Proceedings of the retinal research foundation, Vol 2, Neural mechanisms of visual perception* (Lam DM-K, Gilbert CD, eds), pp 171–198. The Woodlands, TX: Portfolio.
- Optican LM, Richmond BJ (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J Neurophysiol* 57:162–178.
- Parker A, Hawken M (1985) Capabilities of monkey cortical cells in spatial-resolution tasks. *J Opt Soc Am* 2:1101–1114.
- Perkel DH, Gerstein GL, Moore GP (1967) Neuronal spike trains and stochastic point processes. I. The single spike train. *Biophys J* 7:391–418.
- Poggio GF, Viernstein LJ (1964) Time series analysis of impulse sequences of thalamic somatic sensory neurons. *J Neurophysiol* 27:517–545.
- Press HP, Flannery BP, Teukolsky SA, Vetterling WT (1988) Numerical recipes in C, the art of scientific computing. Cambridge: Cambridge UP.
- Richmond BJ, Optican LM (1992) The structure and interpretation of neuronal codes in the visual system. In: *Neural networks for perception* (Wechsler H, ed), pp 104–119. New York: Academic.
- Robinson DA (1963) A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE Trans Biomed Eng* 10:137–145.
- Salzman CD, Murasugi CM, Britten KH, Newsome WT (1992) Microstimulation in visual area MT: effects on direction discrimination performance. *J Neurosci* 12:2331–2355.
- Schild D, Schultens HA (1986) The Fourier transform of a peristimulus time histogram can lead to erroneous results. *Brain Res* 369:353–355.
- Singer W (1994) Putative functions of temporal correlations in neocortical processing. In: *Large scale neuronal theories of the brain* (Koch C, Davis J, eds), in press. Cambridge, MA: MIT Press.
- Smith DR, Smith DK (1965) A statistical analysis of the continual activity of single cortical neurones in the cat unanaesthetized isolated forebrain. *Biophys J* 5:47–74.
- Snowden RJ, Treue S, Andersen RA (1992) The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Exp Brain Res* 88:389–400.
- Softky WR, Koch C (1992) Cortical cells should fire regularly, but do not. *Neural Comput* 4:643–646.
- Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13:334–350.
- Streher BL, Lestienne R (1986) Evidence on precise time-coded symbols and memory of patterns in monkey cortical neuronal spike trains. *Proc Natl Acad Sci USA* 83:9812–9816.
- Teich MC (1992) Fractal neuronal firing. In: *Single neuron computation, Neural nets: foundations to applications* (McKenna T, Davis J, Zornetzer S, eds), pp 589–625. Boston: Academic.
- Tolhurst D, Movshon J, Dean A (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23:775–785.
- Usher M, Stemmler M, Koch C, Olami Z (1994) Network amplification of local fluctuations causes high spike rate variability, fractal firing patterns and oscillatory local field potentials. *Neural Comput*, in press.
- Vogels R, Orban GA (1990) How well do response changes of striate neurons signal differences in orientation—a study in the discriminating monkey. *J Neurosci* 10:3543–3558.
- Vogels R, Spileers W, Orban GA (1989) The response variability of striate cortical neurons in the behaving monkey. *Exp Brain Res* 77:432–436.
- von der Malsburg C (1981) The correlation theory of brain function. MPI for Biophysical Chemistry Göttingen, Department of Neurobiology, internal report 81–2.
- Werner G, Mountcastle VB (1963) The variability of central neural activity in a sensory system and its implications for the central reflection of sensory events. *J Neurophysiol* 26:958–977.
- Zeki SM (1974) Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *J Physiol (Lond)* 236:549–573.
- Zipser D, Kehoe B, Littlewort G, Fuster J (1993) A spiking network model of short-term active memory. *J Neurosci* 13:3406–3420.
- Zohary E, Hillman P, Hochstein S (1990) Time course of perceptual discrimination and single neuron reliability. *Biol Cybern* 62:475–486.
- Zohary E, Shadlen MN, Newsome W (1992) Correlated activity of neurons in area MT. *Soc Neurosci Abstr* 18:110.1