

# Formal and Attribute-Specific Information in Primary Visual Cortex

DANIEL S. REICH,<sup>1,2</sup> FERENC MECHLER,<sup>2</sup> AND JONATHAN D. VICTOR<sup>2</sup>

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

Received 10 April 2000; accepted in final form 27 September 2000

**Reich, Daniel S., Ferenc Mechler, and Jonathan D. Victor.** Formal and attribute-specific information in primary visual cortex. *J Neurophysiol* 85: 305–318, 2001. We estimate the rates at which neurons in the primary visual cortex (V1) of anesthetized macaque monkeys transmit stimulus-related information in response to three types of visual stimulus. The stimuli—randomly modulated checkerboard patterns, stationary sinusoidal gratings, and drifting sinusoidal gratings—have very different spatiotemporal structures. We obtain the overall rate of information transmission, which we call *formal* information, by a *direct method*. We find the highest information rates in the responses of simple cells to drifting gratings (median: 10.3 bits/s, 0.92 bits/spike); responses to randomly modulated stimuli and stationary gratings transmit information at significantly lower rates. In general, simple cells transmit information at higher rates, and over a larger range, than do complex cells. Thus in the responses of V1 neurons, stimuli that are rapidly modulated do not necessarily evoke higher information rates, as might be the case with motion-sensitive neurons in area MT. By an extension of the direct method, we parse the formal information into *attribute-specific* components, which provide estimates of the information transmitted about contrast and spatiotemporal pattern. We find that contrast-specific information rates vary across neurons—about 0.3 to 2.1 bits/s or 0.05 to 0.22 bits/spike—but depend little on stimulus type. Spatiotemporal pattern-specific information rates, however, depend strongly on the type of stimulus and neuron (simple or complex). The remaining information rate, typically between 10 and 32% of the formal information rate for each neuron, cannot be unambiguously assigned to either contrast or spatiotemporal pattern. This indicates that some information concerning these two stimulus attributes is confounded in the responses of single neurons in V1. A model that considers a simple cell to consist of a linear spatiotemporal filter followed by a static rectifier predicts higher information rates than are found in real neurons and completely fails to replicate the performance of real cells in generating the confounded information.

## INTRODUCTION

Recent studies of the responses of visual neurons to stimuli with rich temporal structure, such as flickering checkerboard patterns and drifting gratings that abruptly change direction, have pointed to overall information transmission rates of between 5 and 100 bits/s (Berry and Meister 1998; Buračas et al. 1998; Reich et al. 2000a; Reinagel and Reid 2000; Ruyter van Steveninck et al. 1997). The sensory systems analyzed in these studies range from blowfly lobula plate to primate cortex. The information calculations are based on the *direct method* (Ruyter van Steveninck et al. 1997; Strong et al. 1998), which

estimates the overall rate of information transmission in a set of responses to a single stimulus.

Earlier studies on neurons in primary visual cortex, based on slowly fluctuating stimuli, report information rates an order of magnitude lower (Gershon et al. 1998; Heller et al. 1995; Mechler et al. 1998b; Richmond and Optican 1990; Tolhurst 1989; Victor and Purpura 1996). These studies use a variety of methods other than the direct method to calculate the information rates in responses to sets of stimuli that vary along some particular parameter, such as contrast or spatial pattern. All of these methods calculate information as a measure of the degree to which responses can be clustered into the appropriate stimulus classes.

Comparing these and similar results, Buračas and Albright (1999) argue that neurons, especially cortical neurons, more effectively convey information about stimuli with rich temporal structure than about stimuli with simpler structure. This argument is incomplete, however, because the two sets of studies use qualitatively different approaches to measuring transmitted information, both in terms of the richness of the stimuli and in terms of the analysis method. It is therefore impossible to draw conclusions about the types of stimulus that evoke the highest information rates from such a comparison. Here we link the results of these two sorts of studies by recording the responses of V1 neurons to a battery of stimuli of different spatiotemporal structure and by analyzing the responses in a uniform fashion (a variant of the direct method). Our major finding is that the overall rate of information transmission—which we dub formal information—does vary with stimulus type but that responses to rapidly modulated stimuli do not necessarily convey the most information, particularly in the case of simple cells.

Next we draw a distinction between formal and attribute-specific information rates. Formal information concerns all aspects of the response that depend on the stimulus. Attribute-specific information concerns only aspects of the response that allow the discrimination between stimuli that differ in some particular attribute, such as contrast, in the face of variation in other attributes, such as spatiotemporal pattern. The attribute-specific information is a measure of the degree to which responses to different stimuli cluster according to a particular stimulus attribute and is thus more comparable to the information measured in the second type of study mentioned above. By presenting each type of stimulus at multiple contrasts and

Address for reprint requests: D. Reich, The Rockefeller University, 1230 York Ave., Box 200, New York, NY 10021 (E-mail: reichd@rockefeller.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

appropriately modifying the direct method, we parse the formal information rate into attribute-specific components relating to contrast and spatiotemporal pattern. Here, spatiotemporal pattern refers to a broad category of stimulus attributes that includes temporal fluctuations as well as variations in spatial phase.

Overall, we find that information about contrast is transmitted at a significantly slower rate than information about spatiotemporal pattern, although not for every type of neuron and stimulus. The rate of contrast-specific information transmission depends little on stimulus type. Contrast-specific information rates estimated by the direct method are very similar to contrast-specific information rates estimated by a method based on computing the distances between pairs of spike trains (Victor and Purpura 1996).

We find that contrast- and spatiotemporal pattern-specific information rates together account for less than the full formal information rate—typically 68–90%. This indicates that a significant portion of the information in V1 responses relates to a *confounded* representation of contrast and spatiotemporal pattern by which the spatiotemporal pattern of the stimulus is encoded in a contrast-dependent fashion and the contrast of the stimulus in a spatiotemporal pattern-dependent fashion. An observer who is only aware of the portion of a single neuron's response that contains the confounded information cannot draw conclusions about contrast or spatiotemporal pattern in isolation. The possibility still exists—though it is not addressed in this paper—that the confounded information may be parsed into individual components by considering the simultaneous responses of additional neurons.

We ask whether a basic model of a V1 simple cell can account for our results. This model consists of a linear spatiotemporal filter, which we derive from the responses of a real neuron, followed by a static rectifier and a Poisson spike generator. The responses of such a model to the same stimuli presented to real neurons transmit formal information at rates comparable to, but higher than, those of real responses. However, unlike in the responses of real neurons, all of the formal information in model responses can be parsed into attribute-specific components: the model does not confound the encoding of contrast and spatiotemporal pattern. We have shown that this discrepancy is not due to differences in the underlying dynamics of spike generation, which do not strongly determine information rates (Reich et al. 2000a). Instead, the discrepancy occurs because this basic model lacks certain features of real cortical neurons, such as contrast gain control (Ohzawa et al. 1982), contrast normalization (Albrecht and Geisler 1991; Heeger 1992), and pattern gain control (Carandini et al. 1997a), by which variations in one stimulus attribute can affect the encoding of another.

Portions of this work have appeared in abstract form (Reich et al. 2000b).

## METHODS

We present data from recordings of individual neurons in the primary visual cortices of sufentanil-anesthetized macaque monkeys. Our detailed experimental procedures have been described elsewhere (Reich et al. 2000a; Victor and Purpura 1998). We use three types of stimulus: flickering checkerboards modulated by m-sequences, drifting sinusoidal gratings, and transiently presented, stationary sinusoidal gratings. All stimuli are presented on a Tektronix 608 monitor

with a mean luminance of approximately 150 cd/m<sup>2</sup> and a frame rate of 270.329 Hz.

### *M-sequence stimuli*

The principles and methodology of the m-sequence checkerboard stimuli have been extensively described (Reich et al. 2000a; Reid et al. 1997; Sutter 1992). In the experiments reported here, we use two such stimuli: a 12th-order m-sequence (4,095 stimulus frames) modulating 249 stimulus checks and a 9th-order m-sequence (511 frames) modulating 25 checks. In both cases, each frame lasts for 14.8 ms (four monitor refreshes), so that the total 12th-order stimulus lasts 60.6 s and the 9th-order sequence lasts 7.6 s. Individual checks typically span  $16 \times 16$  arc-min of visual angle and are arranged in a square. The size and orientation of the array are sometimes adjusted based on the neuron's spatial-frequency preferences, but only in cases where the adjustment is expected to produce a dramatically larger response.

The 249-check (*long*) stimulus is surrounded by a black circular aperture, and the 25-check (*short*) stimulus is surrounded by a uniform field at the mean luminance. In both stimuli, every check is modulated by the same m-sequence, but the starting point in the sequence varies from check to check. The minimum offset between starting points is 237 ms (64 samples of the m-sequence). The use of m-sequences in this way ensures that there are essentially no pair-wise correlations in time within individual checks, or in space across checks, that are relevant to the neuron's response. The long stimulus is presented at a single contrast (1), and the short stimulus is presented at each of five geometrically spaced contrasts (0.0625, 0.125, 0.25, 0.5, and 1). Both standard and contrast-inverted (reversed dark and light checks) sequences, each repeated 12–16 times, are presented in the long stimulus. Standard and inverted sequences within a repeat are separated by a period of 23 s during which a uniform field at the mean luminance is presented; repeats are separated by 18 s. For the short stimulus, no inverted sequences are presented. Contrasts are presented in increasing order, separated by uniform-field presentations lasting 10 s, and the entire set of contrasts is repeated 25–100 times with 10 s between repeats.

### *Drifting-grating stimuli*

We use “optimal” sinusoidal gratings with spatial frequency, temporal frequency, and orientation chosen to maximize either firing rate (for complex cells) or response modulation at the driving frequency (for simple cells) (Skottun et al. 1991). For simultaneously recorded neurons, we optimize the gratings for at least one of the cells, usually the one with the most distinct extracellularly recorded waveform (since this is the neuron most easily monitored during the experiment). The parameter choices based on the response of this cell are likely to be similar to the parameters that would have been chosen for the other simultaneously recorded neurons (DeAngelis et al. 1999), as we occasionally verified empirically. We present the gratings for 4 s at each of six geometrically spaced contrasts (0, 0.0625, 0.125, 0.25, 0.5, and 1). We repeat the entire set of contrasts five to eight times, with the order of contrast presentation randomized within blocks. For 4-Hz gratings, this yields 80–128 stimulus cycles at each contrast. Within each block, grating presentations at different contrasts are separated by presentation of a uniform field of the same mean luminance for 8 s, and blocks are separated by presentation of the uniform field for 13 s.

### *Stationary-grating stimuli*

We present stationary sinusoidal gratings at the same spatial frequency and orientation as the drifting gratings. The spatial phase of the stationary gratings is the one that maximizes the neuron's firing rate in response to stationary gratings of unit contrast. For each neuron, we present in increasing order either seven geometrically

spaced (0, 0.03125, 0.0625, 0.125, 0.25, 0.5, 1) or nine arithmetically spaced (0, 0.125, 0.25, 0.375, 0.5, 0.625, 0.75, 0.875, 1) contrasts. Gratings replace a uniform field of the same mean luminance for a period of 237 ms, after which the uniform field reappears for a minimum of 710 ms. The amount of time between grating presentations increases as a function of the contrast of the preceding grating. For example, the amount of off time following the 0.5 contrast presentation is 2.84 s and following the 0.875 contrast presentation is 4.26 s. This strategy is used to approximate a uniform state of contrast adaptation (Sclar et al. 1989) prior to the presentation of each different grating. The entire series of contrasts is usually presented 100 times. We analyze only spikes that occur between 30 and 300 ms after stimulus onset.

### Information rates

We use extensions of the direct method of calculating information rates (Ruyter van Steveninck et al. 1997; Strong et al. 1998) to evaluate the responses to all three types of stimulus. This method is based on a comparison of the response variability across time to the response variability across trials. The underlying principle behind the approach is that the portion of variability that cannot be explained by intrinsic variations in the response to a particular stimulus must represent stimulus-related information.

### Direct method

The straightforward application of the direct method, diagrammed in Fig. 1A, evaluates what we call the formal information rate. The spike train recorded during each trial is divided into time bins, and the spike counts in each bin are tabulated. These spike counts are considered letters in the neuron's response alphabet. Several letters in a row constitute a word, and each word has a probability, possibly stimulus-dependent, of being "spoken" by the neuron. In this paper, we choose one-letter words (single time bins) because our data sets are not typically large enough to obtain reliable multi-letter-word information estimates. Based on others' results in different systems (Reinagel and Reid 2000; Strong et al. 1998), as well as on analysis of a limited number of V1 neurons from which large amounts of data were collected, we estimate that information rates are likely to differ by at most 25% in the two cases, but that the qualitative results (formal vs. attribute-specific information and amount of confounded information) are not likely to change greatly.

We perform our calculations at a variety of bin (letter) sizes, ranging from 0.9 to 59.2 ms. We choose the bin size that yields the highest information rate. This choice is justified because the actual information rate cannot decrease as the bin size decreases (Strong et al. 1998), even though our information-rate estimates may not increase because of limitations set by the amount of available data. It is important to emphasize that the quantities estimated here are information rates for brief samples of the responses (single bins or letters), not the total information contained in an extended response. In general, the conversion between information-rate estimates and information estimates over extended responses is subadditive, in part because the information encoded at different times in the response may be redundant (DeWeese and Meister 1999).

From the set of binned spike counts, we extract two quantities. The first, called the *total entropy* ( $H_T$ ), is a measure of the response variability across time—that is, the uncertainty in spike count across all bins. We calculate  $H_T$  from the distribution of spike counts in individual bins across all time and trials (represented by the light gray rectangle in Fig. 1A). We obtain a direct estimate of the probabilities  $p_n$  that  $n$  spikes are observed directly from the spike count statistics. We then apply Shannon's formula (Cover and Thomas 1991) to obtain the entropy

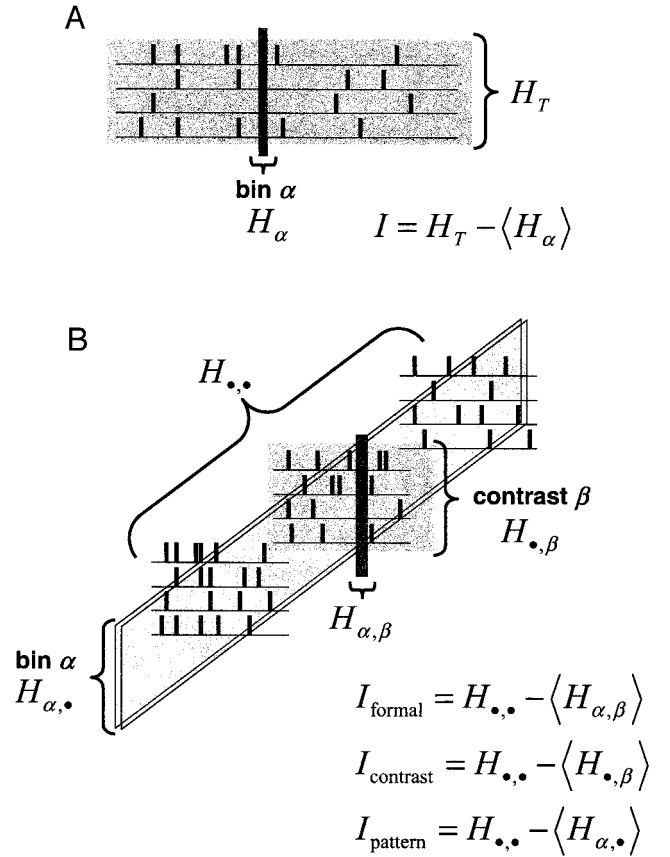


FIG. 1. *Direct method* of information-rate estimation. Stimulus-related information  $I$  is calculated as the difference between the response variability across time (*total entropy*,  $H_S$  or  $H_{..}$ ) and the response variability across trials or stimulus conditions (*noise entropy*). *A*: straightforward application of the direct method to multiple trials of a single stimulus in which the noise entropy is calculated as the average entropy in each bin ( $H_\alpha$ ) (Ruyter van Steveninck et al. 1997). *B*: modification of the direct method for calculation of attribute-specific information in an experiment with multiple contrasts. The total entropy  $H_{..}$  is unchanged, but the noise entropy depends on the particular type of information being estimated (formal or attribute-specific).

$$H = - \sum_{n=0}^{\infty} p_n \log_2 p_n \quad (1)$$

The second quantity that we extract is the *noise entropy*, which is a measure of the response variability across trials at a fixed time—that is, the uncertainty in spike count in bins at a particular time. Unlike the situation with total entropy, which is derived from all bins taken together, the number of noise-entropy estimates is equal to the number of time bins in a single trial. The set of bin-specific entropies  $\{H_\alpha\}$  is obtained from the distribution of spike counts at fixed time bins  $\alpha$  across trials (dark gray in Fig. 1A) and Shannon's formula. The transmitted information is taken to be the difference between the total entropy and the averaged noise entropy:  $I = H_S - \langle H_\alpha \rangle$ . Information values are calculated in bits, which are normalized by the bin size to obtain bits/s and by the total spike count to obtain bits/spike.

An important caveat of the direct method is that it is only sensitive to fluctuations *within* an analyzed response. If there is little variation in the local spike-count distribution during the course of a response, then the direct method yields a low information rate. Such is the case for complex cells when the stimulus is a drifting sinusoidal grating at fixed contrast. Since complex cells respond to these stimuli primarily by elevating their discharge rates (Skottun et al. 1991), the direct method only detects information if the analyzed response includes both background and stimulus-driven firing, because the appearance

of the stimulus causes a change in firing rate. In this chapter, the unit-contrast drifting-grating responses, unlike the unit-contrast stationary-grating responses, do not include any period of background firing. In the past, the direct method has always been applied to the responses of neurons to rapidly varying stimuli, which typically evoke a wide range of firing rates that fluctuate over time (Buračas et al. 1998; Reinagel and Reid 2000; Ruyter van Steveninck et al. 1997). However, there is no a priori reason to limit the application of the direct method to rapidly varying stimuli, and we will show that the method can yield useful results even when applied to other sorts of responses.

### Attribute-specific information

The formal information rate calculated by the direct method evaluates the overall rate of information transfer about all time-varying aspects of the stimulus. It does not evaluate the rates at which information about individual stimulus attributes is transmitted without the confounding influence of other attributes. We refer to these latter quantities as the attribute-specific information rates, and we now describe an extension of the direct method that allows us to estimate them. In our experiments, we concentrate on contrast and spatiotemporal pattern, but the idea can be applied to any situation in which two (or more) attributes are varied independently. As used here, spatiotemporal pattern is an omnibus term that refers to aspects of the stimulus that do not change as the contrast is varied. Our stimuli are completely defined by contrast and spatiotemporal pattern: indeed, each stimulus is defined by a spatiotemporal pattern and a contrast value by which it is multiplied.

As shown in Fig. 1B, attribute-specific information is calculated by a procedure that elaborates on the one depicted in Fig. 1A. The total entropy  $H_{\bullet,\bullet}$  is an overall quantity that represents the uncertainty in spike count across all bins in the entire data set—that is, across time, trials, and contrasts—and is used in both formal and attribute-specific information calculations. When restricted to a single contrast,  $H_{\bullet,\bullet}$  reduces to  $H_T$ , the total entropy from Fig. 1A. [We use the dot ( $\bullet$ ) notation to denote inclusion of either all time or all contrasts in the information calculations.] To obtain the overall noise entropy used in the calculation of formal information rates, we average together all the estimates of the noise entropy  $H_{\alpha,\beta}$  taken from spike counts measured on different trials at fixed times ( $\alpha$ ) and contrasts ( $\beta$ ) (dark gray region). This is a direct extension of the procedure used when only a single stimulus type is presented, as in Fig. 1A, in which  $H_{\alpha,\beta}$  is equivalent to  $H_\alpha$ .

Attribute-specific information rates are obtained by averaging over one or another stimulus attribute and over trials. Intuitively, by ignoring the value of one stimulus attribute, we are considering stimulus variations along the ignored attribute to be a source of “noise.” This potentially adds variability to the spike count and could reduce the ability of stimulus-induced variations in spike count to transmit information about the nonignored attribute. To obtain the contrast-specific noise entropy  $H_{\bullet,\beta}$ , we consider spike counts recorded at contrast  $\beta$ , regardless of time bin or trial number (medium gray), and then average those entropy estimates across contrasts. This essentially represents the uncertainty in spike count at each contrast, averaged across all contrasts. The contrast-specific information is then the difference between the total entropy and the averaged contrast-specific noise entropy:  $H_{\bullet,\bullet} - \langle H_{\bullet,\beta} \rangle$ . To obtain the spatiotemporal pattern-specific noise entropy  $H_{\alpha,\bullet}$ , we average across time bins the entropies derived from spike counts at time  $\alpha$ , regardless of contrast (light gray). This corresponds to the uncertainty in spike count at each time relative to the stimulus, averaged across all times. The spatiotemporal pattern-specific information is the difference between the total entropy and the averaged spatiotemporal pattern-specific noise entropy:  $H_{\bullet,\bullet} - \langle H_{\alpha,\bullet} \rangle$ . The sum of the two pattern-specific information rates cannot exceed the formal information except for measurement errors, and equality can only hold under circumstances in which

the two attributes are independently represented. Proof of this statement and further background concerning attribute-specific information can be found in the APPENDIX.

### Bias in the information estimates

Because we only have access to a limited amount of data, our estimates of the total and noise entropies are both subject to a downward bias. This is a generic property of information estimates from limited data sets (Carlton 1969; Miller 1955). Since the transmitted information is the difference between these two entropies, the resulting information rate will be either underestimated or overestimated depending on the relative magnitude of the bias in the two entropy estimates. When the probabilities of each word are directly estimated from the observed probabilities, an asymptotic estimate of the bias is  $(k - 1)/2N \ln(2)$ , where  $k$  is the number of distinct, observed words (here, spike counts per bin) and  $N$  is the total number of observations (Panzeri and Treves 1996; Victor 2000). Because  $N$  is large for the total entropy (number of bins times number of trials), the correction is quite small (for m-sequence responses, about 0.01%). On the other hand, in the calculation of the individual bin-specific noise entropies,  $N$  can itself be small (as low as 12 for m-sequences), making the correction much larger (sometimes on the order of 10% or more).

In many cases, particularly with short bins, only one distinct spike count—zero—is observed. These bins contribute an entropy of 0 to the averaged noise entropy, even with the asymptotic bias correction (because  $k = 1$ ). Because these bins are so common and because entropy is a logarithmic function of probability, as in Eq. 1, the noise entropy is potentially severely underestimated. This can result in an overestimation of the transmitted information. We address the problem of the zero-count bins by assuming that the noise entropy varies slowly when the number of spikes is very low. This assumption allows us to group several consecutive bins together to generate a single estimate of the bin-specific noise entropy. Specifically, when we encounter a bin with no spikes in any trial, we sequentially consider subsequent bins until we find one that has at least one spike. For these  $m$  bins, we calculate the noise entropy as described, applying the analytic bias correction with  $N$ , the number of observations, equal to  $m$  times the number of trials. We then assign this value of the noise entropy to each of the bins that are grouped together in this way. The final value of the noise entropy is again the average of the individual bin-specific entropies, where some of those entropies have been calculated by grouping several bins together. This grouping occurs most commonly in the calculation of the formal noise entropy for which the number of observations is simply equal to the number of trials. The effective number of trials in the calculation of the spatiotemporal pattern-specific noise entropy is higher because time bins are grouped together across contrasts. The effective number of trials in the calculation of the contrast-specific noise entropy is vastly higher because time bins and trials are grouped together at a single contrast so that we never encounter a bin with only one type of spike count.

We have verified that our assumption—that the noise entropy varies slowly when the number of spikes is very low—allows us to obtain accurate information-rate estimates for synthetic spike trains that are examples of modulated (inhomogeneous) Poisson processes with fewer than half the number of trials that would otherwise be required (simulations not shown). Indeed, for long stimuli (such as the m-sequences), accurate information estimates can sometimes be obtained with as few as four trials of the stimulus—we typically have at least 16. For briefer stimuli (such as the gratings), more trials are required. Because real neuronal responses are not examples of modulated Poisson processes (Reich et al. 1998), we impose an additional criterion to eliminate data sets that do not contain enough trials. In particular, we insist that the formal information rate obtained from half the data (randomly chosen) be within 10% of the corresponding information rate obtained from the full data set. This requirement is

quite strict and eliminates between 20–50% of the neurons recorded with each stimulus type. The neurons that are retained tend to convey information at slightly higher rates.

Finally, to obtain an estimate for the scatter of individual information estimates, we use the jackknife procedure (Efron 1998). Specifically, for each data set, we sequentially remove  $1/16$  of the trials and recalculate the information rates;  $1/16$  is chosen because we have access to only 16 trials for some m-sequence data sets. From the resulting distribution of information rates, we estimate the standard error of the information rate obtained from the full data set. The jackknife estimate of the standard error is

$$\sigma_J = \sqrt{\frac{N-1}{N} \sum_{i=1}^N [\hat{I}_i - \langle \hat{I} \rangle]^2} \quad (2)$$

where  $N$  is the number of trials (16) and  $\hat{I}_i$  is the  $i$ th jackknife information-rate estimate.

### Simple-cell model

To help explain our findings about formal and attribute-specific information rates—in particular, the confounded information rate (see RESULTS)—we model a V1 simple cell as a linear spatiotemporal filter the output of which is subject to a static nonlinear rectification and a Poisson spike generating mechanism (Carandini et al. 1996). Rather than assuming any particular form for the linear filter, we use first-order kernels derived from responses of real V1 neurons to unit-contrast m-sequence stimuli (Reich et al. 2000a; Reid et al. 1997; Sutter 1992); the calculation of this kernel, as well as its normalization, is extensively discussed in the references. The subset of 11 neurons modeled here includes five simple and six complex cells. Complex cells, which often yield significant linear kernels (Reich et al. 2000a), are modeled here in exactly the same way as simple cells. However, because we only use the linear kernels in the model and because the model does not include any full-wave rectification (Movshon et al. 1978a), the model neurons derived from complex cells respond like simple cells. Although we only model neurons with robust linear kernels or receptive-field maps, we verified that these neurons have firing and information rates not significantly different from the firing and information rates of the entire population of simple cells, for unit contrast responses (Kolmogorov-Smirnov test,  $P > 0.05$ ).

To derive the parameters of the rectification (threshold and linear gain), we first predict the linear response by convolving the first-order kernel with the unit-contrast m-sequence stimulus. We then find the constant offset and linear gain that, when applied to the predicted linear response histogram, yield the best least-squares fit to the histogram of the observed unit-contrast m-sequence response for the neuron being modeled. We present the model neurons with the same three stimuli that we present to real V1 neurons: m-sequences, stationary sinusoidal gratings, and drifting sinusoidal gratings. Kernels and stimuli are binned at 1.8-ms resolution (half of a single display frame for the actual visual stimuli). Because the spatiotemporal patterns in grating stimuli—particularly drifting gratings—change faster than the spatiotemporal pattern in m-sequences (even though such changes are not necessarily detected by the neuron), we integrate the gratings over each m-sequence check and time bin before presenting them to the model neuron.

The model yields a response histogram that serves as the modulation envelope of an inhomogeneous Poisson process, which is then used to determine the spike times in each trial. For the purpose of calculating information rates, the response histogram and the Poisson assumption are sufficient, because knowing the firing rate in each bin gives us the full spike-count probability distribution in that bin. Thus for model responses, we obtain an *exact* value for the information rate in one-letter words. In practice, these exact information rates are

extremely close to the ones obtained by applying the method described in the preceding text for real data; indeed, this similarity is the primary justification for the applicability of our bias-correction methods.

## RESULTS

### Formal information rates

Figure 2 shows the responses of a complex cell in monkey V1 to the three different stimuli used in this paper, each presented at unit contrast. Response histograms, obtained by averaging the number of spikes across all trials in consecutive 7.4 ms bins and then normalizing by the binwidth, are presented atop raster diagrams that show the spike times following stimulus onset in each trial. The *top panel* presents results from a flickering checkerboard stimulus, in which the time course of contrast modulation in each check is determined by a binary m-sequence (see METHODS). The stimulus lasts 60.6 s and is repeated 14 times; here, we only show spike times that occurred between 32 and 33 s after stimulus onset.

A striking feature of these responses is that spike firing tends to be clustered at particular times, which presumably follow transient changes in the stimulus. However, since the stimulus changes every 14.8 ms (67.6 times during the 1 s display period), it is clear that not all stimulus transitions are followed by a consistent change in firing probability—that is, only some of the changes, such as the one that causes firing shortly before 32.5 s, are effective in driving the neuron. This event-like firing in response to stimuli of this sort has been noted in other

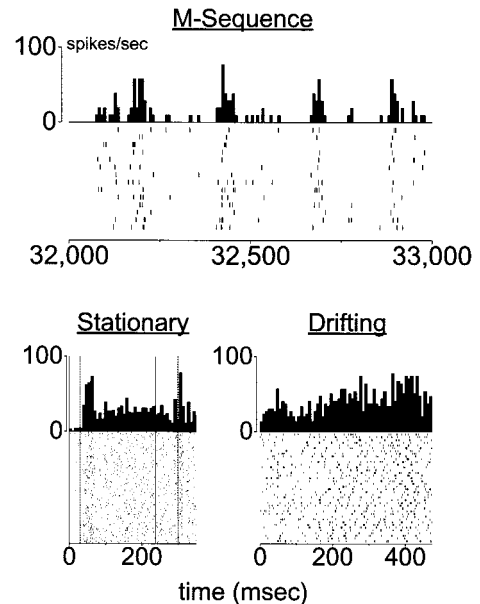


FIG. 2. Responses of a representative complex cell to three types of stimulus at unit contrast. Cell 44/9s. Each panel shows response histograms (7.4-ms bins) atop raster diagrams that depict the spike times fired during each trial. *Top*: flickering binary checkerboard pattern modulated by an m-sequence. The entire sequence lasts 60.6 s. Only spikes that occur in the 1-s period between 32 and 33 s after the start of the stimulus are shown. 14 trials. *Bottom left*: stationary sinusoidal grating presented at the cell's optimal orientation, spatial frequency, and spatial phase. The stimulus appears at *time 0* and is removed 237 ms later; solid vertical lines mark these times. Spikes occurring between 30 and 300 ms after stimulus onset (dashed vertical lines) are analyzed. 100 trials. *Bottom right*: drifting sinusoidal grating presented at the cell's optimal orientation, spatial frequency, and temporal frequency. 40 trials.

species and visual areas as well (Bair and Koch 1996; Berry et al. 1997; Ruyter van Steveninck et al. 1997). As estimated by the direct method (METHODS), the response to the m-sequence conveys 7.8 bits/s (0.75 bits/spike) of stimulus-related information, which is within the range reported both for MT neurons (Buračas et al. 1998) and for salamander and rabbit retinal ganglion cells (Berry and Meister 1998) in responses to similar stimuli.

Figure 2 also shows the responses of the same neuron to 100 presentations of a stationary sinusoidal grating of optimal orientation and spatial frequency. The grating appears at *time 0* and disappears 237 ms later (times marked by solid vertical lines). The response histogram reveals three distinct firing levels. High firing rates begin abruptly 40–45 ms after the stimulus is presented and reappear at about 310 ms, after the grating is removed. In between, the firing rate decays to a lower level that is still higher than the baseline. The high firing rates that occur in the response transients resemble, in terms of peak rate and duration, the brief periods of high firing rate in the m-sequence response. However, the neuron spends a great deal more time firing spikes at the lower rate than at the higher rate. This results in a lower information rate of 4.3 bits/s (0.14 bits/spike), despite the fact that the mean firing rate is 72% higher than for the m-sequence response (29.8 vs. 17.3 spikes/s). Although both total and noise entropy are higher in the stationary-grating response, the noise entropy, which reflects the spike count variability across trials at particular times, is subject to a proportionately greater increase than the total entropy. Thus the variability in firing is larger for the stationary-grating response than for the m-sequence response, and the stimulus-induced modulation is relatively small. Note that information-rate calculations on stationary-grating responses only involve spikes that occur from 30 to 300 ms following stimulus onset (dashed lines), meaning that the off response of this neuron is effectively ignored.

The third stimulus is a sinusoidal grating, again presented at the optimal orientation and spatial frequency, that drifts uniformly at 2.1 Hz. For this neuron, we recorded the responses to 40 cycles of the drifting grating. The analysis treats individual cycles of the grating as separate stimulus trials. As is the case for most complex cells, the response to the drifting grating is only weakly modulated (Skottun et al. 1991), and the most prominent feature is an elevation of the mean firing rate (compare the average response level in the *bottom-right panel* to the response level just after stimulus onset in the *bottom-left panel*). There are no periods of abruptly increased firing rate because, unlike the m-sequence and stationary-grating stimuli, the drifting-grating stimulus contains no temporal transients. The information rate in the drifting-grating response is 4.4 bits/s (0.11 bits/spike). This is in the top 15% of information rates measured in bits/s in our sample of 43 complex-cell responses to drifting gratings but near the median information rate measured in bits/spike. The firing-rate elevation that is such an important feature of the responses of complex cells to drifting sinusoidal gratings does not contribute to the information rate calculated by the direct method, which reflects only reproducible modulations in the spike count probability during the course of the response. Such slow modulation, superimposed on a relatively high overall firing rate of 41.5 spikes/s, is evident in the response histogram.

Figure 3 and Table 1 summarize the results of similar

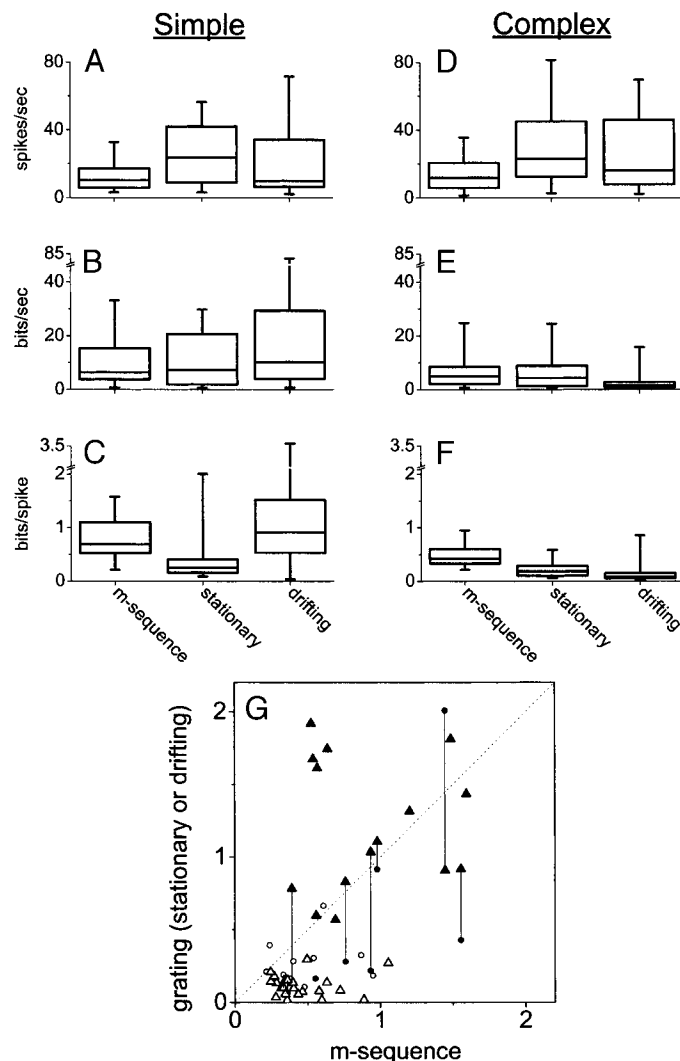


FIG. 3. Summary across V1 neurons of firing and information rates in responses to three types of stimulus. All stimuli are presented at unit contrast. Median values and population sizes are given in Table 1. A–C: simple cells. D–F: complex cells. A and D: mean firing rate, spikes/s. B and E: information rate, bits/s. C and F: information rate, bits/spike. Boxes represent the 25–75% range of the data, whiskers represent the 5–25% and 75–95% range, and horizontal lines represent the medians. G: information rates (bits/spikes) for neurons that conveyed significant amounts of information in response to at least two stimuli. Filled symbols: simple cells. Open symbols: complex cells. Circles: stationary gratings. Triangles: drifting gratings. Lines connect points corresponding to neurons that convey significant amounts of information in response to all three stimulus types. Data from individual cells, plotted in G, follow the population trends evident in A–F.

experiments performed on our entire population of V1 neurons, all at unit contrast. Panels A–F summarize across-neuron results, and G shows within-neuron comparisons for the neurons that convey significant information rates in response to at least two of the stimuli. The results reveal that differences between simple and complex cells are most pronounced in the responses to drifting gratings, which evoke the highest information rates in simple cells but the lowest information rates in complex cells. For simple cells (Fig. 3, A–C), the unexpected finding is that information rates are higher in drifting-grating responses than in m-sequence responses (Fig. 3G, filled triangles are typically above the line of equality). This is evidence against the hypothesis that high information rates in V1 neurons are

TABLE 1. Summary of unit-contrast firing and information rates

Stimulus	Type	$n$	Median		
			spikes/s	bits/s (RE)	bits/spike
m-sequence	simple	11	10.38	6.41 (0.038)	0.69*
	complex	72	11.97†	4.99 (0.053)‡	0.42§
stationary	simple	12	23.69	7.29 (0.16)	0.25*
	complex	26	23.15†	4.38‡ (0.20)	0.19§
drifting	simple	25	10.14	10.28 (0.13)	0.92*
	complex	43	16.16†	1.58‡ (0.33)	0.08§

Table gives medians of distributions plotted in Fig. 3. Median values that differ significantly between simple and complex cells are shaded gray ( $P < 0.05$ , Mann-Whitney  $U$  test). Medians of distributions that vary significantly with stimulus type are denoted by symbols (\*, †, ‡, §;  $P < 0.05$ , Kruskal-Wallis nonparametric ANOVA). RE: relative error, or standard error derived from jackknife resamplings (see METHODS) divided by information rate from the full data set.  $n$ : number of neurons in data set.

more likely to be evoked by stimuli that change rapidly in time than by stimuli that change slowly, as might be the case in the motion-sensitive area MT (Buračas and Albright 1999; Buračas et al. 1998). Complex-cell responses to drifting gratings behave differently: the filled triangles in Fig. 3G are typically below the line of equality. This is because complex cells, unlike simple cells, transmit very little information about the aspect of the drifting-grating stimulus—its spatial phase—that varies during the course of the experiment. We found (Table 1) that stimulus type has a significant effect on firing and information rates for complex cells but only on information rates in bits/spike for simple cells. However, we believe that the lack of significance for simple cells is simply due to the smaller number of simple cells in the sample.

#### Attribute-specific information rates

The direct method is usually applied only to neurons' responses to rapidly modulated stimuli, such as the unit-contrast

m-sequence responses in Figs. 2 and 3. Most stimulus sets used in neurophysiology experiments can be classified along two (or more) attributes: one (or more) features, such as contrast and spatial phase, that are explicitly varied from one stimulus presentation to the next; and the time course of the stimulus itself. For responses to these sorts of stimulus, straightforward implementation of the direct method measures only the overall rate of information transmission, which we call the formal information rate, and not information rates for individual stimulus features, which we call the attribute-specific information rates. Earlier methods, based on firing rates (Tolhurst 1989), principal components (Richmond and Optican 1990), stimulus reconstruction (Bialek et al. 1991), and time structure of individual responses (Panzeri and Schultz 2000; Victor and Purpura 1996), are expressly designed to measure attribute-specific information.

Here, we focus on contrast and spatiotemporal pattern. Formal information rates are parsed into components specific to contrast and spatiotemporal pattern. Responses to five contrasts (0.0625, 0.125, 0.25, 0.5, and 1) are analyzed. The m-sequence stimulus used in the contrast experiments is shorter, running for 7.6 s instead of 60.6 s and modulating 25 checks instead of 249.

Figure 4 shows the responses (raster diagrams and histograms) of a second complex cell, which has a maintained discharge (response to a uniform field at the mean luminance) of 2.4 spikes/s. For all three types of stimulus, the responses generally become more reproducible as contrast increases. Some features, such as the spikes that occur around 7,100 ms in the m-sequence response, or the transient firing-rate elevation at the beginning of the stationary-grating response, are particularly reproducible and give rise to peaks in the histograms.

Figure 5A plots the neuron's firing rate as a function of contrast for all three stimulus types. Clearly, the curves are

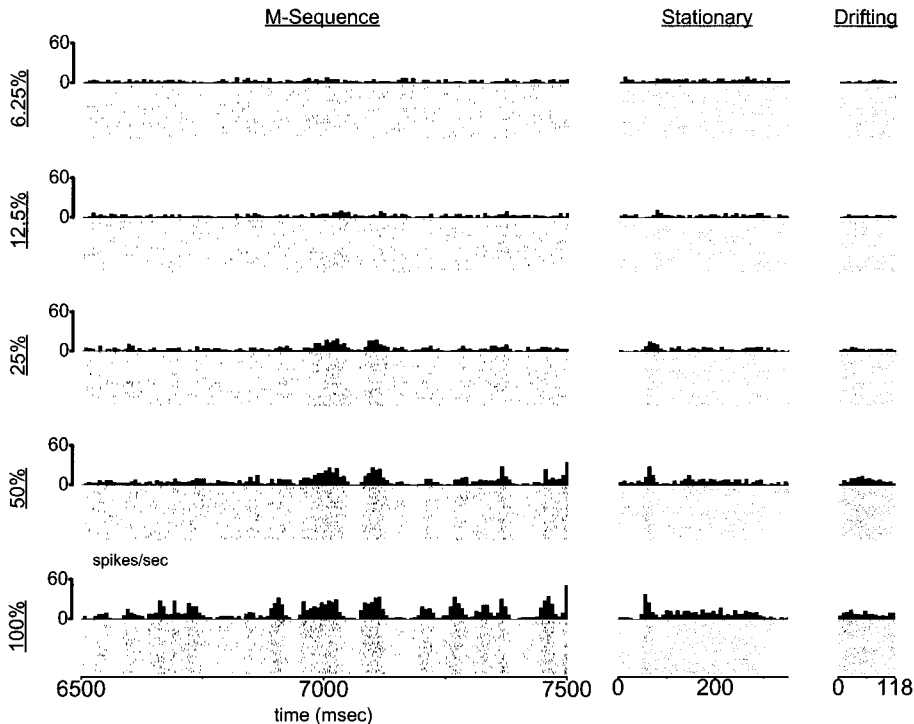


FIG. 4. Responses of a representative complex cell to three types of stimulus at five contrasts. Raster diagrams and response histograms (7.4-ms bins). Cell 45/2s. Contrasts 0.0625, 0.125, 0.25, 0.5, and 1. *Left*: 1-s snippets of the responses to 25 presentations of a 7.6-s m-sequence checkerboard stimulus. *Middle*: responses to 100 presentations of a stationary-grating stimulus, showing the first 350 ms after stimulus onset. *Right*: responses to 192 cycles of a sinusoidal grating drifting uniformly at 8.4 Hz.

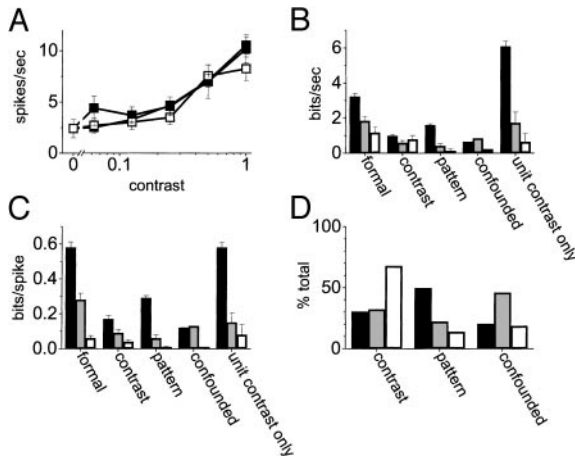


FIG. 5. Firing and information rates for the responses of the neuron from Fig. 4. *A*: mean firing rate (spikes/s) as a function of contrast: m-sequences (■), stationary sinusoidal gratings (▨), and drifting sinusoidal gratings (□). Error bars represent two standard errors of the mean. *B*: information rate (bits/s). *C*: information rate (bits/spike). Error bars in *B* and *C* represent estimates of the standard error derived from jackknife resampling (see METHODS). *D*: information rate (% of total). All 5 contrasts are used for these calculations, except for the bars labeled “unit contrast only” in *B* and *C*.

quite similar, to within the error bars (95% confidence limits of the mean). As shown in Fig. 5, *B* and *C*, the formal information rate is 3.2 bits/s (0.58 bits/spike) in the m-sequence response, 1.8 bits/s (0.28 bits/spike) in the stationary-grating response, and 1.2 bits/s (0.06 bits/spike) in the drifting-grating response. Responses to lower contrasts can contribute extra information as well as extra variability so that including them in the information calculations can both raise and lower formal information rates. In this example, including those contrasts decreases the information rate measured in bits/s (but not bits/spike) in the m-sequence responses (■), increases the information rate measured in bits/spike (but not bits/s) in the stationary-grating responses (▨), and leaves virtually unchanged the information rates in the drifting-grating responses (□).

To isolate the relative amount of information transmitted about different aspects of the stimulus, we modify the direct method by selectively changing our definition of total entropy while leaving unchanged the definition of total entropy (see METHODS). For these stimuli, information is conveyed about contrast and spatiotemporal pattern. Spatiotemporal pattern refers to aspects of the stimuli that affect the response variation across time at a fixed contrast. The results are displayed in Fig. 5, *B* and *C*, together with error estimates derived from jackknife resampling (see METHODS).

The information rates due to contrast and spatiotemporal pattern alone do not sum to the full formal information rate, as they would if the two stimulus features were encoded independently (see APPENDIX). Instead, for this neuron, the sum of the two attribute-specific information rates fails to account for 19–46% of the formal information rate, depending on stimulus type. We call the information not accounted for by the attribute-specific information rates *confounded*. The presence of confounded information means that the dynamics of contrast- and spatiotemporal pattern-encoding are interdependent. The confounded information cannot be used to determine either contrast or spatiotemporal pattern based on the response of this neuron alone.

The results of the contrast and spatiotemporal pattern experiment over the population of neurons are shown in Fig. 6 and Table 2. Here, we have combined data from simple and complex cells for m-sequence and stationary-grating responses because we could not find significant differences among the distributions, probably due in part to the limited number of simple cells in our sample. Overall, as with unit-contrast responses, drifting gratings evoke the highest formal information rates in simple cells but the lowest formal information rates in complex cells. The same is true for spatiotemporal pattern-specific information rates.

The strikingly high formal and pattern-specific information rates in simple cell responses to drifting gratings are explained by the fact that simple cells are exquisitely sensitive to spatial phase (Hubel and Wiesel 1962; Movshon et al. 1978b; Victor and Purpura 1998), which is the only aspect of the stimulus that varies at fixed contrast. The spatial-phase variation causes the firing rate to be deeply modulated, which results in high spatiotemporal pattern-specific, and hence formal, information rates. This is emphatically not the case for complex cells, as discussed above. On the other hand, simple and complex cells transmit contrast-specific information at the same rates in response to drifting gratings. Indeed, the contrast-specific information-rate distributions are relatively independent of stimulus type, measured either in bits/s or bits/spike ( $P > 0.05$ , Kruskal-Wallis ANOVA).

For the neuron presented in Figs. 4 and 5, the attribute-specific information rates do not account for all the formal information in the neuron’s response. This is also the case for the population results (Fig. 6, *right*). The median confounded information—the portion of information that cannot be resolved into contrast- or spatiotemporal pattern-specific components—represents a substantial fraction of the formal informa-

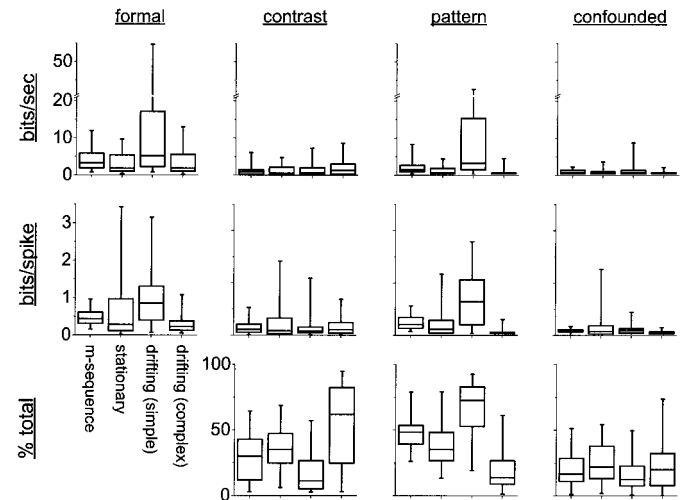


FIG. 6. Summary across V1 neurons of overall and attribute-specific information rates. Data taken from 44 neurons (m-sequences); 50 neurons (stationary gratings); 31 simple cells and 118 complex cells (drifting gratings). Data from simple and complex cell responses to m-sequences and stationary gratings are collapsed together because their distributions are not significantly different from one another. Boxes represent the 25–75% range of the data, whiskers represent the 5–25% and 75–95% range, and horizontal lines represent the medians. *Top*: information rate, bits/s. *Middle*: information rate, bits/spike. *Bottom*: attribute-specific information rate as a percentage of the formal information rate. *First column*: formal information rate. *Second column*: contrast-specific information rate. *Third column*: pattern-specific information rate. *Fourth column*: confounded information rate.



TABLE 2. Summary of five-contrast formal and attribute-specific information rates

Stimulus	Type	n	Median Information Rate										
			bits/s				bits/spike				percent		
			F	C	SP	CNF	F	C	SP	CNF	C	SP	CNF
m-sequence	all	35	3.23	0.83	1.50	0.68	0.46	0.14	0.21	0.086	30.3	45.8	16.6
stationary	all	29	1.84	0.59	0.63	0.74	0.28	0.087	0.12	0.08	35.2	35.2	22.1
drifting	simple	27	5.18	0.69	3.31	0.76	0.85	0.085	0.64	0.092	11.3	72.6	12.5
	complex	57	1.86	1.31	0.30	0.34	0.23	0.11	0.028	0.045	62.0	13.8	20.0

Values are medians of distributions plotted in Fig. 6. Simple and complex cell results were indistinguishable for m-sequence and stationary-grating responses and are therefore combined. Medians of distributions that vary significantly with stimulus type are shaded gray ( $P < 0.05$ , Kruskal-Wallis nonparametric ANOVA, simple- and complex-cell drifting-grating responses separately compared with m-sequence and stationary-grating responses).  $n$ : number of neurons in data set;  $F$ : formal;  $C$ : contrast-specific;  $SP$ : spatiotemporal pattern-specific;  $CNF$ : confounded.

tion. Indeed, across all neurons and stimulus types, the confounded information rate typically accounts for 10–32% (interquartile range) of the formal information rate.

### Information transmission in a model V1 simple cell

The model described here considers the responses of a V1 simple cell to derive from a linear spatiotemporal filter followed by a static rectifier and a Poisson spike generator (see METHODS). The shape and size of the filter, as well as the parameters of the rectifier, are drawn from the responses of actual neurons to the long m-sequence stimulus at unit contrast. Similar models have been used in the past to describe the responses of visual neurons to various kinds of stimuli (Carandini et al. 1996), even though it is well known that V1 neurons, even simple cells, display many features that are not captured in the model, including nonlinearities of spatiotemporal summation (Movshon et al. 1978a) and contrast response (Albrecht and Hamilton 1982).

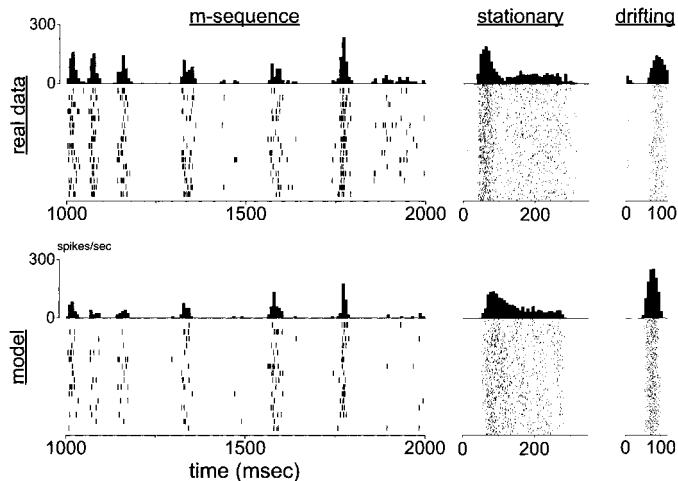


FIG. 7. Comparison of model and simple-cell responses. Cell 35/1s, unit-contrast stimuli. Model parameters are derived from the m-sequence response of the real neuron. The model consists of a linear spatiotemporal filter, a static rectifier, and a Poisson spike generator (see METHODS). Each panel shows response histograms (7.4-ms bins) atop raster diagrams. *Left*: flickering binary checkerboard pattern modulated by an m-sequence. The entire sequence lasts 60.6 s, and only the spikes that occurred in the 1 s period between 1 and 2 s after the start of the stimulus are shown. 16 trials. *Middle*: stationary sinusoidal grating presented at the cell’s optimal orientation, spatial frequency, and spatial phase. The stimulus appears at time 0 and is removed 237 ms later. 100 trials. *Right*: drifting sinusoidal grating presented at the cell’s optimal orientation, spatial frequency, and temporal frequency (8.4 Hz). 256 trials.

Figure 7 shows the responses of a real simple cell (*top*) and its corresponding model neuron (*bottom*) to the three types of stimulus, each presented at unit contrast. Despite the rudimentary nature of the model, it successfully captures many of the features of the real data. In particular, the model replicates the location of the peaks in the m-sequence response and the existence of a transient period of elevated firing rate at the beginning of the stationary-grating response. Notable differences, especially in the grating responses, do exist. For example, the transient portion of the stationary-grating response is briefer in the real data than in the model, and the drifting-grating response is narrower in the model than in the real data. Moreover the distinct, brief period of very low firing rate that immediately follows the real neuron’s transient response to the stationary grating is absent from the model’s response. These differences are not surprising given that the model’s parameters are fit to the linear part of the m-sequence response and not to the grating responses.

The similarities and differences between real and model responses are also reflected in the rates of information transmission. The m-sequence response of the real neuron transmits

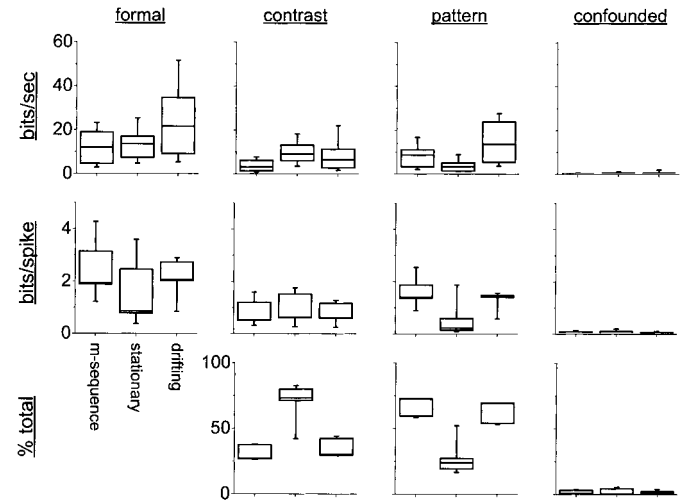


FIG. 8. Summary across 11 model neurons of information rates in the responses to three types of stimulus. Boxes represent the 25–75% range of the data, whiskers represent the 5–25% and 75–95% range, and horizontal lines represent the medians. *Top*: information rate, bits/s. *Middle*: information rate, bits/spike. *Bottom*: attribute-specific information rate as a percentage of the formal information rate. *First column*: formal information rate. *Second column*: contrast-specific information rate. *Third column*: pattern-specific information rate. *Fourth column*: confounded information rate.

TABLE 3. Summary of five-contrast formal and attribute-specific information rates

Stimulus	Median Information Rate										
	bits/s				bits/spike				percent		
	<i>F</i>	<i>C</i>	<i>SP</i>	<i>CNF</i>	<i>F</i>	<i>C</i>	<i>SP</i>	<i>CNF</i>	<i>C</i>	<i>SP</i>	<i>CNF</i>
m-sequence	13.1	3.5	9.5	0.10	1.9	0.52	1.4	0.015	27.1	72.4	0.77
stationary	13.6	9.1	3.5	0.077	0.86	0.63	0.23	0.004	73.2	26.5	0.55
drifting	21.7	6.4	13.6	0.23	2.1	0.62	1.4	0.022	30.1	69.2	1.3

Data for  $n = 11$  model neurons. Medians of distributions plotted in Fig. 8. Compare to the data in Table 2. Medians of distributions that vary significantly with stimulus type are shaded gray ( $P < 0.05$ , Kruskal-Wallis nonparametric ANOVA). Abbreviations as in Table 2.

information at a rate of 15.5 bits/s (1.6 bits/spike), compared with 19.2 bits/s (2.5 bits/spike) for the model. In response to the stationary and drifting-grating stimuli, the real neuron transmits 24.2 bits/s (0.43 bits/spike) and 49.4 bits/s (0.92 bits/spike) of information, respectively, whereas the model neuron transmits 20.8 bits/s (0.42 bits/spike) and 90.8 bits/s (1.2 bits/spike). It should be noted that the information in this real neuron's response to the drifting grating is at the top of the range of such information rates, even among simple cells (Fig. 3).

Figure 8 and Table 3 show that these results generalize to the population of 11 neurons that we modeled; the results should be compared with Fig. 6 (noting the sometimes-different vertical scales) and Table 2. Typically, the information rates of model responses are higher than the information rates of real responses. As is the case with the simple cell modeled in Fig. 7, and with real simple cells, the model neurons convey the most information about drifting gratings. The low spatiotemporal pattern-specific information rates in the modeled stationary-grating responses correspond to the low spatiotemporal pattern-specific information rates in the responses of real neurons to these stimuli. However, unlike in real neurons, the contrast-specific information rate conveyed by the model neurons does depend significantly on the type of stimulus.

The most striking difference between real and model responses is in the confounded information rate, which is nearly zero in model responses. In model responses, contrast and spatiotemporal pattern can be independently determined from the response time course and depth of modulation. The prominence of the confounded information in real responses suggests that an interaction between the coding of contrast and spatiotemporal pattern constitutes one of the primary differences between real and model neurons.

## DISCUSSION

We used three types of stimulus to evaluate the ways in which the spatiotemporal features of a stimulus affect the rates at which V1 neurons transmit information. The first stimulus is based on pseudo-random m-sequences and appears as a rapidly modulated checkerboard pattern, in which the spatial pattern changes every 14.8 ms. The second stimulus is a sinusoidal grating at the cell's optimal orientation, spatial frequency, and spatial phase; it appears abruptly and is removed 237 ms later. The third stimulus is a sinusoidal grating that drifts at the cell's optimal temporal frequency and has the same orientation and spatial frequency as the stationary grating. We find that V1 simple cells typically transmit information at the highest rates

in response to high-contrast, drifting-grating stimuli. The same stimuli evoke the lowest information rates in complex cell responses because those responses are not modulated in time.

### *Do cortical neurons transmit information at high or low rates?*

Our results are surprising in light of the fact that comparisons of information rates in a variety of neural systems suggest that stimuli that change rapidly in time drive neurons to encode information at rates often more than an order of magnitude higher than the corresponding rates for slowly changing stimuli (Buračas and Albright 1999). Based on the formal and attribute-specific information rates calculated for V1-neuron responses to different types of stimulus, it is likely that the major cause of this discrepancy is not simply the rate of variation of the stimulus, as some have argued (Buračas and Albright 1999). Moreover, the cause does not lie in the number of transient changes in a stimulus, because such transients are absent from drifting gratings (which evoke the highest information rates in simple cells) but present in both m-sequences (which evoke the highest information rates in complex cells) and stationary gratings.

Instead, the results suggest that the magnitude of a measured information rate has a complicated dependence on the type of attribute-specific information that is being estimated and the sensitivity of the neuron under study to that stimulus attribute. For example, a stimulus (such as a drifting grating) that consists of rapid changes in spatial phase evokes high spatiotemporal pattern-specific information rates in simple cells and low spatiotemporal pattern-specific information rates in complex cells, but a stimulus (such as an m-sequence checkerboard) that consists of rapid changes in luminance evokes indistinguishable information rates in simple and complex cells.

### *Information rates and channel capacity*

The information rates for complex cells can be compared with estimates of the channel capacities of complex cells in the supragranular layers of alert monkeys, which range from 6.7 to 8.5 bits/s (Wiener and Richmond 1999). The channel capacity is a measure of the maximum information rate that a communications channel can transmit (Cover and Thomas 1991). In response to the stimuli used here, which differ from the stimuli used by Wiener and Richmond, complex cells transmit information at approximately half of this estimated channel capacity. Since our stimuli were not designed to evoke information rates that approach the channel capacity, we consider this result to be rather impressive.

### *Relevance of the confounded information for visual processing*

Our use of the direct method to calculate both formal and attribute-specific information rates reveals a hitherto-overlooked aspect of the information conveyed by V1 neurons. We found that a substantial fraction of the information (typically 10–32%) cannot be attributed to either contrast or spatiotemporal pattern alone. This portion of the information, which we call *confounded*, arises from an interdependence of contrast and spatiotemporal pattern in generating neuronal responses. Confounded information is not present in the model responses, where there is no such interdependence (see APPENDIX). In other words, the amount of confounded information quantifies the effects of changes in the spatiotemporal profile of the stimulus on the contrast response and sensitivity functions of V1 neurons (Albrecht 1995; Gawne et al. 1996; Maffei and Fiorentini 1973; Tolhurst and Movshon 1975). Such changes may be mediated by variation in the “adaptive state” of a neuron under different spatiotemporal stimulus conditions, particularly between rapidly varying stimuli, like the m-sequence checkerboard, and gratings (Gaska et al. 1994).

Another way in which contrast and spatiotemporal pattern can potentially interact is through a refractory period that is intrinsic to a neuron’s spike generating mechanism but that influences the responses to both stimuli. However, we do not believe that refractory periods contribute significantly to the confounded information reported here. This is because as much or more confounded information is present in data sets that have been resampled to effectively eliminate the refractory period while preserving the overall rate modulation and distribution of spike counts per trial (Reich et al. 2000a).

Whatever its basis, the finding that a substantial portion of the total information is confounded means that downstream neurons cannot use all of the information in the responses of their inputs to draw conclusions about either one of those stimulus attributes in isolation. But perhaps the task of the visual system is not simply to decompose stimuli into components relating to contrast and spatiotemporal pattern. For example, if we had determined attribute-specific information along the visual system’s preferred axes, we might not have found any confounded information.

Alternatively, or perhaps in addition, it is possible that the messages that correspond to the confounded information would separate into contrast- and spatiotemporal pattern-specific components if the concurrent responses of other neurons to the same stimulus were considered (deCharms 1998). It is known, for example, that spikes that are synchronous across two cat LGN neurons can convey additional information beyond what can be obtained from each neuron’s individual response (Dan et al. 1998), and similar results have been obtained in various cortical systems (Maynard et al. 1999; Riehle et al. 1997; Vaadia et al. 1995). However, it is important to point out that simply averaging together the responses of redundant neurons, or even neurons that have some degree of correlated variability but identical average responses (Shadlen and Newsome 1998), would not help to disambiguate the confounded information. Thus, whether concurrent decoding of responses of a cluster of neurons can reduce the amount of confounded information is an issue that must be resolved experimentally. It is relatively

straightforward to do this by an extension of the direct method to the responses of multiple neurons recorded together.

### *What do we expect from a simple model?*

We evaluated the degree to which a simple model of V1 simple cells can replicate our experimental results. This model is quasi-linear and therefore fails to account for many of the interesting nonlinearities displayed by V1 neurons, particularly complex cells (Movshon et al. 1978a) but also, to some extent, simple cells (DeAngelis et al. 1993; Mechler et al. 1998a). In particular, the model does not account for contrast-specific nonlinearities (Albrecht and Hamilton 1982; Carandini et al. 1997b; Dean 1981). We find that this model significantly overestimates the magnitude of formal and attribute-specific information rates, in particular the contrast-specific information rates in stationary-grating responses (compare Fig. 8 to Fig. 6). Most significantly, the model responses contain no confounded information, in stark contrast to the prominent confounded information found in real responses.

In earlier work (Reich et al. 2000a), we used a spike train resampling technique to show that, for V1 neurons, the details of spike generation do not have a large effect on the magnitude of formal information rates. That result, together with additional resamplings done in connection with the present study (not shown), indicates that the discrepancies between real and model information rates (including confounded information) are not likely to be due to the assumption of a Poisson spike generating mechanism in the model. Moreover, these discrepancies are also not likely to be due to cell-to-cell variation in the shape of the linear filter or kernel, since such variation is similar for real and model responses and in both cases has little impact on information rates.

Instead the discrepancies between real and model responses almost certainly relate to the fact that real responses to stimuli that differ only in contrast are not simply related by a scaling factor but rather depend strongly on factors such as spatiotemporal pattern and the level of adaptation (Albrecht 1995; Bonds 1991; Ohzawa et al. 1982). It is possible that a single mechanism—nonlinear suppression that is sometimes called *contrast normalization* (Albrecht and Geisler 1991; Heeger 1992)—can account for all of these discrepancies but only if the mechanism is sensitive to the spatiotemporal parameters of the stimulus and can affect the dynamics of the response. Indeed, such mechanisms are known to exist in the retina (Shapley and Victor 1981), lateral geniculate nucleus (Sclar 1987), and primary visual cortex (Reid et al. 1992). The sensitivity can be intrinsic to the suppressive mechanism itself or, alternatively, might be derived from a pooling of the responses of other V1 neurons with different stimulus-response properties.

### *Summary*

The major finding in this paper is that V1 neurons transmit formal information at high rates for a variety of stimulus types and that the amount of attribute-specific information is much lower. Contrast-specific information rates depend little on stimulus and cell type, whereas spatiotemporal pattern-specific information rates depend strongly on these factors. A substantial fraction of the formal information cannot be attributed to either contrast or spatiotemporal pattern if only the responses

of single neurons are taken into account, and this confounded information is likely to be a result of dynamic interactions between stimulus attributes during response generation. Further work may determine the degree to which the confounded information can be sorted into stimulus-specific components on the basis of the simultaneous responses of groups of neurons.

#### APPENDIX

In this appendix, we make rigorous the statement that there is no confounded information if and only if a neuron's response depends independently on the two stimulus attributes (in our experiments, contrast and spatiotemporal pattern). We assume that stimuli are defined by an independent choice of a stimulus  $s_1$  out of a set  $S_1$  (corresponding to the 1st attribute) and a stimulus  $s_2$  out of a set  $S_2$  (corresponding to the 2nd attribute). A corollary of this demonstration is that there is no confounded information if and only if, for each possible response  $r$ , there is no mutual information between the conditional distributions  $\{S_1|r\}$  and  $\{S_2|r\}$ . This is in turn equivalent to the statement that, for each response  $r$ , the conditional probability  $p(s_1, s_2|r)$  is a separable function of  $s_1$  and  $s_2$ :  $p(s_1, s_2|r) = p(s_1|r)p(s_2|r)$ .

For notation, we follow the conventions of Cover and Thomas (1991). The mutual information between two variables  $X$  and  $Y$  is

$$I(X;Y) = H(X) + H(Y) - H(X,Y) \quad (A1)$$

where  $H(X)$  and  $H(Y)$  are, respectively, the entropies of  $X$  and  $Y$ , and  $H(X,Y)$  is the joint entropy of  $X$  and  $Y$ . For example, if  $x \in X$  is distributed according to the probabilities  $p(x)$ , then

$$H(X) = - \sum_{x \in X} p(x) \log p(x)$$

By the Chain Rule for entropies

$$H(X,Y) = H(X) + H(Y|X) \quad (A2)$$

where  $H(Y|X)$  denotes the conditional entropy of  $Y$  given  $X$  and is defined as

$$H(Y|X) = - \sum_{x \in X} p(x) \sum_{y \in Y} p(y|x) \log p(y|x) \quad (A3)$$

Here,  $p(y|x)$  is the conditional probability of  $y$ , given the occurrence  $x$ .

We define the confounded information  $C$  as

$$C = I(R;S_1,S_2) - I(R;S_1) - I(R;S_2) \quad (A4)$$

where  $S_i$  is the  $i$ th set of stimuli and  $R$  is the set of responses. By substitution of Eqs. A1 and A2 into Eq. A4 and using the fact that  $H(S_1,S_2) = H(S_1) + H(S_2)$  (since  $S_1$  and  $S_2$  are independent), we find that

$$C = H(S_1|R) + H(S_2|R) - H(S_1,S_2|R) \quad (A5)$$

Further substitution of Eq. A3 into Eq. A5 yields

$$C = \sum_{r \in R} p(r) \left[ - \sum_{s_1 \in S_1} p(s_1|r) \log p(s_1|r) - \sum_{s_2 \in S_2} p(s_2|r) \log p(s_2|r) + \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} p(s_1, s_2|r) \log p(s_1, s_2|r) \right]. \quad (A6)$$

Each  $p(r)$  is a probability and therefore nonnegative. The term in brackets in Eq. A6 must also be nonnegative, since it is the mutual information of  $S_1$  and  $S_2$ , given the occurrence of  $r$ , and quantities of mutual information cannot be less than zero (Cover and Thomas 1991). Thus  $C = 0$  if and only if the term in brackets in Eq. A6 is zero for every  $r$  that occurs with nonzero probability  $p(r)$ . But this is true if and only if, for every  $r$ ,  $I(S_1|r; S_2|r) = 0$ . This latter condition in turn requires that

$$p(s_1, s_2|r) = p(s_1|r)p(s_2|r) \quad (A7)$$

meaning that  $p(s_1, s_2|r)$  is a separable function of  $s_1$  and  $s_2$ .

It is straightforward to show that Eq. A7 is equivalent to our independence condition, given that the stimulus probabilities are independent, i.e.,  $p(s_1, s_2) = p(s_1)p(s_2)$ . By Bayes's rule

$$p(s_1, s_2|r) = \frac{p(r|s_1, s_2)p(s_1)p(s_2)}{p(r)} \quad (A8)$$

Combining Eqs. A7 and A8 and the fact that  $p(s_i|r)p(r) = p(r|s_i)p(s_i)$ , we find that

$$p(r|s_1, s_2) = \frac{p(r|s_1)p(r|s_2)}{p(r)}, \quad (A9)$$

precisely the independence condition.

#### Examples

We now discuss eight simple examples to illustrate the concept of confounded information. In these examples, we consider systems with two independent inputs, each of which can take a value of 0 or 1 with equal probability, and a single output. There are therefore four dif-

TABLE 4. Eight simple systems that feature two binary inputs and different output distributions

Stimuli		Responses [Conditional Probability $p(r s_1, s_2)$ ]							
$s_1$	$s_2$	1	2	3	4	5	6	7	8
0	0	0 (1)	0 (1)	1 (1)	0 (1)	0 (0.5) 1 (0.5)	0 (1)	0 (1)	0 (1)
0	1	0 (1)	1 (1)	0 (1)	0 (1)	0 (1)	0 (0.5) 1 (0.5)	1 (1)	1 (1)
1	0	1 (1)	0 (1)	0 (1)	0 (1)	0 (0.5) 1 (0.5)	0 (0.5) 1 (0.5)	1 (1)	1 (1)
1	1	1 (1)	1 (1)	1 (1)	1 (1)	1 (1)	1 (1)	2 (1)	3 (1)
<i>Information (bits)</i>									
$I(S_1;R)$		1	0	0	0.31	0.19	0.19	0.5	1
$I(S_2;R)$		0	1	0	0.31	0	0.19	0.5	1
$I(S_1, S_2;R)$		1	1	1	0.81	0.5	0.5	1.5	2
<i>CNF</i>		0	0	1	0.19	0.31	0.12	0.5	0

The conditional probability of each possible output is given in parentheses next to the output value itself. Information is calculated by the methods detailed in the APPENDIX.  $I(X;Y)$  refers to the information between two variables  $x$  and  $y$ , and *CNF* refers to the confounded information.

ferent input configurations of the stimuli  $(s_1, s_2)$ : (0,0), (0,1), (1,0), and (1,1). In the first six examples, the output is also binary, whereas in the last two examples, the output can take on more than two distinct values. The example systems are displayed in Table 4, together with the corresponding information values.

In *example 1*, the response is independent of  $s_1$  and is completely determined by  $s_2$ . In *example 2*, the response is independent of  $s_2$  and is completely determined by  $s_1$ . Of the 2 bits of entropy in the stimuli, only 1 bit is conveyed as information, but it is transmitted perfectly. Moreover, the response to one stimulus is independent of the value of the other, so that there is no confounded information.

In *example 3*, the system responds with a 1 if the two inputs are identical, and with a 0 if they are different. There is still 1 bit of information transmitted, but in this case the response to one stimulus depends completely on the value of the second, so that all the information is confounded—that is, there is no information transmitted about either stimulus unless the value of the other stimulus is known.

In *example 4*, the system responds only if the values of both stimuli are 1. Here, the response is symmetric in the two stimuli, so that the response to each stimulus alone conveys the same amount of information. However, because the response depends jointly on the values of both stimuli, some of the information—23% of the total—is confounded.

In the fifth and sixth examples, the system's response again has a complicated dependence on the two stimuli, so that the confounded information is nonzero. In *example 5*, the system responds at random if  $s_2 = 0$ , and identically reflects  $s_1$  if  $s_2 = 1$ . Confounded information arises because although no information is conveyed about  $s_2$ , the response to  $s_1$  is more informative if  $s_2$  is known. In *example 6*, the system responds at random if the two stimuli are different and reflects their shared value if they are the same. The response thus conveys equal amounts of information about the two stimuli, but there is still some confounded information.

The first six examples illustrate that confounded information can arise even in very simple systems, so long as the conditional response probabilities are separable in the two stimuli, as in Eq. A9. The last two examples demonstrate that this requirement is not equivalent to a requirement that stimulus encoding be linear or additive. *Example 7* is a system that simply sums the values of the two stimuli (and thus has three possible responses); the confounded information in this system is 33% of the total. On the other hand, *example 8* is a system that generates distinct responses to each of the four stimuli, but its response is also additive in the sense that the response to any pair of stimuli is the sum of the responses to two pairs of stimuli that add up to the same input. For instance, the response to (1,1) is equal to the sum of the responses to (1,0) and (0,1). As with any system that maps each input to a distinct output, even a system that is not additive at all, the system in *example 8* does not produce confounded information.

Thus it is not surprising that we observe substantial amounts of confounded information in real neuronal responses: real responses depend in a complicated way on both contrast and spatiotemporal pattern. In the simple examples considered here, in fact, only systems that ignore one or the other stimulus, as in *examples 1* and *2*, or systems that respond differently to each stimulus pair, as in *example 8*, can feature responses that convey information about both stimuli without confounding them.

We thank B. Knight and K. Purpura for much useful advice.

This work was supported by National Institutes of Health Grants GM-07739 and EY-07138 (D. S. Reich) and EY-9314 (J. D. Victor).

## REFERENCES

- ALBRECHT DG. Visual cortex neurons in monkey and cat: effect of contrast on the spatial and temporal phase transfer functions. *Vis Neurosci* 12: 1191–1210, 1995.

- ALBRECHT DG AND GEISLER WS. Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Vis Neurosci* 7: 531–546, 1991.
- ALBRECHT DG AND HAMILTON DB. Striate cortex of monkey and cat: contrast response function. *J Neurophysiol* 48: 217–237, 1982.
- BAIR W AND KOCH C. Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput* 8: 1185–1202, 1996.
- BERRY MJ II AND MEISTER M. Refractoriness and neural precision. *J Neurosci* 18: 2200–2211, 1998.
- BERRY MJ II, WARLAND DK, AND MEISTER M. The structure and precision of retinal spike trains. *Proc Natl Acad Sci USA* 94: 5411–5416, 1997.
- BIALEK W, RIEKE F, RUYTER VAN STEVENINCK RR DE, AND WARLAND D. Reading a neural code. *Science* 252: 1854–1857, 1991.
- BONDS AB. Temporal dynamics of contrast gain in single cells of the cat striate cortex. *Vis Neurosci* 6: 239–255, 1991.
- BURAČAS GT AND ALBRIGHT TD. Gauging sensory representations in the brain. *Trends Neurosci* 22: 303–309, 1999.
- BURAČAS GT, ZADOR AM, DEWEESE MR, AND ALBRIGHT TD. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 20: 959–969, 1998.
- CARANDINI M, BARLOW HB, O'KEEFE LP, POIRSON AB, AND MOVSHON JA. Adaptation to contingencies in macaque primary visual cortex. *Philos Trans R Soc Lond B Biol Sci* 352: 1149–1154, 1997a.
- CARANDINI M, HEEGER DJ, AND MOVSHON JA. Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci* 17: 8621–8644, 1997b.
- CARANDINI M, MECHLER F, LEONARD CS, AND MOVSHON JA. Spike train encoding by regular-spiking cells of the visual cortex. *J Neurophysiol* 76: 3425–3441, 1996.
- CARLTON AG. On the bias of information estimates. *Psychol Bull* 71: 108–109, 1969.
- COVER TM AND THOMAS JA. *Elements of Information Theory*. New York: Wiley, 1991.
- DAN Y, ALONSO JM, USREY WM, AND REID RC. Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature Neurosci* 1: 501–507, 1998.
- DEAN AF. The relationship between response amplitude and contrast for cat striate cortical neurones. *J Physiol (Lond)* 318: 413–427, 1981.
- DEANGELIS GC, GHOSE GM, OHZAWA I, AND FREEMAN RD. Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons. *J Neurosci* 19: 4046–4064, 1999.
- DEANGELIS GC, OHZAWA I, AND FREEMAN RD. Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. II. Linearity of temporal and spatial summation. *J Neurophysiol* 69: 1118–1135, 1993.
- DECHARMS RC. Information coding in the cortex by independent or coordinated populations. *Proc Natl Acad Sci USA* 95: 15166–15168, 1998.
- DEWEESE MR AND MEISTER M. How to measure the information gained from one symbol. *Network* 10: 325–340, 1999.
- EFRON B AND TIBSHIRANI RJ. *An Introduction to the Bootstrap*. Boca Raton, FL: Chapman & Hall/CRC Press, 1998.
- GASKA JP, JACOBSON LD, CHEN HW, AND POLLEN DA. Space-time spectra of complex cell filters in the macaque monkey: a comparison of results obtained with pseudowhite noise and grating stimuli. *Vis Neurosci* 11: 805–821, 1994.
- GAWNE TJ, KJAER TW, AND RICHMOND BJ. Latency: another potential code for feature binding in striate cortex. *J Neurophysiol* 76: 1356–1360, 1996.
- GERSHON ED, WIENER MC, LATHAM PE, AND RICHMOND BJ. Coding strategies in monkey V1 and inferior temporal cortices. *J Neurophysiol* 79: 1135–1144, 1998.
- HEEGER DJ. Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9: 181–197, 1992.
- HELLER J, HERTZ JA, KJAER TW, AND RICHMOND BJ. Information flow and temporal coding in primate pattern vision. *J Comput Neurosci* 2: 175–193, 1995.
- MAFFEI L AND FIORENTINI A. The visual cortex as a spatial frequency analyser. *Vision Res* 13: 1255–1267, 1973.
- MAYNARD EM, HATSPOULOS NG, OJAKANGAS CL, ACUNA BD, SANES JN, NORMANN RA, AND DONOGHUE JP. Neuronal interactions improve cortical population coding of movement direction. *J Neurosci* 19: 8083–8093, 1999.
- MECHLER F, REICH DS, AND VICTOR JD. Receptive field nonlinearities in macaque primary visual cortex (V1) are not selectively activated by edges. *Soc Neurosci Abstr* 24: 1979, 1998a.
- MECHLER F, VICTOR JD, PURPURA KP, AND SHAPLEY R. Robust temporal coding of contrast by V1 neurons for transient but not for steady-state stimuli. *J Neurosci* 18: 6583–6598, 1998b.

- MILLER GA. Note on the bias on information estimates. *Information Theory in Psychology: Problems and Methods* II-B: 95–100, 1955.
- MOVSHON JA, THOMPSON ID, AND TOLHURST DJ. Receptive field organization of complex cells in the cat's striate cortex. *J Physiol (Lond)* 283: 79–99, 1978a.
- MOVSHON JA, THOMPSON ID, AND TOLHURST DJ. Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J Physiol (Lond)* 283: 53–77, 1978b.
- OHZAWA I, SCLAR G, AND FREEMAN RD. Contrast gain control in the cat visual cortex. *Nature* 298: 266–268, 1982.
- PANZERI S AND SCHULTZ SR. A unified approach to the study of temporal, correlational and rate coding. *Neural Comput.* In press.
- PANZERI S AND TREVES A. Analytical estimates of limited sampling biases in different information measures. *Network* 7: 87–100, 1996.
- PURPURA K, CHEE-ORTS MN, AND OPTICAN LM. Temporal encoding of texture properties in visual cortex of awake monkey. *Soc Neurosci Abstr* 19: 315.15, 1993.
- REICH DS, MECHLER F, PURPURA KP, AND VICTOR JD. Interspike intervals, receptive fields, and information encoding in primary visual cortex. *J Neurosci* 20: 1964–1974, 2000a.
- REICH DS, MECHLER F, AND VICTOR JD. Information rates in the responses of V1 neurons to temporally rich stimuli. *Invest Ophthalmol Vis Sci (Suppl)* 41: 53, 2000b.
- REICH DS, VICTOR JD, AND KNIGHT BW. The power ratio and the interval map: spiking models and extracellular recordings. *J Neurosci* 18: 10090–10104, 1998.
- REID RC, VICTOR JD, AND SHAPLEY RM. Broadband temporal stimuli decrease the integration time of neurons in cat striate cortex. *Vis Neurosci* 9: 39–45, 1992.
- REID RC, VICTOR JD, AND SHAPLEY RM. The use of m-sequences in the analysis of visual neurons: linear receptive field properties. *Vis Neurosci* 14: 1015–1027, 1997.
- REINAGEL P AND REID RC. Temporal coding of visual information in the thalamus. *J Neurosci* 20: 5392–5400, 2000.
- RICHMOND BJ AND OPTICAN LM. Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. II. Information transmission. *J Neurophysiol* 64: 370–380, 1990.
- RIEHLE A, GRUN S, DIESMANN M, AND AERTSEN A. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278: 1950–1953, 1997.
- RUUYTER VAN STEVENINCK RR DE, LEWEN GD, STRONG SP, KOBERLE R, AND BIALEK W. Reproducibility and variability in neural spike trains. *Science* 275: 1805–1808, 1997.
- SCLAR G. Expression of “retinal” contrast gain control by neurons of the cat's lateral geniculate nucleus. *Exp Brain Res* 66: 589–596, 1987.
- SCLAR G, LENNIE P, AND DEPRIEST DD. Contrast adaptation in striate cortex of macaque. *Vision Res* 29: 747–755, 1989.
- SHADLEN MN AND NEWSOME WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18: 3870–3896, 1998.
- SHAPLEY RM AND VICTOR JD. How the contrast gain control modifies the frequency responses of cat retinal ganglion cells. *J Physiol (Lond)* 318: 161–179, 1981.
- SKOTTUN BC, DE VALOIS RL, GROSOF DH, MOVSHON JA, ALBRECHT DG, AND BONDS AB. Classifying simple and complex cells on the basis of response modulation. *Vision Res* 31: 1079–1086, 1991.
- STRONG SP, KOBERLE R, RUUYTER VAN STEVENINCK RR DE, AND BIALEK W. Entropy and information in neural spike trains. *Phys Rev Lett* 80: 197–200, 1998.
- SUTTER EE. A deterministic approach to nonlinear systems analysis. In: *Nonlinear Vision: Determination of Neural Receptive Fields, Function, and Networks*, edited by Pinter RB and Nabet B. Boca Raton, FL: CRC, 1992, 171–220.
- TOLHURST DJ. The amount of information transmitted about contrast by neurones in the cat's visual cortex. *Vis Neurosci* 2: 409–413, 1989.
- TOLHURST DJ AND MOVSHON JA. Spatial and temporal contrast sensitivity of striate cortical neurones. *Nature* 257: 674–675, 1975.
- VADIA E, HAALMAN I, ABELES M, BERGMAN H, PRUT Y, SLOVIN H, AND AERTSEN A. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* 373: 515–518, 1995.
- VICTOR JD. Asymptotic bias in information estimates and the exponential (Bell) polynomials. *Neural Comput.* In press.
- VICTOR JD AND PURPURA KP. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J Neurophysiol* 76: 1310–1326, 1996.
- VICTOR JD AND PURPURA KP. Spatial phase and the temporal structure of the response to gratings in V1. *J Neurophysiol* 80: 554–571, 1998.
- WIENER MC AND RICHMOND BJ. Using response models to estimate channel capacity for neuronal classification of stationary visual stimuli using temporal coding. *J Neurophysiol* 82: 2861–2875, 1999.