Interspike Intervals, Receptive Fields, and Information Encoding in Primary Visual Cortex

Daniel S. Reich,^{1,2} Ferenc Mechler,² Keith P. Purpura,² and Jonathan D. Victor²

¹Laboratory of Biophysics, The Rockefeller University, New York, New York 10021, and ²Department of Neurology and Neuroscience, Weill Medical College of Cornell University, New York, New York 10021

In the primate primary visual cortex (V1), the significance of individual action potentials has been difficult to determine, particularly in light of the considerable trial-to-trial variability of responses to visual stimuli. We show here that the information conveyed by an action potential depends on the duration of the immediately preceding interspike interval (ISI). The interspike intervals can be grouped into several different classes on the basis of reproducible features in the interspike interval histograms. Spikes in different classes bear different relationships to the visual stimulus, both qualitatively (in terms of the average stimulus preceding each spike) and quantitatively (in terms of the amount of information encoded per spike and per second). Spikes preceded by very short intervals (3 msec or less) convey

information most efficiently and contribute disproportionately to the overall receptive-field properties of the neuron. Overall, V1 neurons can transmit between 5 and 30 bits of information per second in response to rapidly varying, pseudorandom stimuli, with an efficiency of ~25%. Although some (but not all) of our results would be expected from neurons that use a firing-rate code to transmit information, the evidence suggests that visual neurons are well equipped to decode stimulus-related information on the basis of relative spike timing and ISI duration.

Key words: spike train; primary visual cortex; V1; information theory; white noise; receptive field; synaptic depression; synaptic facilitation; temporal coding; m-sequence; interspike interval

Cortical sensory neurons have high intrinsic temporal precision (Mainen and Sejnowski, 1995; Nowak et al., 1997) and can encode information on the scales of milliseconds and tens of milliseconds (Buračas and Albright, 1999). Three questions arise immediately. (1) What kinds of stimuli are encoded on the different time scales in a neuron's response? (2) How much information is encoded on each time scale? (3) How might this information be decoded by relatively simple components of neurons and neural circuits?

Here, we answer the first two questions experimentally by measuring the responses of neurons in the primary visual cortex (V1) of macaque monkeys to rapidly varying, pseudorandom ("m-sequence") stimuli. The interspike intervals (ISIs) of spike trains fired by these neurons fall into three subsets, distinguished on the basis of ISI duration, in a stereotyped manner across neurons. We use a reverse-correlation procedure to generate receptive field (RF) maps from the full responses as well as from response subsets that only contain spikes that follow ISIs of particular durations. Finally, we use information theory to quantify the rate and efficiency with which full responses and response subsets convey messages about the visual stimulus.

Our results indicate that spikes in different ISI subsets are fired in response to different visual stimuli. In particular, spikes preceded by ISIs <3 msec, which occur during periods of very high firing rate, tend to be evoked by stimuli that have several subregions of opposite contrast covering the neuron's receptive field. Each of these spikes also tends to convey more stimulus-related

information than the average spike. On the other hand, spikes preceded by ISIs >38 msec are often fired in response to spatially uniform stimuli that reverse contrast over time.

The third question, concerning the ways in which these different messages are decoded, is not addressed directly by our experiments. We note at the outset that this question is conceptually independent from another much-debated question in cortical physiology: whether cortical neurons encode information through a rate code or a temporal code. In fact, both types of code can generate receptive-field maps and information rates similar to what is described here. However, the existence of stereotyped ISI durations in V1 (described in this paper), together with recently described synaptic and dendritic machinery (such as depression, facilitation, and coincidence detection) that can selectively increase or decrease the importance of particular spikes in shaping a postsynaptic response, suggests that real-time decoding of neuronal signals may rely on known biophysical mechanisms specifically sensitive to ISI duration.

MATERIALS AND METHODS

Recording and stimuli. We recorded the responses of single V1 neurons in opiate-anesthetized macaque monkeys (Reich et al., 1998; Victor and Purpura, 1998). All experimental procedures complied with the guidelines of the National Eye Institute and our institution. We measured spike times to the nearest 0.1 msec for 135 neurons, and to the nearest 3.7 msec (one frame of the visual display) for 36 neurons. We include in our analysis only the 99 neurons [32 simple, 60 complex, and 7 unclassified (Skottun et al., 1991)] that responded with firing rates higher than 3 spikes/sec and that had significantly modulated RF maps (see below).

Our stimuli look to human observers like random and rapidly flickering checkerboards. In fact, however, the stimuli are highly structured. They consist of a grid in which the temporal sequence of the luminance levels (0 or 300 cd/m²) in each of 249 pixels (typically 16 × 16 arc-min) is determined by a pseudorandom, binary m-sequence (Sutter, 1992; Victor, 1992; Reid et al., 1997). The same 4095 step m-sequence is used in each pixel, but the starting position in the sequence is different. Because

Received Sept. 27, 1999; revised Dec. 9, 1999; accepted Dec. 17, 1999.

This work was supported by National Institutes of Health Grants GM07739 and EY07138 (D.S.R.), NS36699 (K.P.), and EY9314 (J.D.V.). We thank Bruce Knight, David Reich, Jason Eisner, Steve Kalik, and Anne-Marie Canel.

Correspondence should be addressed to Daniel Reich, The Rockefeller University, 1230 York Avenue, Box 200, New York, NY 10021. E-mail: reichd@rockefeller.edu.

Copyright © 2000 Society for Neuroscience 0270-6474/00/201964-11\$15.00/0

the minimum offset between any pair of pixels is 237 msec, and because an m-sequence is uncorrelated with temporal shifts of itself, there is little danger that the same stimulus sequence would simultaneously affect different parts of the receptive field. The sequence is advanced simultaneously in all pixels at a stimulus frame rate of 67.58 Hz, so that the luminance of each stimulus pixel can potentially change once every four frames of the 270.3 Hz visual display. This stimulus frame rate has been shown to evoke good receptive-field maps from cortical neurons (Reid et al., 1997). The entire m-sequence is repeated 8–16 times, as is its contrast-inverse, which is presented to eliminate spurious effects of residual correlations in the stimulus on our receptive field maps (Sutter, 1992).

Receptive field maps. Cross-correlation of an evoked spike train with the m-sequence stimulus—a process also known as “spike-triggered averaging”—yields a detailed map of the neuron’s spatiotemporal receptive field (see Fig. 4). This map essentially represents the average stimulus preceding each spike, and it is rendered as a series of contour plots depicting spatial snapshots that are sequential in time. Each map shows the average change in contrast in each of the stimulus pixels that significantly modulated the response for at least two consecutive 3.7 msec time bins ($p < 0.01$ in each bin; the estimated impulse response was at least 2.6 SEs from 0, where standard errors were determined empirically from multiple trials), as well as in surrounding pixels. We smooth the contour maps by cubic spline interpolation.

To the degree that the neuron is a linear system, the derived RF map can also be considered to depict its “spatiotemporal impulse response”—that is, its average response to an incremental flash of light at each spatial position (Victor, 1992). Of course, V1 neurons are not linear systems, and there are often higher-order components of the RF map that make the impulse response interpretation imprecise. In this situation, the RF maps are the linear functions that best fit the full response, and we retain the spirit of the impulse–response interpretation in the normalization of our RF maps. Contour heights represent, at each time frame, the change in firing rate induced by a luminance step in a particular stimulus pixel averaged over the displayed time window of 14.8 msec. To obtain absolute firing rates for the full response, the RF map of which is denoted by f in Equations 1 and 2 in Results, add these values to the mean firing rates given in the legend of Figure 4. To obtain absolute firing rates for the subset responses (see Results), multiply the RF map values by the fraction of spikes in the appropriate subset and add the product to the subset’s mean firing rate.

Information. To measure the information contained in m-sequence responses, we use a method modified from Strong and colleagues (de Ruyter van Steveninck et al., 1997; Strong et al., 1998). Spike trains are divided into time bins, each of which may be occupied by zero, one, or more than one spike. The possible spike counts in each bin can be thought of as letters in the neuron’s response alphabet, with several letters in a row constituting a word. Each word has a characteristic probability, possibly stimulus dependent, of being “spoken” by the neuron.

To measure the full information in a spike train, we would need to use a limitless sample of infinitely long words containing infinitesimally short letters. This is clearly impossible, so we choose our word and letter lengths according to physiological criteria, taking into account factors such as integration time and temporal precision, and also according to the amount of available data, which limits the accuracy with which the word probabilities can be estimated. We use 14.8 msec words [the stimulus frame time, but also similar to the time constant of cortical neurons (Ogawa et al., 1981; Shadlen and Newsome, 1998)] and 3.7 msec letters (the frame time of the visual display, but also close to the cutoff time between short and medium ISIs), so that each word is four letters long. For these word and letter lengths, the 16 repeats of our 60.6 sec stimulus provided more than adequate amounts of data to robustly estimate the information. Other values—word lengths ranging from 3.7 to 59.2 msec and letter lengths ranging from 1.8 to 14.8 msec—gave qualitatively similar results. In general, information values were highest when we used the shortest words and letters.

The information that the neuron transmits about a particular stimulus is defined as the signal entropy, or variability, minus the noise entropy. The signal entropy is derived from the total set of words spoken by the neuron during the course of its response. The noise entropy is derived from the set of the words spoken at each particular time in the response, and it is averaged across the entire response duration. The signal entropy is calculated by constructing a probability table of all words in the response and applying Shannon’s formula, $H_s = -\sum p_j \log_2 p_j$ (Cover and

Thomas, 1991), where p_j indicates the estimated probabilities of occurrence of each word. Because we only have access to a limited amount of data, this estimate of the signal entropy is subject to a downward bias, the correction for which is estimated by $(k - 1)/2N \ln(2)$, in bits, where k is the number of possible words and N is the total number of words observed (Carlton, 1969; Panzeri and Treves, 1996). In practice, the average correction to the signal entropy is 0.01% in our data sets, so including it was therefore inconsequential.

It is more difficult to obtain an accurate estimate of the noise entropy because we have access to only 16 trials, at most, for each neuron. However, we can obtain an upper bound on the noise entropy, and a lower bound on the transmitted information, by assuming that the letters in a word are independent of one another, and then adding the noise entropies letter by letter (Cover and Thomas, 1991). In this case, we do apply the correction for limited data, both because it represents a significant fraction of the noise entropy (~10%) and because it ensures that we are calculating a true lower bound on the transmitted information.

RESULTS

Interspike intervals

We report on the responses of 99 V1 neurons in anesthetized macaque monkeys to multiple repeats of pseudorandom (m-sequence) stimuli. These stimuli contain a wide variety of spatial and temporal patterns, none of which dominates the stimulus but some of which are typically effective stimuli for these neurons. The stimuli are therefore well suited to probe a neuron’s ability to convert spatial information into spike trains.

Figure 1 shows interspike interval histograms (ISIHS) constructed from the response of a simple cell to both the m-sequence stimulus (*solid line*) and a uniform field of the same mean luminance (*shaded gray region*). The top panel is the standard ISIH, in which ISIs are collected into equal bins of 1 msec width. For both the m-sequence and uniform-field responses, the standard ISIH features a prominent peak at very short ISIs, which decays rapidly. The peak is higher for the m-sequence response than for the uniform-field response, but it has approximately the same width. After the initial peak and rapid decay, the m-sequence ISIH shows a secondary, slower decay that lasts from ~5 msec until 20 msec; this secondary decay is almost entirely absent from the uniform-field response. Finally, both responses begin a very slow decline, the “long tail” (Gerstein and Mandelbrot, 1964; Smith and Smith, 1965), at ~20 msec.

To more prominently display the short-ISI features, we plot the same standard ISIH on a logarithmic time scale (Fig. 1, *middle panel*). The similar shape of the initial peak and the differential secondary decay are clearly evident in this plot, as well. However, because the binning is relatively coarse at short ISIs and relatively fine at high ISIs, which obscures detail on both time scales, we present the data in yet another way. In the *bottom panel*, the histogram is constructed with logarithmic time bins so that the duration of each successive bin is a fixed multiple of the previous one. This “log-ISIH,” a new way of looking at such data, highlights features such as prominent peaks or shoulders that are at best barely visible in the standard ISIH. The log-ISIH of the m-sequence response has three distinct peaks (corresponding, in the standard ISIH to the initial rapid peak, the secondary decay, and the long tail), whereas the log-ISIH of the uniform-field response has only two peaks (the initial rapid peak and the long tail). The differences between the log-ISIHs of the m-sequence response and the uniform-field response indicate that not all of the log-ISIH features are attributable to m-sequence stimulation. The similarity of the prominent short-ISI peak suggests that this feature, in particular, is largely intrinsic to the neuron.

In Figure 2, we replot the log-ISIH of the m-sequence response

Figure 1. Construction of the log-inter_spike interval histogram (log-ISIH). Each panel shows the ISIHs of the m-sequence response (solid line) and the response to a uniform field at the same mean luminance (shaded) of a simple cell (44/9t) that fired 21.0 spikes/sec, 35,581 ISIs total. *Top*, Standard ISIH (1 msec bins); *middle*, standard ISIH plotted on a logarithmic time axis; *bottom*, log-ISIH consisting of 300 bins spaced logarithmically between 1 msec and 10 sec, which bracketed the distribution of ISIHs found in the data. Each bin is 3.1% larger than the previous one. We obtain a relatively smooth histogram by applying a uniformly distributed, random timing jitter to each ISI, the magnitude of which is at most half the data collection resolution of 0.1 msec. The log-ISIH allows us to easily distinguish three separate ISI peaks, which are not readily visible in the standard ISIH.

from Figure 1 (*top panel*), and we add log-ISIHs from three additional neurons. Each log-ISIH is subdivided, by eye, into its component peaks; the number of peaks varies from neuron to neuron. Thus, the top neuron has three peaks, the next two have two peaks (which differ in position), and the fourth neuron has only a single peak. Peaks are gray scale-coded according to the relative position of the maximum: black for short, light gray for medium, and dark gray for long ISIs. The *bottom panel* is a log-histogram of the estimated boundary between ISI peaks; it summarizes data from 66 neurons and includes 19 short/medium boundary points and 32 medium/long boundary points. Across all 99 neurons, 34 had a single peak, 46 had evidence of two distinct peaks, and 19 had three peaks. We found no significant difference

Figure 2. Division of log-ISIHs into component peaks. The *top four panels* are the responses of different neurons to the m-sequence stimulus. Log-ISIHs consist of 200 bins spaced logarithmically between 0.1 msec and 2 sec, so that each bin is 5.1% larger than the previous one. Each log-ISIH is divided, by eye, into its component peaks, which are then color-coded according to the position of the maximum (black for <3 msec, light gray for 3–38 msec, and dark gray for >38 msec). *Top*, Same simple cell as in Figure 1; three peaks. *Second panel*, Simple cell (38/6), 5.9 spikes/sec, 11,343 ISIs; two peaks (short and long). *Third panel*, Simple cell (34/12t), 5.2 spikes/sec, 9,997 ISIs; two peaks (medium and long). *Fourth panel*, complex cell, 29.9 spikes/sec, 57,943 ISIs; one peak (medium). The boundary positions across 66 neurons are collected into a log-ISIH (*bottom panel*), which reveals a tightly clustered, bimodal distribution.

between simple and complex cells in terms of the number of log-ISIH peaks (χ^2 test).

Two very surprising results emerge from this analysis. First, the summary histogram is bimodal. This reflects the fact that there are at most three log-ISIH peaks in any neuron's response. Moreover, neurons with only two log-ISIH peaks (the majority) have boundary points that match either the short/medium or medium/long boundary of the three-peak neurons, rather than some intermediate value. The medians of the two boundary points are 3.0 msec (90% range: 1.8–3.7 msec) and 38.1 msec (90% range: 20.3–64.8 msec). The second surprising feature of the summary histogram is that the two modes are quite sharp, which indicates that log-ISIH peaks are remarkably consistent across neurons. Because the boundary points do not correspond to temporal features of the stimulus, which occur in multiples of 14.8 msec, or to the frame time of the visual display, which was 3.7 msec, we suspect that they reflect some aspect of the intrinsic biophysical hardware of cortical neurons and/or their connections.

The consistency of the boundary points across neurons also raises the possibility that the subsets delimited by those boundaries play distinct roles in information encoding. To address this possibility, we divided the ISIs in each neuron's response into three subsets, applying the median boundary points across neurons regardless of the number of peaks in the particular neuron's log-ISIH. We considered ISIs <3 msec to be part of the short-ISI subset, ISIs between 3 and 38 msec to be part of the medium-ISI subset, and ISIs >38 msec to be part of the long-ISI subset. Averaged across neurons, we found that 10% of spikes fell into the short-ISI subset, 45% into the medium, and 45% into the long. Box plots showing the distributions across neurons of the percentage of spikes in each subset are drawn in Figure 5 (*top*).

The first step in demonstrating that the individual subsets play important roles in information encoding is to show that they are not epiphenomena of the firing rate modulation. To do this, we performed an “exchange resampling” procedure (Victor and Purpura, 1996), which assigns to each trial in the resampled spike train the same number of spikes as had occurred in that trial in the real response. The spike times themselves are drawn at random, without replacement, from the entire set of actual spikes. Exchange resampling exactly preserves the firing rate modulation and the number of spikes in each response trial but randomizes the relationship between consecutive spikes. Log-ISIHs of exchange-resampled responses are shown as *solid lines* in Figure 3, superimposed on the log-ISIHs of real responses. If the ISI structure were completely determined by fast firing rate modulations and slow variations in responsiveness, the log-ISIH of real and resampled spike trains would superimpose. In fact, this superposition test succeeds for the longest ISIs but fails for short and medium ISIs. One cause of the failure is likely to be the presence, in real neurons, of a refractory period on the order of 1 msec. However, a refractory period alone cannot account for features such as a sharp notch between short and medium ISI peaks, or for amplitude differences such as the one in the *top left panel*.

Receptive field maps

For each neuron, we calculated RF maps or spike-triggered average stimuli (see Materials and Methods) from the full response and from response subsets, where the spikes in each subset were preceded by ISIs from a single ISI subset (short, medium, or long). Figure 4 shows example RF maps from two neurons,

Figure 3. Real versus exchange-resampled log-ISIHs. We compare the log-ISIHs (same parameters as in Fig. 2) of real spike trains (*shaded*) with those obtained from a resampling procedure that preserves the firing rate modulation and distribution of spikes per trial from the original data (*solid lines*). If the original data were consistent with a modulated Poisson process, then the original and resampled log-ISIHs in all panels would superimpose. *Top row, left panel (asterisk)*, Complex cell (33/1), 31.7 spikes/sec, 61,556 ISIs. *Top row, right panel (plus sign)*, Simple cell (39/9), 34.7 spikes/sec, 67,269 ISIs. *Middle row, left panel*, Simple cell (35/1), 10.0 spikes/sec, 19,283 ISIs. The other three log-ISIHs are from Figure 2.

depicting representative examples of the changes that we saw across ISI subsets. RF maps are shown as a series of interpolated contour plots, each averaged over 14.8 msec of the response, at different time lags; taken together, they depict the dynamics of the RF. Red signifies on-subregions of the RF, in which bright stimuli were associated with a higher firing rate than dark stimuli, and blue signifies off-subregions. The ratio of the RF map scales of two subsets is equal to the ratio of the numbers of spikes in the two subsets.

In general, we found that RF maps derived from spike subsets differed from RF maps derived from the full responses, and from each other, in both shape and amplitude. Figure 4*A* shows the RF map of a complex cell that primarily displayed both shape and amplitude changes. This cell fired an average of 31.7 spikes/sec in response to our stimulus, and its log-ISIH is shown in Figure 3 (*asterisk*). The RF map derived from all spikes (*top row*) is dominated by the off-subregion, although there is evidence of a weaker on-subregion. When only short-ISI spikes are considered (*second row*), however, the on-subregion is selectively enhanced in two ways: (1) its duration, measured with a resolution finer

Figure 4. Receptive-field maps. *A*, Complex cell (33/1), average firing rate 31.7 spikes/sec, spatial extent $1^{\circ}22' \times 1^{\circ}22'$, information rate 25.2 bits/sec or 0.79 bits/spike, efficiency 17.8%, log-ISIH in Figure 3 (*asterisk*). Calibration: -4.9 to 4.9 spikes/sec. *B*, Simple cell (39/9), average firing rate 34.7 spikes/sec, spatial extent $1^{\circ}22' \times 1^{\circ}39'$, information rate 31.2 bits/sec or 0.90 bits/spike, efficiency 20.5%, log-ISIH in Figure 3 (*plus sign*). Calibration: -12.9 to 12.9 spikes/sec. In each panel, the *top row* is the RF map of the full response, the *second row* is the RF map of the short-ISI subset (10% of spikes in *A*, 24% in *B*), the *third row* is the RF map of the medium-ISI subset (68 and 50%), and the *bottom row* is the RF map of the long-ISI subset (22 and 26%). Values on the calibration bar refer to the RF map of the full response only (f in Eq. 1 and 2); to obtain values for any of the ISI subsets (s), multiply the full-response values by the fraction of spikes in that subset.

than what is shown in Figure 4*A*, is 42 msec instead of 19 msec; and (2) its peak amplitude, after correcting for the different numbers of spikes, is 2.9 times greater than the corresponding amplitude in the all-spikes RF map. The relative amplitude of the off-subregion is also enhanced in the short-ISI RF map, but only by a factor of 1.8 and not at all time lags. The shape of the short-ISI RF map suggests that these spikes may selectively encode stimuli that are characterized by strong spatial opponency.

A second prominent RF map change in Figure 4*A* occurs in the long-ISI subset (*bottom row*). For this subset, the single RF subregion has a biphasic time course, changing from off to on at ~ 89 msec, even while the RF maps of other subsets are still dominated by an off-subregion. The time course of the long-ISI

RF map, and the fact that it is for the most part spatially uniform at each time lag, suggests that these spikes may primarily encode temporal features in the stimulus and that they typically follow, by at least 89 msec, a period during which the neuron was inhibited by the stimulus. Ignoring the ISI structure and treating all spikes equally obscures the effect of this inhibition.

Figure 4*B* is the RF map of a simple cell that fired an average of 34.7 spikes/sec in response to the m-sequence stimulus; the log-ISIH of its response is shown in Figure 3 (*plus sign*). The short- and medium-ISI RF maps show evidence of large amplitude changes but less evidence of large shape changes. In other words, scaling these subset RF maps by some factor (which we call the “efficacy”; see below) would make them look very similar to

the all-spikes RF map. The long-ISI RF map, however, primarily exhibits a shape change: the off-subregion is relatively diminished between 30 and 59 msec, and the on-subregion is relatively enhanced between 59 and 74 msec.

Two related indices can be used to quantify these shape and amplitude changes. To obtain these indices, we treat the RF maps as vectors in space and time, with dimension equal to the product of the number of stimulus pixels and the number of time bins. We define the *similarity index* (DeAngelis et al., 1999) to be the correlation coefficient between vectors derived from two different RF maps:

$$\text{similarity index} = \frac{\langle \mathbf{s}, \mathbf{f} \rangle}{\|\mathbf{s}\| \|\mathbf{f}\|} \quad (1)$$

where $\mathbf{f}$ is the vectorized RF map of the full spike train, $\mathbf{s}$ is the vectorized RF map of the subset spike train, $\langle \dots \rangle$ denotes the inner product, and $\|\dots\|$ denotes the norm, or the square root of a vector's inner product with itself. The similarity index has a value close to 1 for maps that have similar shapes (e.g., 0.98 for the full- and short-ISI RF maps in Fig. 4*B*), 0 for maps that are nearly orthogonal, and -1 for maps that have opposite polarities (on-subregions become off-subregions and vice versa). The similarity index is an omnibus measure that averages over both space and time and may therefore dilute the effects of prominent but localized RF map changes. Thus, even a qualitatively large shape change such as occurs in the short-ISI RF map of Figure 4*A* has a similarity index of 0.83, which is near the 75th percentile of short-ISI similarity indices across neurons (Fig. 5, *middle*). At the other extreme, the shape change seen in the long-ISI RF map of Figure 4*A*, which is among the most dramatic in our sample, has a similarity index of 0.60.

To measure amplitude changes, we calculate the *efficacy*, the factor by which the amplitude of the subset RF map must be scaled to best match the full RF map in a least-squares sense, after correcting for differences in the number of spikes. Mathematically:

$$\text{efficacy} = \left(\frac{n_f}{n_s} \right) \left(\frac{\langle \mathbf{s}, \mathbf{f} \rangle}{\|\mathbf{f}\|^2} \right) \quad (2)$$

where n_f is the number of spikes in the full spike train, n_s is the number of spikes in the subset spike train, and other variables are as in Equation 1. An efficacy >1 signifies that subset spikes play a larger role in the generation of the all-spikes RF map than expected given the number of spikes in that subset, whereas an efficacy <1 signifies the opposite. An efficacy <0 would indicate that the RF map polarity must be flipped and then scaled to obtain the best match. The short-ISI RF map of Figure 4*B* has a particularly high efficacy of 1.77, signifying a large contribution from those spikes to the overall RF properties of that neuron. On the other hand, the medium-ISI RF map of Figure 4*A* has an efficacy of 1.02, indicating that those spikes contribute the expected amount to the overall RF map.

Across all neurons, similarity indices are in the range 0.5–1 (Fig. 5, *middle*), meaning that shape changes, although present, are only moderate. On the other hand, the efficacy distributions of the three subsets are distinct and largely nonoverlapping (Fig. 5, *bottom*), which indicates that short-ISI spikes make the largest contribution to the all-spikes RF maps, long-ISI spikes make the smallest contribution, and medium-ISI spikes contribute about as much as expected given their frequency.

Figure 5. Across-neuron distributions of RF map indices. Data from 66 neurons. Box plot whiskers show the 5th and 95th percentiles, and box boundaries show the 25th and 75th percentiles. The horizontal line divides each box at the median, and the square represents the mean. *Top*, Percentage of spikes with preceding ISIs <3 msec (*short*), 3–38 msec (*medium*), and >38 msec (*long*). *Middle*, Similarity index (or correlation coefficient) between the subset and full-response RF maps; values can be between -1 and $+1$. *Bottom*, Efficacy, or average contribution made by spikes in each subset to the overall RF map properties. An efficacy of 1 means that the spikes in that subset contribute as much as expected, given their number. See Results for mathematical definitions of the similarity index and efficacy.

Rate versus temporal encoding

To test whether our findings about shape and amplitude changes of RF maps are consistent with a rate code model of neuronal firing, we compared real data with the results of an exchange-resampling procedure. Exchange resampling exactly preserves the firing-rate modulation inherent in the original data, as described above. In fact, this procedure yields the only ensemble of spike trains that match the time-varying firing rate of real data and are rigorously consistent with a firing-rate model. We exchange-resampled each spike train 200 times, and we calculated the similarity indices and efficacies of the subset RF maps of those resampled spike trains. We used the same ISI subset definitions for resampled spike trains as we had for the real spike trains, despite the fact that the log-ISIHs were different (Fig. 3). We then

Figure 6. RF map changes: real versus exchange-resampled. In each panel, we plot the similarity index (*left column*) and efficacy (*right column*) for real data on the horizontal axis, and the mean value of the same parameter derived from 200 exchange resamplings of the real data on the vertical axis. Data are from 66 neurons, and each point represents a different neuron. The exchange-resampling procedure tests the hypothesis that the observed RF map changes could have been generated by a neuron that fires spikes according to a rate code with exactly the same firing rate modulation as the original response. *Filled squares* represent subsets for which the real data were significantly different from the resampled data, that is, points that fell significantly ($p < 0.05$, two-tailed direct comparison) off the diagonal in each panel.

compared the similarity indices and efficacies of the real responses with the means and distributions obtained from exchange-resampled responses.

Figure 6 shows a series of scatter plots that describe the relationship between similarity indices (*left column*) and efficacies (*right column*) for real and resampled responses. Each point represents a different neuron. The similarity index or efficacy of that neuron's real RF map is plotted along the horizontal axis, and the corresponding mean value from 200 exchange resamplings is plotted along the vertical axis. Results from different subsets are in different rows. In nearly all panels, the cloud of points lies near the line of equality, where real and resampled RF maps have the same similarity indices or efficacies. The major exception is the *top right panel*, which shows that the short-ISI efficacies tend to be higher for resampled RF maps than for real RF maps, indicating that short-ISI spikes play an even larger role in generating the RF properties of rate-code-generated spike trains than of real spike trains.

On a neuron-by-neuron basis, the similarity indices and efficacies of the real spike trains tended to fall in the tails of the distributions of similarity indices and efficacies of resampled spike trains. In Figure 6, *solid points* represent cases in which the real data fell within either 2.5% tail of the distribution of resampled data. However, the behavior of the two indices was qualitatively, and for the most part quantitatively, similar for real and resampled RF maps, suggesting that the rate code model does a reasonable job of explaining the indices (although not the ISI statistics).

Information

In addition to comparing RF maps, we calculated the stimulus-related information carried by spikes in the full spike train and in each ISI subset (see Materials and Methods) (de Ruyter van Steveninck et al., 1997; Strong et al., 1998). The results are complementary to the ones obtained by RF map calculation, which present, quite literally, a qualitative and quantitative picture of the stimulus-encoding characteristics of spikes in particular subsets.

The results of the information calculation are shown in Figure 7. Not surprisingly, the full spike train conveys the most stimulus-related information on an absolute scale of bits per second (*top*), more than is carried by any subset of that spike train. Short-ISI spikes, which are rarest, convey the least information. On a neuron-by-neuron basis, the sum of the transmitted information across the three ISI subsets is larger than the information carried by the full spike train. This is perplexing, because the full spike train must contain as much information as the sum of its subsets. However, there are at least two explanations for this finding. First, because we calculated the information transmitted via a code consisting of 14.8 msec words and 3.7 msec letters, and because our method was designed to give us a lower bound on information in the first place, it is possible that we systematically underestimate the information in the full spike train to a greater extent than in the subset spike trains. Second, the three subset spike trains may contain redundant information. This is likely because the process of selecting a particular ISI subset necessarily places constraints on the other, nonoverlapping subsets. These constraints and redundancies are not taken into account in the word/letter code.

On a bits per spike scale (Fig. 7, *middle*), the short-ISI spikes, which convey the least information per second, are actually most informative: on average, each short-ISI spike conveys 4.6 bits of information about the stimulus, corresponding to the high efficacy of these spikes. Medium- and long-ISI subset spikes convey substantially less information, not much more than what is conveyed by each spike in the full spike train. This is also reflected in our measure of efficiency (de Ruyter van Steveninck et al., 1997), which compares the transmitted information to the total signal entropy: the median efficiency of short-ISI spikes is 45%, whereas the median efficiencies of the other spike subsets are near 25%.

Finally, Figure 8 shows the neuron-by-neuron comparison between real and exchange-resampled information values, for the full spike train and for the three ISI subsets. The results are qualitatively similar to the results for the similarity index and efficacy, which are shown in Figure 6. In general, whether measured on a bits per second or bits per spike scale, and regardless of the particular ISI subset, the information in a real spike train is similar to the information in its exchange-resampled counterparts: most points lie along the line of equality. This means that the qualitative differences among subsets would be expected from

Figure 7. Across-neuron distributions of transmitted information values. Data are from 98 neurons. As in Figure 5, box plot whiskers show the 5th and 95th percentiles, and box boundaries show the 25th and 75th percentiles. The horizontal line divides each box at the median, and the square represents the mean. The full response (leftmost distribution in each panel) is compared with each of the three ISI subset responses. *Top*, Transmitted information, measured in units of bits per second (*bits/s*). *Middle*, Transmitted information, bits per spike (*bits/spike*). *Bottom*, Efficiency, or percentage of the total signal entropy (see Materials and Methods) that is used to transmit stimulus-related information.

neurons that use a rate code. On the other hand, nearly all points fall significantly off the line of equality ($p < 0.05$, two-tailed direct comparison, $\blacksquare$ vs $\square$), indicating that on a neuron-by-neuron basis, the results are quantitatively inconsistent with the predictions of a rate code.

DISCUSSION

Our results provide new evidence that different spikes within a single response can convey messages about different stimulus features, a finding that is consistent with earlier reports (McClurkin et al., 1991; Victor and Purpura, 1996). Our results advance this work in three ways. (1) They identify, through the log-ISIH peaks and in the form of ISI subsets, the time scales that are relevant for V1 neurons responding to rapidly varying stimuli. (2) They provide, through the RF maps, a direct picture of the receptive field properties of the different spikes. (3) They suggest

Figure 8. Information: real versus exchange-resampled. In each panel, we plot the information in bits per second (*bits/s*) (left column) and bits per spike (*bits/spike*) (right column) for real data on the horizontal axis, and the mean value of the same parameter derived from 40 exchange resamplings of the real data on the vertical axis. Data are from 98 neurons, and each point represents a different neuron. The exchange resampling procedure tests the hypothesis that the observed RF map changes could have been generated by a neuron that fires spikes according to a rate code with exactly the same firing rate modulation as the original response. *Filled symbols* represent subsets for which the real data were significantly different from the resampled data, that is, points that fell significantly ($p < 0.05$, two-tailed direct comparison) off the diagonal in each panel.

that the information encoded by individual spikes can be decoded by classifying it on the basis of the duration of the immediately preceding ISI.

In a pioneering study of the H1 neuron of the blowfly (de Ruyter van Steveninck and Bialek, 1988), a Gaussian white noise stimulus was used to evaluate the probability distributions of stimuli that typically preceded arbitrary temporal response patterns consisting of up to three spikes and two ISIs. The authors showed that the mean stimulus changed gradually with the duration of the ISI before each spike. In our experiments, we did not

obtain enough data from each neuron to be able to perform an identical analysis. Although it is likely that the same result holds in monkey V1, that small changes in ISI boundary points yield small changes in RF maps, the existence and consistency across neurons of the three ISI subsets appear novel and lead us to hypothesize that the grouping of ISIs into subsets similar to the ones we have described reflects natural modes of information processing. It is tempting to relate the ISI subsets to phenomena such as oscillatory responses (Gray et al., 1989)—the typical interval in the medium-ISI subset corresponds to a frequency of around 40 Hz—but we have no evidence to indicate that the two findings are related.

Bursts

Our results suggest that short-ISI spikes are especially important for the transmission of visual information. These short-ISI spikes likely correspond to the class of spikes known as “bursts” (Connors and Gutnick, 1990), although the correspondence cannot be conclusively established from extracellular recordings alone. In the thalamus, the mechanisms of burst production are well known (Jahnsen and Llinás, 1984; Sherman, 1996), and the relevance of bursts for information encoding and transmission in the lateral geniculate nucleus, in particular, has been studied (Mukherjee and Kaplan, 1995; Reinagel et al., 1999; Usrey and Reid, 1999). Reinagel et al. (1999) found that burst spikes convey 1.5 to 3 times as much information as tonic spikes, a finding that is similar to our own finding about short-ISI spikes.

In the visual cortex, where the mechanisms of burst production are less well understood, several lines of evidence suggest that short-ISI spikes are particularly important for information transmission. Compared with other spikes, short-ISI spikes are differently tuned to certain stimulus attributes (Cattaneo et al., 1981; Legéndy and Salzman, 1985; Livingstone et al., 1996; DeBusk et al., 1997) and tend to be more likely to evoke a postsynaptic response (Alonso et al., 1996; Lisman, 1997; Snider et al., 1998). In earlier work, we demonstrated that short-ISI spikes tend to be reliable (Victor et al., 1998) in that they are fired at the same time on multiple repeats of a stimulus. Here, we have shown that short-ISI spikes are effective, too, in that they make a larger than expected contribution to a neuron's RF properties. We have also demonstrated that the RF maps of short-ISI spikes tend to have regions of stark spatial opponency, suggesting that these spikes may preferentially extract features such as bars and lines from visual stimuli. Finally, we have shown that short-ISI spikes, in part because of their relative paucity, convey more information per spike than spikes in the other ISI subsets and that they do so with higher efficiency.

Information

We found that most V1 neurons convey between 5 and 30 bits/sec, or between 1 and 3 bits/spike. This range is consistent with information rates calculated from responses to similar stimuli in other systems, ranging from fly to primate (Buračas and Albright, 1999). It is, however, at least an order of magnitude higher than the information rate in V1 responses to flashed stimuli, which ranges from 0.1 to 0.5 bits/sec (Richmond and Optican, 1990; Victor and Purpura, 1996; Gershon et al., 1998). A primary cause of this discrepancy may be that V1 neurons very efficiently convey information about stimulus transients, which occur 64 times/sec in the m-sequence stimulus but only once in the flashed stimuli. However, this does not rule out the possibility that these neurons are intrinsically more efficient in conveying information about

elementary stimulus features such as contrast and orientation when the stimulus is rapidly varying than when it is constant.

Another reason that the information rates are so different for the two kinds of stimuli is methodological: the “direct method” of calculating information used here makes no assumptions about the stimuli, but rather compares the response variability over time with the response variability across trials. Methods that typically yield lower information rates, on the other hand, evaluate a neuron's ability to discriminate between N particular stimuli, which limits the total amount of information that can be encoded to $\log_2 N$. In this regard, it is important to realize that our stimuli are not optimized to drive V1 neurons to convey information at a rate close to their channel capacity.

Encoding and decoding

The processes of encoding and decoding information are logically distinct. It is possible, for example, that V1 neurons encode information into their rapidly modulated firing rates by means of a Poisson spike generator, consistent with a rate-coding model, even while they decode information by measuring presynaptic ISI durations. In this paper, we compare real responses with exchange-resampled responses, which we consider to be rate-coded because the spike times are determined only from the firing rate. Because exchange-resampling exactly preserves spike times, the firing rate fluctuations occur at the same rate as they do in real data.

According to one definition (Borst and Théunissen, 1999), simply demonstrating that information is carried on time scales more rapid than the time scale of stimulus fluctuations, as we do here in our information calculations with short word and letter lengths, constitutes a demonstration that information is temporally encoded. By this criterion, even exchange-resampled spike trains, not to mention real ones, are temporally encoded (Fig. 8). It is not surprising, therefore, that the RF maps and information values of exchange-resampled spike trains are similar to, albeit significantly different from, the ones derived from real neurons. It is more surprising that their log-ISIHs are so different (Fig. 3), although they share some features such as multiple peaks and sometimes even peak positions. Moreover, another study that used different stimuli and different analysis methodology has shown that real V1 spike trains are not fully consistent with the predictions of rate-coding models (Reich et al., 1998).

In our view, the consistent presence of distinct ISI subsets across neurons, and the different types of visual information that can be extracted by examining spikes from the various subsets in isolation, suggest that ISI decoding may be an important feature of V1 neurons, just as it has been shown to be in much simpler systems such as the visceral ganglia of *Aplysia* (Segundo et al., 1963). To accomplish this type of decoding, neurons need not do anything more sophisticated than be sensitive to the durations of individual ISIs. This sensitivity can be embodied in a single synapse and does not require averaging across stimulus repeats, stretches of time that may be long compared with the time scale of firing rate modulation, or a large population of neurons that carry similar information.

Although our results provide support for the hypothesis that ISI decoding plays a role in information transfer in the visual cortex, they are also consistent with other types of decoding schemes that do not make use of ISIs. Such a scheme includes, for example, the estimation of firing rates through averaging across many neurons that convey similar information (Shadlen and Newsome, 1998). Cortical microstimulation (Salzman et al., 1992)

influences both rate and ISI structure and is thus consistent with both views of neural coding. A direct experimental resolution of the roles of these different kinds of decoding schemes would require manipulation of the ISI structure of neural activity without changing the average firing rate, and observation of the effect of this manipulation on an animal's behavior. In mammalian cortex, the design and execution of such experiments are challenges to current techniques, but it is interesting to note that such manipulations do indeed affect olfactory discrimination in the locust (Stopfer et al., 1997) and gustatory perception in the rat (Di Lorenzo and Hecht, 1993).

Real synapses may accomplish ISI decoding by means of processes such as short-term, real-time synaptic modification, including synaptic depression and facilitation (Gerstner et al., 1997; Markram et al., 1998; Goldman et al., 1999), and dendritic nonlinearities, including coincidence detection (Abeles, 1982; Mel, 1994; Margulis and Tang, 1998; Yuste et al., 1999). These processes can selectively weight individual spikes based on the durations of single ISIs (Maass and Zador, 1999), and their particular form can depend on the types of neurons that are connected by each synapse (Thomson, 1997; Reyes et al., 1998). Synaptic facilitation tends to enhance the synaptic efficacy of short-ISI spikes, whereas synaptic depression tends to diminish it, thereby increasing the relative efficacy of medium- and long-ISI spikes (Gerstner et al., 1997). Thus, we suggest that a primary role of short-term synaptic modification is to aid in the decoding of information about multiple stimulus features that would be missed if all spikes were treated equally.

REFERENCES

- Abeles M (1982) Role of the cortical neuron: integrator or coincidence detector? *Isr J Med Sci* 18:83–92.
- Alonso JM, Usrey WM, Reid RC (1996) Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* 383:815–819.
- Borst A, Theunissen FE (1999) Information theory and neural coding. *Nat Neurosci* 2:947–957.
- Buračas GT, Albright TD (1999) Gauging sensory representations in the brain. *Trends Neurosci* 22:303–309.
- Carlton AG (1969) On the bias of information estimates. *Psychol Bull* 71:108–109.
- Cattaneo A, Maffei L, Morrone C (1981) Patterns in the discharge of simple and complex visual cortical cells. *Proc R Soc Lond B Biol Sci* 212:279–297.
- Connors BW, Gutnick MJ (1990) Intrinsic firing patterns of diverse neocortical neurons. *Trends Neurosci* 13:99–104.
- Cover TM, Thomas JA (1991) Elements of information theory. New York: Wiley.
- de Ruyter van Steveninck R, Bialek W (1988) Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc R Soc Lond B Biol Sci* 234:379–414.
- de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W (1997) Reproducibility and variability in neural spike trains. *Science* 275:1805–1808.
- DeAngelis GC, Ghose GM, Ohzawa I, Freeman RD (1999) Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons. *J Neurosci* 19:4046–4064.
- DeBusk BC, DeBruyn EJ, Snider RK, Kabara JF, Bonds AB (1997) Stimulus-dependent modulation of spike burst length in cat striate cortical cells. *J Neurophysiol* 78:199–213.
- Di Lorenzo PM, Hecht GS (1993) Perceptual consequences of electrical stimulation in the gustatory system. *Behav Neurosci* 107:130–138.
- Gershon ED, Wiener MC, Latham PE, Richmond BJ (1998) Coding strategies in monkey V1 and inferior temporal cortices. *J Neurophysiol* 79:1135–1144.
- Gerstein GL, Mandelbrot B (1964) Random walk models for the spike activity of a single neuron. *Biophys J* 4:41–68.
- Gerstner W, Kreiter AK, Markram H, Herz AVM (1997) Neural codes: firing rates and beyond. *Proc Natl Acad Sci USA* 94:12740–12741.
- Goldman MS, Nelson SB, Abbott LF (1999) Decorrelation of spike trains by synaptic depression. *Neurocomputing* 26–27:147–153.
- 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.
- Jahnsen H, Llinás R (1984) Electrophysiological properties of guinea-pig thalamic neurones: an in vitro study. *J Physiol (Lond)* 349:205–226.
- Legéndy CR, Salzman M (1985) Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons. *J Neurophysiol* 53:926–939.
- Lisman JE (1997) Bursts as a unit of neural information: making unreliable synapses reliable. *Trends Neurosci* 20:38–43.
- Livingstone MS, Freeman DC, Hubel DH (1996) Visual responses in V1 of freely viewing monkeys. *Cold Spring Harb Symp Quant Biol* 61:27–37.
- Maass W, Zador A (1999) Dynamic stochastic synapses as computational units. *Neural Comput* 11:903–917.
- Mainen ZF, Sejnowski TJ (1995) Reliability of spike timing in neocortical neurons. *Science* 268:1503–1506.
- Margulis M, Tang CM (1998) Temporal integration can readily switch between sublinear and supralinear summation. *J Neurophysiol* 79:2809–2813.
- Markram H, Gupta A, Uziel A, Wang Y, Tsodyks M (1998) Information processing with frequency-dependent synaptic connections. *Neurobiol Learn Mem* 70:101–112.
- McClurkin JW, Optican LM, Richmond BJ, Gawne TJ (1991) Concurrent processing and complexity of temporally encoded neuronal messages in visual perception. *Science* 253:675–677.
- Mel BW (1994) Information processing in dendritic trees. *Neural Comput* 6:1031–1085.
- Mukherjee P, Kaplan E (1995) Dynamics of neurons in the cat lateral geniculate nucleus: in vivo electrophysiology and computational modeling. *J Neurophysiol* 74:1222–1243.
- Nowak LG, Sanchez-Vives MV, McCormick DA (1997) Influence of low and high frequency inputs on spike timing in visual cortical neurons. *Cereb Cortex* 7:487–501.
- Ogawa T, Ito S, Kato H (1981) Membrane characteristics of visual cortical neurons in vitro slices. *Brain Res* 226:315–319.
- Panzeri S, Treves A (1996) Analytical estimates of limited sampling biases in different information measures. *Network* 7:87–100.
- Reich DS, Victor JD, Knight BW (1998) The power ratio and the interval map: spiking models and extracellular recordings. *J Neurosci* 18:10090–10104.
- Reid RC, Victor JD, Shapley RM (1997) The use of m-sequences in the analysis of visual neurons: linear receptive field properties. *Vis Neurosci* 14:1015–1027.
- Reinagel P, Godwin D, Sherman SM, Koch C (1999) Encoding of visual information by LGN bursts. *J Neurophysiol* 81:2558–2569.
- Reyes A, Lujan R, Rozov A, Burnashev N, Somogyi P, Sakmann B (1998) Target-cell-specific facilitation and depression in neocortical circuits. *Nat Neurosci* 1:279–285.
- Richmond BJ, Optican LM (1990) Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. II. Information transmission. *J Neurophysiol* 64:370–380.
- Salzman CD, Murasugi CM, Britten KH, Newsome WT (1992) Microstimulation in visual area MT: effects on direction discrimination performance. *J Neurosci* 12:2331–2355.
- Segundo JP, Moore GP, Stensaas LJ, Bullock TH (1963) Sensitivity of neurones in *Aplysia* to temporal patterns of arriving impulses. *J Exp Biol* 40:643–667.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
- Sherman SM (1996) Dual response modes in lateral geniculate neurons: mechanisms and functions. *Vis Neurosci* 13:205–213.
- Skottun BC, De Valois RL, Grosof DH, Movshon JA, Albrecht DG, Bonds AB (1991) Classifying simple and complex cells on the basis of response modulation. *Vision Res* 31:1079–1086.
- Smith DR, Smith GK (1965) A statistical analysis of the continual activity of single cortical neurones in the cat unanesthetized isolated forebrain. *Biophys J* 5:47–74.
- Snider RK, Kabara JF, Roig BR, Bonds AB (1998) Burst firing and modulation of functional connectivity in cat striate cortex. *J Neurophysiol* 80:730–744.

- Stopfer M, Bhagavan S, Smith BH, Laurent G (1997) Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 390:70–74.
- Strong SP, Koberle R, de Ruyter van Steveninck RR, Bialek W (1998) Entropy and information in neural spike trains. *Phys Rev Lett* 80:197–200.
- Sutter EE (1992) A deterministic approach to nonlinear systems analysis. In: *Nonlinear vision: determination of neural receptive fields, function, and networks* (Pinter RB, Nabet B, eds), pp 171–220. Boca Raton, FL: CRC.
- Thomson AM (1997) Activity-dependent properties of synaptic transmission at two classes of connections made by rat neocortical pyramidal axons in vitro. *J Physiol (Lond)* 502:131–147.
- Usrey WM, Reid RC (1999) Synchronous activity in the visual system. *Annu Rev Physiol* 61:435–456.
- Victor JD (1992) Nonlinear systems analysis in vision: overview of kernel methods. In: *Nonlinear vision: determination of neural receptive fields, function, and networks* (Pinter RB, Nabet B, eds), pp 1–37. Boca Raton, FL: CRC.
- Victor JD, Purpura KP (1996) Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J Neurophysiol* 76:1310–1326.
- Victor JD, Purpura KP (1998) Spatial phase and the temporal structure of the response to gratings in V1. *J Neurophysiol* 80:554–571.
- Victor JD, Reich DS, Purpura KP, Mechler F (1998) Bursts and “reliable” spikes generated by neurons in V1: relationships and spatiotemporal origins. *FASEB: Retinal Neurobiology and Visual Perception*.
- Yuste R, Majewska A, Cash SS, Denk W (1999) Mechanisms of calcium influx into hippocampal spines: heterogeneity among spines, coincidence detection by NMDA receptors, and optical quantal analysis. *J Neurosci* 19:1976–1987.