

Natural Stimulation of the Nonclassical Receptive Field Increases Information Transmission Efficiency in V1

William E. Vinje^{1,2} and Jack L. Gallant^{1,3}

¹Neuroscience Program and Departments of ²Molecular and Cellular Biology and ³Psychology, University of California at Berkeley, Berkeley, California 94720-1650

We have investigated how the nonclassical receptive field (nCRF) affects information transmission by V1 neurons during simulated natural vision in awake, behaving macaques. Stimuli were centered over the classical receptive field (CRF) and stimulus size was varied from one to four times the diameter of the CRF. Stimulus movies reproduced the spatial and temporal stimulus dynamics of natural vision while maintaining constant CRF stimulation across all sizes. In individual neurons, stimulation of the nCRF significantly increases the information rate,

the information per spike, and the efficiency of information transmission. Furthermore, the population averages of these quantities also increase significantly with nCRF stimulation. These data demonstrate that the nCRF increases the sparseness of the stimulus representation in V1, suggesting that the nCRF tunes V1 neurons to match the highly informative components of the natural world.

Key words: information theory; nonclassical receptive field; V1; sparse coding; efficiency; natural vision

The classical receptive field (CRF) of a visual neuron is traditionally defined as the region of space where stimuli evoke action potentials. Surrounding the CRF is the nonclassical receptive field (nCRF), where stimuli can modulate the responses evoked by CRF stimulation (Allman et al., 1985). The nCRF may serve to mediate contrast gain control through divisive modulation of the responses evoked by CRF stimulation (Heeger, 1992; Wilson and Humanski, 1993). However, several experiments suggest that the nCRF may also be critical for representing extended contours (Gilbert and Wiesel, 1990; Fitzpatrick, 2000), corners (Sillito et al., 1995), or local curvature (Wilson and Richards, 1992; Krieger and Zetzsche, 1996), and may aid in figure-ground segmentation (Knierim and Van Essen, 1992). Together, these results demonstrate that the nCRF plays an important role in the functioning of V1 neurons.

In a previous study, we showed that natural stimulation of the nCRF increases the selectivity of V1 neurons and decorrelates their responses (Vinje and Gallant, 2000). Those results suggested that nCRF stimulation increases the sparseness of stimulus representation in V1. Sparseness refers to the coding density of a neural representation. In a maximally dense representation, every neuron responds to every stimulus and information is fully distributed across the population. In a maximally sparse representation, each neuron responds to a single stimulus and acts as a “grandmother cell.” Extremely dense and extremely sparse codes

are biologically implausible; any real neural code will fall somewhere between these two extremes.

In a sparse representation, neurons are narrowly tuned and relatively few are active at any moment. A central tenet of sparse coding is that information should be translated without loss into an efficient representation where the responses of a few active neurons are rich in information content. Reducing the number of active neurons is metabolically economical, thus easing a major constraint on information processing in the brain (Laughlin et al., 1998; Sibson et al., 1998). In addition, the relatively large information content per neuron potentially influences many aspects of brain function, including pattern recognition capability and memory capacity (Barlow, 1961; 2001).

The optimal level of sparseness is a function of the goals of the system and the resources available. Recent theoretical work suggests that natural images can be efficiently represented by a sparse code (Srinivasan et al., 1982; Barlow, 1989; Field, 1993; Bell and Sejnowski, 1997; Olshausen and Field, 1997, 2000; Simoncelli and Olshausen, 2001). Field (1987) demonstrated that linear filters can produce a highly kurtotic, sparse output distribution in response to natural images. However, some nCRF functions might only be realizable with nonlinear mechanisms [e.g., biologically plausible curvature/corner detectors must be substantially nonlinear (Zetzsche and Barth, 1990; Krieger and Zetzsche, 1996)]. Therefore, nonlinear operations such as those implemented by the nCRF are likely to play an important role in increasing the sparseness of neural coding (Olshausen and Field, 1997).

The hypothesis that nCRF stimulation increases the sparseness of individual V1 neurons leads to numerous predictions. Several of these predictions were confirmed in a previous report (Vinje and Gallant, 2000). As sparseness increases, individual neurons become more selective in their responses to complex stimuli, the kurtosis of the firing rate distribution increases, and the responses of neuron pairs are decorrelated.

The hypothesis that nCRF stimulation increases sparseness also leads to four additional predictions. First, the average response rate should decrease as sparseness increases in order to

Received Sept. 21, 2001; revised Jan. 9, 2002; accepted Jan. 9, 2002.

This work was supported by grants from the Whitehall and Sloan Foundations and from the National Eye Institute (J.L.G.). W.E.V. was partially supported by a National Institutes of Health training grant. We thank Rob de Ruyter van Steveninck and Daniel Reich for the encouragement to undertake an information theoretic analysis of nCRF function and for valuable discussions regarding information theory. We also thank Kathleen Bradley, Stephen David, Kate Gustavsen, Ben Hayden, James Mazer, and Scott Perkins for infrastructure support and many helpful comments.

Correspondence should be addressed to Jack L. Gallant, University of California at Berkeley, 3210 Tolman Hall #1650, Berkeley, CA 94720-1650. E-mail: gallant@socrates.berkeley.edu.

Copyright © 2002 Society for Neuroscience 0270-6474/02/222904-12\$15.00/0

reduce the metabolic demands of visual processing. Second, the reduction in spiking activity should not reduce the information carried by the population of V1 neurons. Third, the average information content per spike should increase. Finally, as sparseness increases, individual neurons should become more efficient at information transmission.

MATERIALS AND METHODS

Subjects and physiological procedures. All animal procedures were approved by oversight committees at the University of Washington (St. Louis, MO) and the University of California at Berkeley and conformed to or exceeded all relevant National Institutes of Health and United States Department of Agriculture standards. Surgical procedures were conducted under appropriate anesthesia using standard sterile techniques (Connor et al., 1997).

Extracellular, single-neuron recordings were made with epoxy-coated tungsten electrodes (AM Systems, Everett, WA and FHC, Bowdoinham, ME) from two awake, behaving monkeys (*Macaca mulata*). Signals were amplified, band-pass filtered, and isolated with a hardware window discriminator. Spike triggers were monitored at 8 kHz. Only clearly isolated single units were included in the data set.

Chambers were located over putative V1 by means of external cranial landmarks. To confirm that recordings were obtained from V1 neurons, we compared measured receptive field sizes and electrophysiological response properties with those expected from the literature.

Receptive-field estimation. The boundaries of the CRF were estimated using bars and gratings for which characteristics and placement were manually controlled. We estimated the size of the CRF as the diameter of the circle that circumscribed the minimum response field of the neuron. For most neurons, these manual estimates were confirmed by reverse correlation analysis using a dynamic (72 Hz) sequence of small white squares flashed randomly in and around the CRF. Reliable CRF estimates were typically obtained from 100–300 sec of data, representing 20–60 behavioral fixation trials. In most cases there was excellent agreement between CRF profiles estimated using the two methods. In those cases in which the methods disagreed, the reverse correlation size estimates were used. CRF diameters ranged from ~20 to 50 min of arc, consistent with other studies (Snodderly and Gur, 1995).

Simulated eye-movement model. During natural vision, primates make stereotyped eye movements consisting of relatively long, stable fixations interspersed with rapid saccades from one point to another (Keating and Keating, 1982; Burman and Seagraves, 1994). The temporal structure of natural visual stimulation is strongly influenced by these underlying eye movements. We simulated natural macaque eye movements using a statistical model. Eye-movement distributions were acquired during free-viewing experiments using a scleral search coil. These data were used to model the distribution of saccade lengths and the velocity profiles appropriate for each saccade. For each simulated eye-movement sequence, fixation durations were chosen randomly from a Gaussian distribution with a mean of 350 msec and an SD of 50 msec. Saccade directions were chosen randomly from a uniform distribution of angles.

Natural-vision movies. Natural-vision movies were constructed by extracting image patches from natural scenes along the simulated eye-scan path. Scenes were chosen from a commercial, high-resolution photo-CD image library of landscapes, structures, people, and animals (Corel Corp., Ottawa, Ontario, Canada) and were converted to grayscale before display. Image patches were extracted along a simulated scan path that was sampled at ~1 kHz. Each 13.8 msec (72 Hz) movie frame was constructed by averaging 14 separate image patches. Individual frames were then concatenated to form movies. This over-sampling followed by averaging minimized the potential of introducing temporal aliasing artifacts into the movie.

Patches of one, two, three, or four times the diameter of the CRF were used to create a set of natural-vision movies. Movies of different sizes were not scaled versions of one another. Instead, the patch boundary was changed to reveal more or less of the underlying natural scene. Thus, the region of the natural vision movie covering the CRF was identical across all movie sizes, and any response modulation attributable to stimulus size should reflect the effects of nCRF stimulation. Figure 1 illustrates the stimulus generation method.

Flashed natural-image patches. An additional stimulus set was constructed by extracting image patches from along an eye-scan path that was recorded during free viewing of natural scenes (Gallant et al., 1998). Eye positions during fixations were identified using an automated pro-




Figure 1. Natural-vision movies reproduce the stimulation that occurs during free viewing of natural scenes. To construct a natural vision movie, a saccadic scan path (white line) is generated using a model derived from previously recorded eye movements. Image patches centered on the scan path coordinates (white circles) are then extracted from the underlying image. Image patches were from one to four times the size of the CRF. (The small circle indicates $1 \times$ CRF diameter, whereas the large circle indicates $4 \times$ CRF diameter.) Note that although nCRF stimulation varied substantially with stimulus size, the stimulus falling on the CRF was the same for all sizes.

cedure that registered a fixation whenever the eye remained within a 0.3 CRF diameter window for at least 70 msec; a change in fixation was registered when the eye moved >0.3 receptive field diameters from its original location. These fixation locations were used as center points for patch extraction, and patches of $1 \times$ CRF and $3 \times$ CRF diameter were extracted from the natural scene in the manner described above. Each patch data set contained responses from 10–25 such patches.

The image patches were presented in grayscale under behavioral conditions similar to those used for natural vision movies (see below). Patches were shown at either the same size as the estimated CRF or three times larger than the CRF. Each behavioral trial included four random patches flashed for 500 msec each and separated by 700 msec interstimulus intervals.

Stimulus presentation. Stimulus presentation and behavioral control were handled by an Indigo² workstation (SGI, Mountain View, CA) using custom software. Stimuli were presented on a high-quality video monitor (Sony Trinitron; Sony, Tokyo, Japan) at 1280×1024 pixel resolution. Movies were broken into 5 sec segments (trials) and were shown centered on the CRF center of the recorded neuron. During movie display, the animal fixated on a small target spot near the center of the monitor. Eye position was monitored using a scleral search coil, and trials were aborted if the eye deviated from fixation by $>0.35^\circ$. At the end of each successful trial, the animal earned a liquid reward. Only one stimulus size was shown on each trial; stimuli of different sizes were randomly interleaved across trials.

Response modulation ratio. The nCRF modulation produced by stimuli of a given diameter is quantified in terms of the response modulation ratio:

$$R_i^{m \times \text{CRF}} = \frac{\langle r_i^{m \times \text{CRF}} \rangle}{\langle r_i^{\text{CRF}} \rangle}. \quad (1)$$

In our analysis, the fundamental quantity of interest is the average number of action potentials occurring in each time bin of the natural-vision movie. In Equation 1, $\langle r_i^{\text{CRF}} \rangle$ is the average response recorded during the i th time bin for stimuli confined to the CRF, and $\langle r_i^{m \times \text{CRF}} \rangle$ is the average response recorded during the i th time bin for stimuli m times the diameter of the CRF. Responses are averaged across repeated stimulus presentation trials.

Selectivity index for natural-vision movies. We define a selectivity index based on the responses of a neuron across a stimulus set:

$$S^{m \times \text{CRF}} = \frac{1 - \left\{ \frac{\left(\sum_{i=1}^n (r_i^{m \times \text{CRF}})^2 \right)^2}{n \left(\sum_{i=1}^n (r_i^{m \times \text{CRF}})^2 \right)^2} \right\}}{1 - \left(\frac{1}{n} \right)} = \frac{1 - \left\{ \frac{\mu^2}{\mu^2 + \sigma^2} \right\}}{1 - \left(\frac{1}{n} \right)} \quad (2)$$

Here μ is the mean response of the cell, σ is its SD, and the number of time bins is given by n .

The terms in braces define the activity fraction of the neuron across the stimulus set (Tovee et al., 1993). It is easy to anticipate the asymptotic behavior of the activity fraction (consider the expanded form of the activity fraction in the middle expression of Eq. 2). If a neuron were nonselective, then $r_i^{m \times \text{CRF}}$ would be constant across stimuli and the numerator and denominator of the activity fraction would be equal. In contrast, if a neuron responded to only the k th stimulus then the numerator would be given by $(r_k^{m \times \text{CRF}})^2$, whereas the denominator would be larger by a factor of n , $n(r_k^{m \times \text{CRF}})^2$. Thus, the activity fraction ranges from 1, when the cell is nonselective, to $1/n$, when the cell responds to a single stimulus frame.

Equation 2 rescales the activity fraction so that it conveniently ranges from 0 to 1. S will be 0 if a neuron is completely nonselective and 1 if it responds only to a single stimulus. For convenience, we express S as a percentage. In a previous publication, the selectivity index was referred to as the sparseness index (Vinje and Gallant, 2000). In this paper, we use sparseness as an adjective describing how stimuli are represented by sensory neurons; therefore, increasing sparseness should produce numerous effects, including increasing selectivity.

Information transmission in sensory neurons. From the perspective of information theory, an axon is a biological communication channel. Consider an observer who is monitoring the axon of a sensory neuron with known filtering properties. Before the neuron responds, the observer is uncertain about the nature of the stimulus. After observing the responses of the neuron, the observer can determine the overlap between the stimulus and the neural filtering properties. Thus, the response of the neuron reduces the observer's uncertainty about the stimulus. The amount by which a response reduces uncertainty is referred to as the mutual information carried between stimulus and response.

The total stimulus entropy, $H(s)$, quantifies the observer's uncertainty regarding the stimulus before the response of the neuron is observed. The conditional stimulus entropy $H(s|r)$, gives the stimulus uncertainty that remains after response observation. If the response is reliably influenced by the stimulus, then the conditional stimulus entropy will necessarily be less than the total stimulus entropy. The transmitted mutual information is given by (Cover and Thomas, 1991):

$$I(s, r) = H(s) - H(s|r) \quad (3)$$

The stochastic nature of spike generation means that neural responses are variable even when a stimulus is repeated exactly. The response variability caused by noise (noise entropy) limits the amount of information that can be transmitted about a stimulus. Figure 2A illustrates the relationship between stimulus entropy, noise entropy, and mutual information. The reduction in stimulus uncertainty is equal to the reduction in uncertainty regarding the response:

$$I(s, r) = H(s) - H(s|r) = H(r) - H(r|s) \quad (4)$$

Here, $H(r)$ is the total response entropy, which quantifies the overall variability of the responses of a neuron across the stimulus ensemble. $H(r|s)$ is the conditional response entropy, describing the average variability in responses evoked by a single stimulus. The conditional response entropy is equivalent to the noise entropy. In practice it is often easier to evaluate response entropies than stimulus entropies.

Calculation of total response entropy and conditional response entropy. It is straightforward to compute the total response entropy via the direct method (de Ruyter van Steveninck et al., 1997; Borst and Theunissen, 1999; Reich et al., 2000). All direct information estimation methods begin by translating the spike train into discrete words that represent local spike patterns. The choice of translation process is equivalent to choosing a hypothesis about how neurons encode and decode information. The detailed nature of the encoding/decoding process is still unre-

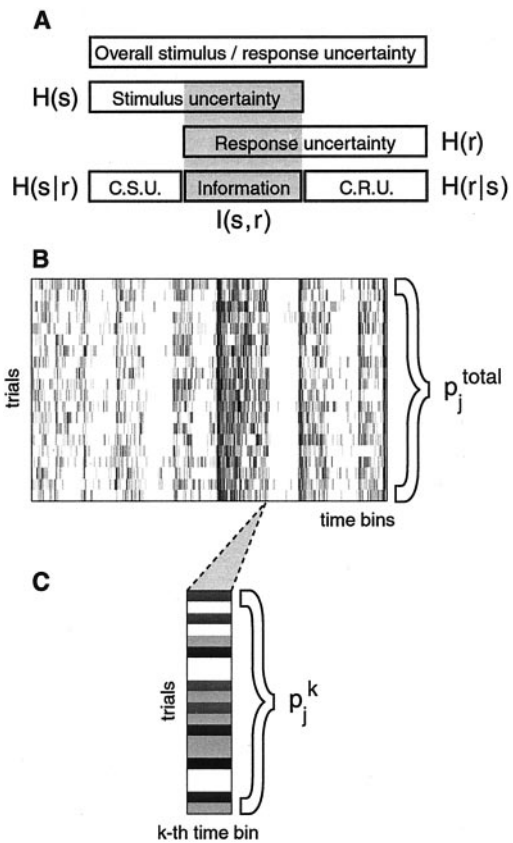


Figure 2. Sensory neurons transmit information when their responses allow an observer to reduce uncertainty regarding the nature of the stimulus. *A*, Diagram illustrating relationship between total uncertainty present in the set of all stimulus–response pairings for a given neuron; the *second rectangle* represents the observer's *a priori* uncertainty about the stimuli in a natural-vision movie; the *third rectangle* represents the uncertainty in the observed responses of the neuron. These uncertainties can be translated into entropies by means of Equation 5. The single number that summarizes overall stimulus uncertainty is the total stimulus entropy, $H(s)$, while the total response entropy is $H(r)$. The remaining *rectangles* are the conditional stimulus uncertainty (C.S.U.) and the conditional response uncertainty (C.R.U.) (quantified by the entropies $H(s|r)$ and $H(r|s)$, respectively). The *gray-shaded region* denotes correlations between the stimulus and the responses of the neuron; this correlation is what allows information, $I(s, r)$, to be transmitted. If every stimulus evokes a unique and repeatable response, then response uncertainty will be entirely determined by stimulus uncertainty. In this case the gray-shaded region would completely overlap both stimulus and response uncertainties. In real neurons, repeated presentation of a stimulus produced a range of responses, so $H(r) > I(s, r)$. The remaining uncertainty, $H(r|s)$, is attributable to noise in the encoding and transmission process. *B*, Grayscale rastergram of single neuron responses to repeated movie presentations. *Rows* represent repeated presentations of the movie, whereas *columns* represent individual time bins. Each time bin contains a single response word whose identity is determined by the number of action potentials (identity is indicated by the shading of each bin). The total response entropy, $H(r)$, is a function of the frequency with which each word is observed, p_j^{total} . *C*, Magnified view of responses to one stimulus repeated 20 times. Variation in the identity of the response words is clearly visible across trials and is quantified as noise entropy, $H(r|s = k)$. Noise entropy is a function of the probability that each word occurs in response to the k th stimulus, p_j^k .

solved for V1 neurons, but the most common assumption is that neurons in V1 employ a memory-less rate code. Under this assumption, information is carried by the number of spikes occurring in each time bin. All bins are treated independently, so there is no possibility that information is carried (or lost) by patterns in the firing rate that extend across

multiple bins. This rate-coding assumption also ignores highly precise temporal patterns that may occur within a single time bin (i.e., tight coupling of spike times to external events or internal oscillations).

If the firing rate possesses temporal correlations extending across multiple time bins, then the assumption of a memory-less rate code may lead to overestimation of the information transmission rate of the neuron. Conversely, if information is carried by the temporal structure of the spiking activity within a time bin, then the memory-less rate code assumption may lead to underestimation of the information transmission rate. Clearly, the assumption of a memory-less rate code has strengths and weaknesses. Many more complex neural codes have been proposed (for example, see Optican and Richmond, 1987; Richmond et al., 1987; Meister, 1996; de Ruyter van Steveninck et al., 1997), but their existence is controversial. More complex coding schemes are also more difficult to assess experimentally; this is especially true for codes involving extended temporal correlations. For these reasons, we have restricted the current analysis to the hypothesis of memory-less rate coding.

After the spike train is translated into discrete words, the probability of word occurrence is determined empirically from the data. After determining the occurrence probability of each word, the entropy can be found using (Shannon and Weaver, 1949):

$$H = - \sum_{j=0}^{\infty} p_j \log_2(p_j). \quad (5)$$

The summation runs over discrete words, and p_j is the probability of the occurrence of the j th word.

Under the assumption of a memory-less rate code, the spike train is divided into nonoverlapping time bins that are treated as independent words. Each word is uniquely identified via the number of spikes that it contains (Reich et al., 2000). The total response entropy is given by:

$$H(r) = - \sum_{j=0}^{\infty} p_j^{\text{total}} \log_2(p_j^{\text{total}}). \quad (6)$$

Where p_j^{total} is the number of time bins containing exactly j action potentials divided by the total number of time bins. The total response entropy is a function of both the number of distinct response words and their frequency of occurrence. Total response entropy is therefore related to the dynamic range of a neuron; neurons with larger dynamic ranges will be able to generate a larger variety of spike patterns in response to a given stimulus set.

The noise entropy describes the average variability of responses to single stimuli. Let p_j^k be the probability that the j th word occurred in response to the k th stimulus. The noise entropy for stimulus k is given by:

$$H(r|s = k) = - \sum_{j=0}^{\infty} p_j^k \log_2(p_j^k). \quad (7)$$

The probability of word occurrence for each stimulus, p_j^k , is equal to the number of stimulus-repetition trials on which the k th stimulus produces j action potentials, divided by the overall number of repetitions. (In the experiments reported in this paper, each stimulus was repeated between 10 and 40 times.)

The overall noise entropy of the neuron is found by averaging across the noise entropies of the individual stimuli:

$$H(r|s) = \langle H(r|s = k) \rangle_k. \quad (8)$$

Given $H(r)$ and $H(r|s)$, Equation 4 provides information transmission per time bin. Figure 2B, C provides a graphical overview of how the response probabilities are determined from the data. From Equation 4 it can be seen that the fundamental quantity in the analysis is the information per time bin. However, to allow comparison with other studies, it is desirable to report information transmission rates per second or per spike. Information per second is found by dividing the information per time bin by the duration of each time bin. Information per spike is found by dividing the information per second by the mean number of spikes per second.

In the current study, two experimental factors may result in underestimation of information transmission rates. First, visual stimuli were presented while animals performed a simple visual fixation task. During fixation the eye is not entirely steady; a small degree of ocular drift and corrective microsaccadic eye movements are inevitable. These small eye

Table 1. Population averages as a function of stimulus size and time bin duration

	Stimulus size	Time bin duration (msec)			
		4.6	13.8	25	50
Information per second (bits/sec)	1 × CRF	15.6	10.6	8.4	6.3
	2 × CRF	18.3	13.2*	10.7*	7.9*
	3 × CRF	18.7	13.6*	11.1*	8.0*
	4 × CRF	17.7	13.1	10.2	7.5
Information per spike (bits/spike)	1 × CRF	1.0	0.7	0.5	0.4
	2 × CRF	1.3	0.9*	0.7*	0.5*
	3 × CRF	1.6*	1.2*	1.0*	0.7*
	4 × CRF	1.6*	1.2*	1.0*	0.7*
Efficiency (%)	1 × CRF	17.5	16.5	16.8	17.4
	2 × CRF	21.8*	21.7*	22.5*	22.9*
	3 × CRF	26.0*	26.1*	26.9*	26.6*
	4 × CRF	26.2*	26.6*	26.7*	26.5*

Asterisks indicate mean values that significantly exceed 1 × CRF mean values ($p < 0.05$).

movements introduce variability in retinal stimulation that in turn increases response variability. This artificially inflates our estimates of $H(r|s)$ and thereby decreases our estimates of $I(s, r)$. A second source of bias might arise from the absence of top-down influences that could influence V1 responses during natural vision. For example, during natural vision the ocular-motor system might provide V1 with an efference signal denoting eye movements that could allow V1 to process information more efficiently. This possibility is clearly speculative, but the possible role of extraretinal influences is poorly understood in V1. Both factors suggest that our experimental estimates of information transmission should be interpreted as *de facto* lower bounds on the true capabilities of V1 neurons. Fortunately, our analysis centers on the changes in information transmission that result from differential stimulation of the nCRF. These factors should be common across nCRF stimulation conditions and have little or no effect on our results.

Choice of time-bin duration. Because our analysis assumes a rate code, the duration of the time bins should match the true integration time of the target neurons. Unfortunately, this critical time constant is unknown. To compensate, we analyze the data using several different binning times (4.6, 13.8, 25, and 50 msec) that span the range of plausible integration times (Bair, 1999). A summary of the results obtained with other bin lengths is given in Table 1 and also discussed in Results. To facilitate comparison with previous work (Vinje and Gallant, 2000), we focus on the results obtained with 13.8 msec time bins. With the exception of Table 1, all figures and results come from 13.8 msec binning unless stated otherwise.

Correction for finite data bias in the response entropies. The values for p_j^{total} and p_j^k are estimated from the experimental data, leading to uncertainty in $H(r)$ and $H(r|s)$. The uncertainties in the entropy estimates contain both random error (because of sampling) and systematic biases. Error attributable to sampling is handled conventionally, by considering whether results are statistically significant. The bias, however, can be removed explicitly. In particular, the noise entropy is strongly affected by limitations in the number of trial repetitions. In general, this results in potential underestimation of the noise entropy (which would produce an overestimation of information transmission).

For both $H(r)$ and $H(r|s)$, the relationship between the true entropy and the experimental estimate of the entropy is given by (Treves and Panzeri, 1995; Strong et al., 1998):

$$H_{\text{exp}} = H_{\text{true}} + \sum_{\alpha=1}^{\infty} \frac{c_{\alpha}}{N^{\alpha}}, \quad (9)$$

where c_{α} is an empirically determined weighting coefficient for the α th correction term and N denotes the number of times each stimulus was repeated. In our data, the linear bias term dominates the sum in Equation




Figure 3. Finite repetitions of each stimulus lead to systematic biases in the experimentally measured entropies. Response entropy (circles) and noise entropy (squares) are plotted versus $1/N$, where N is the number of trials used in the calculation of p_j^{total} and p_j^k . Error bars were determined by the Jackknife method (Efron and Tibshirani, 1993). Because of the small size of the error bars, they are shown in white atop the data points. Lines represent the least-squares fit of Equation 10 to the response entropy data (solid line) and noise entropy data (dashed line). The true entropy values are the ordinate intercepts of the best-fit functions.

9. In light of this, we consider only the first- and second-order correction terms:

$$H_{exp}^N = H_{true} + \frac{c_1}{N} + \frac{c_2}{N^2}. \quad (10)$$

All of the analyses presented in this report used the bias-corrected entropies, H_{true} .

To find H_{true} , we divide the original data set into several subsets, each containing N trials, and evaluate H_{exp}^N for each subset. Subsets contain, respectively, one-quarter, one-third, one-half, or all of the original trials. Second, we fit Equation 10 to these data via least-squares minimization. The value of H_{true} is the ordinate intercept of the best-fit function. Figure 3 illustrates this process for an example neuron.

Testing for excessive finite data bias. The number of trials required for accurate bias correction depends on time-bin duration and the response properties of the neuron under study. If there are too few repeated stimulus presentations, higher-order correction terms become important and Equation 10 fails to sufficiently describe the finite data bias. Fortunately excessive levels of bias contamination can be detected by testing whether the experimentally estimated entropies violate the Ma bound (Strong et al., 1998).

The Ma bound is a lower bound on response entropy and can be estimated for both the total response entropy, $H(r)$, and the noise entropy $H(r|s)$. For words composed of single time bins, the general expression for the Ma bound is given by (Ma, 1981):

$$H_{Ma} = -\log_2(\langle p_j \rangle) = -\log_2 \left(\sum_{j=0}^{\infty} p_j^2 \right). \quad (11)$$

H_{Ma} is useful because it is less susceptible to finite data bias than H_{exp} . The response entropy can sink below the Ma bound only if H_{exp} is strongly contaminated by finite data bias (Strong et al., 1998). Because finite data bias affects experimental estimates of the noise entropy more strongly than estimates of the total response entropy, the noise entropy is more likely to violate the Ma bound.

We computed the Ma bounds on both total response entropy and noise entropy to allow exclusion of any neurons with gross levels of bias contamination. During responses to natural-vision movies, both response entropies were greater than H_{Ma} for all neurons. This satisfies the Ma bound criterion and indicates that our entropy estimates are free from excessive finite data bias.

Significance testing. We determined whether the descriptive statistics of two sample sets are significantly different via randomized, two-tailed t tests (Manly, 1991). In all cases randomization was performed to rule out the null hypothesis that the two sets of observations come from the same underlying population distribution. Thus, significance implies that the value of the descriptive statistic for nCRF data is significantly different from the corresponding value obtained with CRF data. The standard

significance criterion of $p \leq 0.05$ is sufficient when comparing two collections of neurons. However, when judging significant differences in single neurons or time bins, we use a more restrictive significance criterion, $p \leq 0.01$.

RESULTS

Response modulation by the nCRF in area V1 during natural vision

Many studies have demonstrated that the nCRF has pronounced, generally suppressive effects on responses (Hubel and Wiesel, 1965; Blakemore and Tobin, 1972; Bishop et al., 1973; Nelson, 1991). However, nCRF modulation can also enhance responses (Jones, 1970; Hirsch and Gilbert, 1991; Knierim and Van Essen, 1992; Levitt and Lund, 1997; Kapadia et al., 2000). We have found that nCRF stimulation during natural vision can both enhance and suppress responses. Figure 4A shows the peristimulus time histogram (PSTH) obtained from one V1 neuron in response to stimulation by a natural-vision movie confined to the CRF. When the stimulus size is increased to four times the diameter of the CRF ($4 \times \text{CRF}$) some responses are enhanced while others are suppressed (Fig. 4B). The modulation ratio, R_i , summarizes the influence of the nCRF on the i th stimulus time bin (see Materials and Methods). In Figure 4, those time bins with significant R_i values ($p \leq 0.01$) are shown in white (significant suppression) or black (significant enhancement). Modulation by the nCRF depends on both image content and the elapsed time from fixation onset. These intrafixation temporal dynamics may reflect presynaptic depression (Abbott et al., 1997; Chance et al., 1998), some form of short-term adaptation, or perhaps the influences of intracortical feedback (Rao and Ballard, 1999).

To quantify the observed modulation, we calculated R_i values for all time bins in our data set (Fig. 5A–C). Again, R_i values are colored according to their significance: white for significant suppression, black for significant enhancement. (Modulation ratios and histogram values are plotted on logarithmic scales because of the large dynamic range of modulation produced by natural stimulation of the nCRF.) As stimulus size increases, there is a modest increase in the number of significantly modulated time bins. In general, enhancement is always less pronounced than suppression. However, for all stimulus sizes a substantial fraction of modulation is positive. As stimulus size increases, the net modulation becomes steadily more suppressive.

Suppression also significantly decreases the mean spiking rate of individual neurons. The fractions of neurons whose spike rates are significantly suppressed by nCRF stimulation are 50% at $2 \times \text{CRF}$, 59% at $3 \times \text{CRF}$, and 73% at $4 \times \text{CRF}$ ($p \leq 0.01$). The suppression of individual neurons is reflected in the average spike rate of the population, which decreases with increasing stimulus size (Fig. 5D).

In a previous study, we showed that increasing stimulus size decorrelated the responses of neuron pairs (Vinje and Gallant, 2000). The decorrelation index measures the relative overlap of the tuning properties for each neuron pair; as neuron pairs become decorrelated the overlap in their tuning functions is reduced. Thus, for large stimuli, different neurons were unlikely to fire in response to the same space–time stimulus, whereas for stimuli confined to the CRF, there was a significant chance of correlated firing.

Increasing nCRF stimulation produces a net increase in suppressive modulation, a reduction in the overall population activity rate, and a reduction in tuning overlap. These results support the first untested prediction: increasing nCRF stimulation reduces




Figure 4. The nCRF modulates responses during natural vision. *A*, PSTH obtained from one V1 neuron in response to a natural-vision movie confined to the CRF. Responses are weakly modulated by the simulated fixations (information per second, 13.1 bits/sec; information per spike, 0.18 bits/spike; efficiency, 10%; selectivity index, 13%). *B*, Responses of the same cell to a natural-vision movie composed of the CRF stimulation used in *A* plus a circular surrounding region. The overall stimulus size was $4 \times$ CRF diameter. Stimulation of the nCRF dramatically increases variation of responses across fixations (information per second, 28.4 bits/sec; information per spike, 0.67 bits/spike; efficiency, 26%; selectivity index, 51%). Responses to some stimuli are significantly enhanced (*black bins*; $p \leq 0.01$). For this neuron, enhancement is concentrated in the onset transients occurring at the beginning of simulated fixations. Other responses are strongly suppressed (*white bins*; $p \leq 0.01$). The *under-bar* highlights those time bins where significant enhancement and suppression occur.

metabolic load by lowering mean spike rates. Furthermore, these three findings suggest that nCRF stimulation reduces the effective bandwidth of single neurons, thereby restricting the range of stimuli that they represent.

nCRF stimulation increases information transmission rate

Does this shrinkage in effective bandwidth reduce the amount of information represented by V1 neurons? If information is lost, then the stimulus representation will be coarsened rather than made sparser (Foldiak and Young, 1995; Olshausen and Field, 1997; Barlow, 2001). Information must be preserved if nCRF stimulation truly increases sparseness. Information transmission can be preserved in numerous ways. One possibility is that the overall information transmission rate might be preserved at the level of individual neurons. Alternatively, some neurons may increase their information transmission rates while other neurons transmit less information.

Information transmission rates (bits per second) for our sample of V1 neurons are shown in Figure 6*A–D*. For each neuron at each stimulus size, we compared information rates observed with and without nCRF stimulation. Neurons with significantly increased information rates are shown in *black*, while those with significantly decreased rates are shown in *white* ($p \leq 0.01$). The effects of natural nCRF stimulation vary across neurons. Some exhibit decreases in information transmission rates, whereas others exhibit increases. Interestingly, significant increases in information transmission rates occur more frequently than significant decreases. The ratio of significant increases to significant decreases is 3.8:1 at $2 \times$ CRF, 3.4:1 at $3 \times$ CRF, and 3.7:1 at $4 \times$ CRF.

For our sample of neurons, the average information transmission rate also increases with stimulus size (Fig. 6*E*). The increase in mean rate is modest but statistically significant for stimulus sizes of $2 \times$ CRF and $3 \times$ CRF ($p \leq 0.05$) and is marginally significant for stimuli of $4 \times$ CRF diameter ($p \leq 0.07$).

Table 1 shows the average information rate as a function of stimulus size and time-bin duration. In general, the average rate increases as time-bin duration decreases. From 50 msec to 4.6 msec, the information transmission rate increases by $\sim 250\%$. The increase in information rates for short binning times is commonly observed in neurophysiological data sets (Strong et al., 1998) and

occurs because $H(r)$ increases more rapidly than $H(r|s)$ as bin duration shrinks.

Our second prediction is that the average information transmission rate should not decrease as stimulus size increases. Our results demonstrate that information transmission actually increases with stimulus size. This is consistent with the predicted preservation of information. It also suggests that nCRF stimulation may be necessary to fully realize the information-processing potential of V1 neurons.

nCRF stimulation increases information per spike

As discussed in the introductory remarks, sparse coding offers several potential advantages to the nervous system. It may simplify development of neural connections, increase learning rates, and increase memory capacity (Barlow, 1961, 2001). Sparse coding also reduces the number of action potentials required to represent a scene and thereby decreases the metabolic demands of information processing (Srinivasan et al., 1982; Laughlin et al., 1998). If the system is to maintain the fidelity with which a scene is represented, this reduction in spiking activity must be accompanied by an increase in the average amount of information each spike provides about the stimulus. Thus, natural nCRF stimulation should increase the average information carried by each spike.

The average information that a spike transmits about the stimulus is found by simply dividing the information per second by the mean number of spikes per second: $I_{\text{spike}} = I_{\text{sec}}/\mu$, where μ is the mean spike rate of the neuron for all stimuli of a given size.

Information transmission per spike is shown in Figure 7*A–D*. Figure conventions are identical to those used in Figure 6. Stimulation of the nCRF can increase or decrease the information per spike, but the trend is strongly toward increasing the information content of spikes. The ratio of neurons with significant increases to those with significant decreases is 6.5:1 at $2 \times$ CRF and 26:1 at $3 \times$ CRF. For data obtained with stimuli of $4 \times$ CRF diameter, all significantly modulated neurons show increases in their information transmission per spike.

The mean information per spike also increases substantially as a function of stimulus size (Fig. 7*E*, *black circles*). For stimuli of $4 \times$ CRF diameter, the mean information per spike is 1.85 times larger than that of the value obtained with CRF-sized stimuli. All

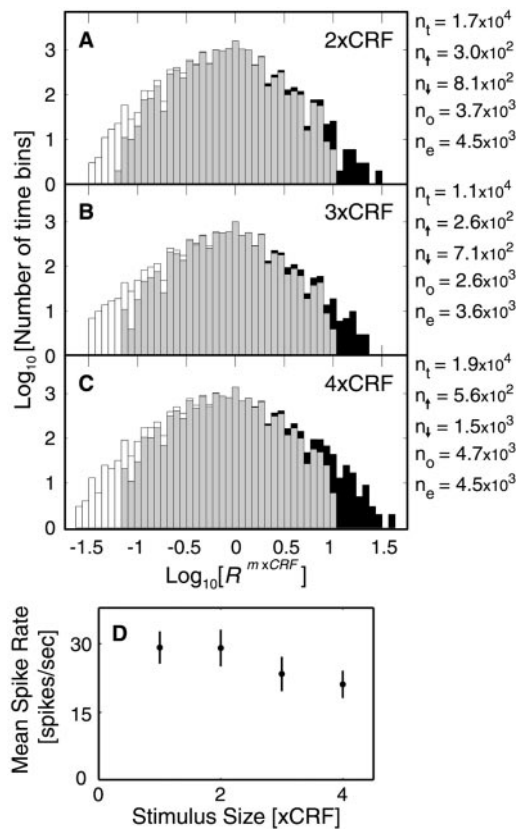


Figure 5. Stimulus size affects the modulation ratio and the mean spike rate. *A–C*, Modulation ratios, R_i , for stimuli of sizes $2 \times \text{CRF}$, $3 \times \text{CRF}$, and $4 \times \text{CRF}$. Both axes are presented on a logarithmic scale (base 10) to facilitate display. Positive values of $\log[R_i]$ indicate relative enhancement, whereas negative values of $\log[R_i]$ indicate relative suppression. Those time bins with significantly enhanced ratios are shown in *black*, whereas those with significantly decreased ratios are shown in *white* ($p \leq 0.01$). The total number of valid time bins in each data set, n_t ; the number of time bins demonstrating significant enhancement, n_{\uparrow} ; and the number of time bins demonstrating significant suppression, n_{\downarrow} , are indicated to the *right* of each histogram. Not all time bins are included in the histogram. Time bins for which $R_i = 0$ are indicated by n_o . These bins are included in the total number of valid bins, but have undefined logarithms. Time bins where CRF-sized stimuli evoke no spikes are indicated by n_e . These bins have undefined modulation ratios and are not included in the total number of valid time bins. Enhancement is less frequent than suppression, especially at larger stimulus sizes. Significant enhancement occurs in 1.7, 2.3, and 3.0% of the time bins for stimulus sizes of $2 \times \text{CRF}$, $3 \times \text{CRF}$ and $4 \times \text{CRF}$, respectively, whereas significant suppression occurs in 4.6, 6.4, and 8.0%. *D*, Average spikes per second versus stimulus size. The mean spike rate falls monotonically as stimulus size increases. The mean spike rate in response to $4 \times \text{CRF}$ stimuli is $\sim 75\%$ of mean the rate observed with CRF stimulation alone.

stimuli of a size $\geq 2 \times \text{CRF}$ produce significant increases in information per spike ($p \leq 0.05$). Because the information-per-spike distributions are positively skewed, we also evaluated the median information transmission per spike (Fig. 7*E*, *gray triangles*). As expected, the medians increase less than the means, but still increase significantly for sizes of $\geq 3 \times \text{CRF}$ ($p \leq 0.05$). Table 1 presents the average information per spike as a function of stimulus size and time-bin duration. As duration decreases, the information per spike increases in a manner similar to that observed for information per second.

Natural nCRF stimulation increases the information content of each spike for most neurons in our sample. This confirms the

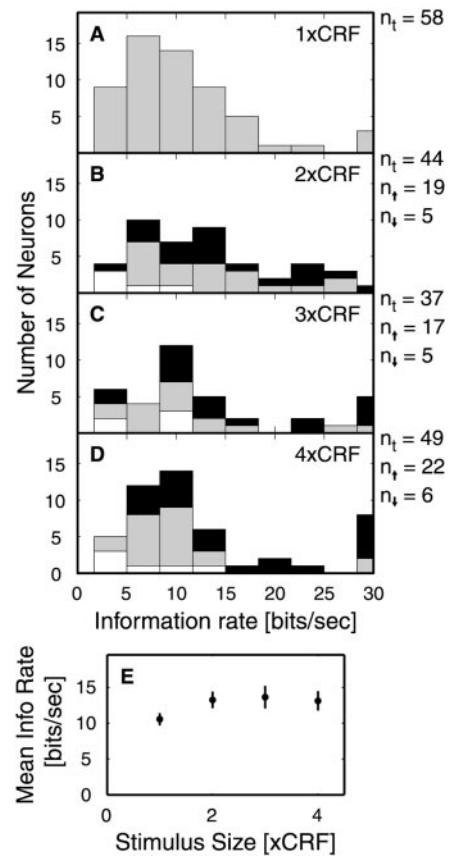


Figure 6. Information transmission in V1 neurons increases with stimulus size. *A–D*, Information transmission rates for stimuli of sizes $1 \times \text{CRF}$, $2 \times \text{CRF}$, $3 \times \text{CRF}$, and $4 \times \text{CRF}$. Neurons in which stimulus size significantly increases information transmission per second are shown in *black*, whereas those neurons with significant decreases in transmission are shown in *white* ($p \leq 0.01$). The total number of neurons in each histogram, n_t ; the number of neurons demonstrating significant increases, n_{\uparrow} ; and the number of time bins demonstrating significant decreases, n_{\downarrow} , are indicated to the *right* of each histogram. Significant increases occur in 43, 39, and 45% and significant decreases occur in 11, 14, and 12% of the neurons for stimulus sizes of $2 \times \text{CRF}$, $3 \times \text{CRF}$, and $4 \times \text{CRF}$, respectively. *E*, Mean information transmission per second versus stimulus size. The mean rates obtained with $2 \times \text{CRF}$ and $3 \times \text{CRF}$ stimuli are significantly higher than the mean rate observed with stimuli confined to the CRF ($p \leq 0.05$). With stimuli of $4 \times \text{CRF}$ diameter the increase in mean rate is marginally significant ($p \leq 0.07$).

third prediction of the hypothesis that nCRF stimulation increases sparseness in V1.

nCRF stimulation during natural vision increases efficiency

As nCRF stimulation increases sparseness, it should also increase the efficiency of information processing. In information theoretic terms, efficiency measures the fraction of available bandwidth that a neuron actually uses to transmit information. Formally this is expressed as the ratio of the amount of information actually transmitted over the theoretical maximum amount of information that could be transmitted (Cover and Thomas, 1991; Borst and Theunissen, 1999):

$$E = \frac{H(r) - H(r|s)}{H(r)} = \frac{I(r, s)}{H(r)}. \quad (12)$$

Figure 8*A–D* shows efficiency versus stimulus size for our sample of neurons. Figure conventions again match those used in

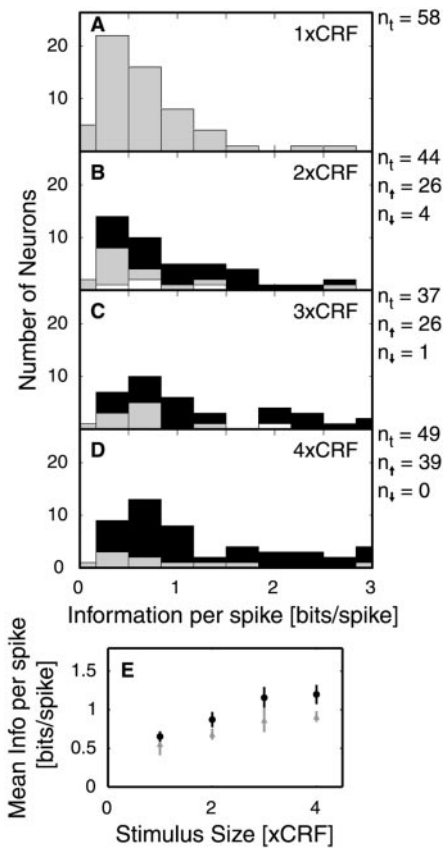


Figure 7. Information per spike increases with stimulus size. *A–D*, Information transmission per spike versus stimulus size. Figure conventions match those used in Figure 6. Significant increases in information per spike occur in 59, 70, and 79% and significant decreases occur in 9, 3, and 0% of the neurons for stimulus sizes of $2 \times$ CRF, $3 \times$ CRF, and $4 \times$ CRF, respectively. *E*, Mean (black circles) and median (gray triangles) information per spike estimates as a function of stimulus size. The mean information per spike obtained with $2 \times$ CRF and larger stimuli is significantly higher than that observed with stimuli confined to the CRF ($p \leq 0.05$).

Figure 6. As stimulus size increases, so does the efficiency of single neurons. The ratio of neurons with significant increases to those with significant decreases is 6.3:1 at $2 \times$ CRF and 26:1 at $3 \times$ CRF. With $4 \times$ CRF stimuli, all significantly modulated neurons show increases in the efficiency of information transmission.

Mean efficiency increases with nCRF stimulation (Fig. 8E); for $4 \times$ CRF-sized stimuli, the mean efficiency is 1.6 times larger than the value obtained with CRF-sized stimuli. The increases in mean efficiency are statistically significant for all stimuli of a size $\geq 2 \times$ CRF ($p \leq 0.05$). Table 1 presents the average efficiency as a function of stimulus size and time-bin duration. In contrast to information rate and information per spike, mean efficiency does not change substantially as bin duration decreases. As bin duration shrinks, increases in $H(r)$ inflate the apparent information transmission per second and per spike. However, in the case of efficiency, the denominator of Equation 12 largely cancels this effect.

Neurons use their available transmission bandwidth more efficiently when the nCRF is stimulated than when stimuli are confined to the CRF. Because efficiency does not explicitly depend on spike rate, this result complements the finding that nCRF stimulation increases the amount of information available in each spike and confirms the last of our predictions.

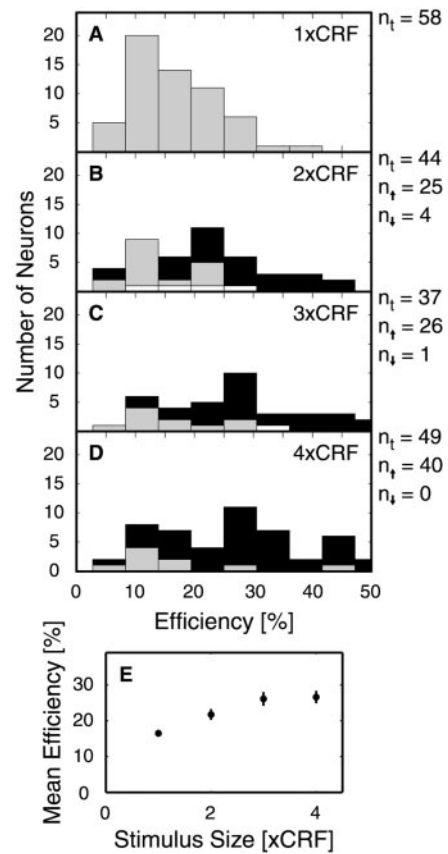


Figure 8. Efficiency of information transmission increases with stimulus size. *A–D*, Efficiency versus stimulus size. Figure conventions match those used in Figure 6. Significant increases in efficiency occur in 57, 70, and 82% and significant decreases occur in 9, 3, and 0% of the neurons for stimulus sizes of $2 \times$ CRF, $3 \times$ CRF, and $4 \times$ CRF, respectively. *E*, Mean efficiency as a function of stimulus size. Efficiency observed with $2 \times$ CRF and larger stimuli is significantly higher than that observed with stimuli confined to the CRF ($p \leq 0.05$).

Information transmission and efficiency correlate with selectivity

Thus far we have shown that nCRF stimulation increases information transmission rates, the information content of single spikes, and processing efficiency in both individual neurons and our sample population. In a previous study (Vinje and Gallant, 2000), we showed that nCRF stimulation increases the selectivity of V1 neurons (Fig. 9). All of these results are consistent with the idea that nCRF stimulation increases the sparseness of the representation of visual information in V1. A supplementary test of the sparse coding hypothesis is to determine whether selectivity is correlated with information transmission in individual neurons. If the nCRF increases sparseness, then cells that show a substantial increase in selectivity contingent on nCRF stimulation should be more informative and more efficient than those that do not show such changes.

Stimulus selectivity is not significantly correlated with information per second in our sample of cells (Fig. 10A–D). However, selectivity is significantly correlated with information per spike for all stimulus sizes (Fig. 10E–H; $p \leq 0.01$). The correlations between information per spike and selectivity are 0.91, 0.90, 0.89, and 0.89 for CRF-, $2 \times$ CRF-, $3 \times$ CRF-, and $4 \times$ CRF-sized stimuli, respectively. Finally, stimulus selectivity is also significantly correlated with efficiency (Fig. 10I–L; $p \leq 0.01$). Correla-

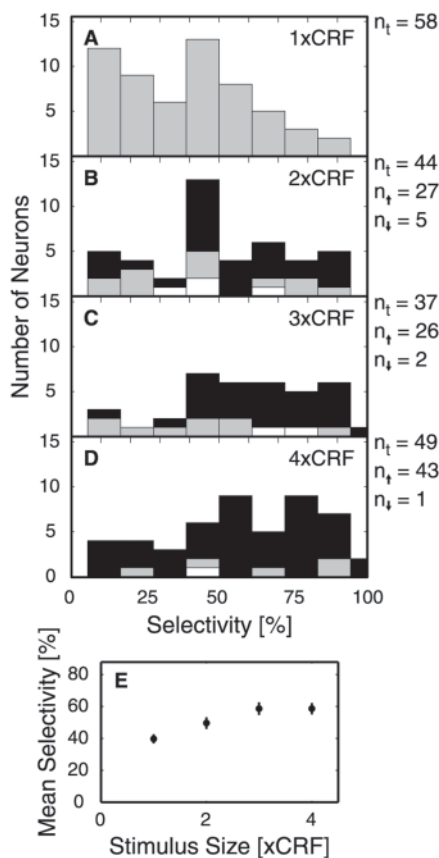


Figure 9. Stimulus selectivity increases with stimulus size. *A–D*, Selectivity versus stimulus size. Figure conventions match those used in Figure 6. Significant increases in selectivity occur in 61, 70, and 88% and significant decreases occur in 11, 5, and 2% of the neurons for stimulus sizes of $2 \times$ CRF, $3 \times$ CRF, and $4 \times$ CRF, respectively. *E*, Mean selectivity versus stimulus size. Selectivity obtained with $2 \times$ CRF and larger stimuli is significantly higher than that found with stimuli confined to the CRF ($p \leq 0.05$).

tions between efficiency and selectivity are 0.89, 0.89, 0.87, and 0.88 for CRF-, $2 \times$ CRF-, $3 \times$ CRF-, and $4 \times$ CRF-sized stimuli, respectively.

The lack of correlation between information transmission per second and selectivity suggests that the observed increases in information rate may not be central to the process of increasing sparseness. This is perhaps unsurprising, given that the prediction was merely that average information transmission should be preserved. In contrast, the correlation of selectivity with information per spike and efficiency suggests that these three measures are related by an underlying causal factor. It seems likely that this causal factor is the sparseness of information representation in V1. As sparseness increases, there are corresponding increases in selectivity, information per spike, and efficiency of information transmission.

Results obtained with flashed natural-image patches

Natural-vision movies are designed to mimic the stimulation that occurs during saccadic vision of a static scene. The majority of the movie consists of fixations where image content is held constant. These fixations are linked by simulated saccades with realistic acceleration profiles. The stimulation contained in saccades blends together image patches and avoids any discontinuous change in stimulus content. Most previous physiology experi-

ments use flashed stimuli that contain instantaneous onset and offset transitions and substantial interstimulus intervals. Clearly the nature of the transitions between image patches is very different in these two procedures.

To facilitate comparisons between our results and those obtained using flashed stimuli, we performed the following control experiment. Image patches were selected from natural scenes and presented as flashed stimuli ($n = 10$ neurons; see Materials and Methods). Responses to the flashed stimuli were concatenated to form a pseudo-movie and analyzed in the same manner as natural-vision movies (responses during interstimulus intervals were discarded).

Unfortunately, flashed stimulus patches were presented only five times; therefore, this data set suffers from larger entropy biases than our main data set. This problem is partially caused by difficulty in accurately estimating the second correction term in Equation 10 and was partially ameliorated by using only the linear correction term. To enable comparison with data from natural-vision movies, we also limited the natural-vision data to the first five trials and applied only the linear bias correction. This approach subjected both data sets to the same bias-producing conditions and thus allowed a fair comparison between the results from the natural vision and the flashed stimuli.

The effects of nCRF stimulation with flashed stimuli are generally similar to those obtained with natural-vision movies. Both the information per spike and the efficiency increase with stimulus size. For the largest flashed stimulus size ($3 \times$ CRF), the average information per spike increases by 25% and the average efficiency increases by 10%. Information transmission per second does not increase.

These results suggest that increasing stimulus size increases response sparseness with flashed stimuli, as it does with natural-vision movies. However, these effects are somewhat smaller with flashed stimuli. This may be an artifact of small sample size, because not all neurons demonstrate strong nCRF modulation effects. Alternatively, this may reflect differences in the transient responses evoked by the two stimulus classes. When the first 200 msec are removed from the response to each flashed-image patch, information transmission more closely matches that obtained with natural-vision movies.

Total entropy and noise entropy both decrease with increasing stimulus size

Information is the difference between two measures of variance, the total response entropy and the noise entropy. An increase in information can reflect a decrease in noise entropy, an increase in total entropy, or some combination of the two. Each of these changes would alter neuronal spiking patterns in ways that allow insight into the specific biophysical mechanisms underlying nCRF modulation. If nCRF stimulation increases the total response entropy, then the nCRF must increase the dynamic range of the neuron and/or the reliability of spikes elicited by the stimulus. In contrast, if the noise entropy decreases consequent to nCRF stimulation, then the nCRF must suppress spikes that are not relevant to encoding the stimulus.

Entropy measures are summarized in Figure 11. Figure 11*A–D* shows total stimulus entropy, and Figure 11*E–H* shows noise entropy. Those neurons with significantly increased entropies are shown in *black*, while those with significantly decreased entropies are shown in *white* ($p \leq 0.01$). It is readily apparent that the nCRF has a large effect on both total entropy and noise entropy. On average, both total entropy and noise entropy decrease with

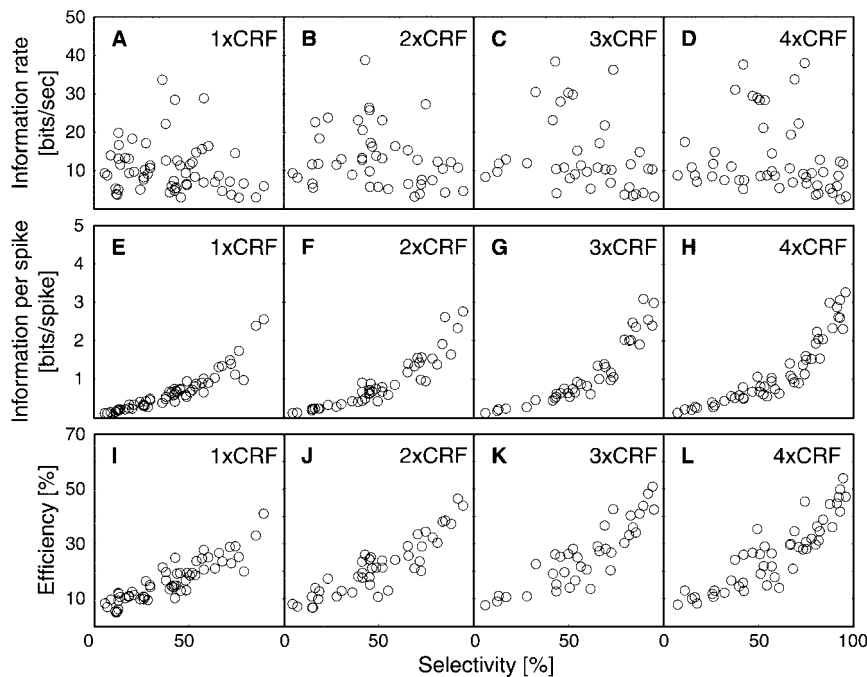


Figure 10. Information per spike and efficiency are correlated with selectivity. *A–D*, Information transmission rate versus selectivity for stimuli $2 \times$ CRF, $3 \times$ CRF, and $4 \times$ CRF. The lack of meaningful correlation is immediately apparent. *E–H*, Information per spike versus selectivity. These measures are strongly correlated for all stimulus sizes. *I–L*, Efficiency versus selectivity. The measures are also correlated for all stimulus sizes.

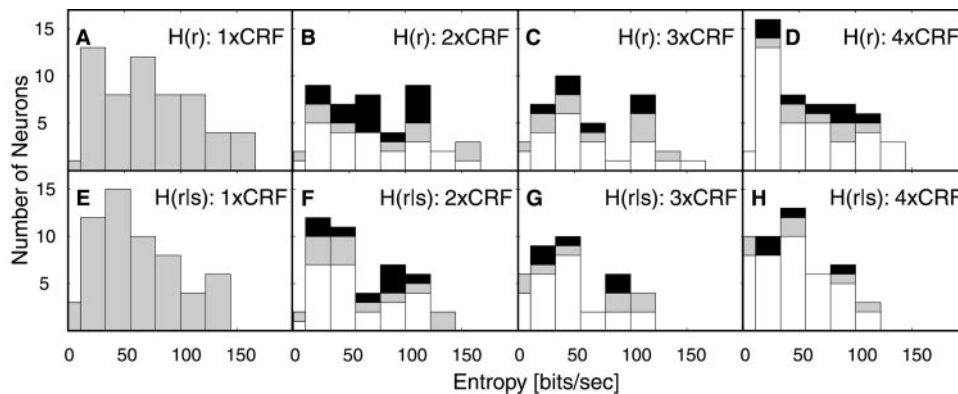


Figure 11. Stimulation of the nCRF decreases both total response entropy and noise entropy. *A–D*, Total response entropy versus stimulus size. Significant increases in total response entropy occur in 30, 16, and 14%, and significant decreases occur in 50, 57, and 71% of the neurons for stimulus sizes of $2 \times$ CRF, $3 \times$ CRF, and $4 \times$ CRF, respectively. *E–H*, Noise entropy versus stimulus size. Significant increases in noise entropy occur in 18, 14, and 8% and significant decreases occur in 55, 65, and 80% of the neurons for stimulus sizes of $2 \times$ CRF, $3 \times$ CRF and $4 \times$ CRF, respectively. Larger stimuli tend to reduce both the total response entropy and the noise entropy of individual neurons.

nCRF stimulation. However, the noise entropy falls faster than the total entropy.

The differential effect of nCRF stimulation on these two entropies underlies the observed increases in information rate, information per spike, and efficiency. The simultaneous decrease of both total and noise entropies explains why nCRF stimulation has a relatively weak effect on information per second: such stimulation decreases both total entropy and noise entropy and dilutes the effective increase in overall information transmission rates.

Information per spike and efficiency are both ratio measures with the weakly increasing information rate in their numerators. However, the denominator terms of both measures (μ and $H(r)$, respectively) shrink with increasing stimulus size. This convergence of a weakly increasing numerator and a decreasing denominator underlies the strong increases in information per spike and efficiency as a function of stimulus size. The nCRF appears to suppress most responses and enhance a select few. As sparseness increases, those action potentials that are not reliably linked to stimulus properties are winnowed from the responses of the neuron.

DISCUSSION

Our results show that nCRF stimulation changes the response entropies of V1 neurons. Stimulation of the nCRF decreases total response entropy but has an even greater effect on the noise entropy. This differential modulation underlies the pattern of results we observed: relative to CRF stimulation alone, naturalistic nCRF stimulation increases selectivity, information per second, information per spike, and efficiency.

Previous theoretical research has shown that the informative components of natural scenes are sparsely distributed (Field, 1987; Olshausen and Field, 1996, 1997; Bell and Sejnowski, 1997). Our results suggest that the nCRF might tune V1 neurons to match the sparsely distributed, informative components of natural scenes. The resulting neural code is also sparse, highly selective, and efficient.

The level of sparseness in the neural code does not necessarily match that of natural images; the sparse components of natural images are determined by the physical structure of the world, while the sparseness of a neural code also reflects biophysical, computational, and behavioral constraints (van Hateren and Ruderman, 1998). Therefore, the neural code might be more or

less sparse than would be expected based simply on the statistics of natural images. Future research should reveal how the level of sparseness in the neural representation compares with the distribution of informative image features in natural scenes.

Stimulus dependence of information transmission

Recently, Reich et al. (2000) have evaluated information transmission by V1 neurons in anesthetized macaques. Their work is particularly relevant because their analysis was similar to ours and their stimulus set is complementary to ours. They used three types of stimuli: drifting sinusoidal gratings, stationary gratings, and checkerboard m-sequences. Each stimulus encompassed both the CRF and nCRF of the neurons under study.

For the V1 complex cells in their study, Reich et al. (2000) obtained the following median values of information transmission using drifting gratings, stationary gratings, and m-sequences, respectively: 1.58, 4.38, and 4.99 bits/sec and 0.08, 0.19, and 0.42 bits/spike. They found higher information rates for simple cells: 10.28, 7.29, and 6.41 bits/sec and 0.92, 0.25, and 0.69 bits/spike. Our natural visual stimuli produce substantially higher information transmission than that reported by Reich et al. (2000). For large stimuli ($4 \times$ CRF), the median values of information transmission are 9.12 bits/sec and 0.91 bits/spike.

There are several factors that may underlie this difference. First, although both studies derive their bias removal methods from Treves and Panzeri (1995), the details of the methods are different. These differences could potentially affect estimated information rates. Second, information transmission rates might be affected by the anesthesia used in the Reich et al. (2000) study. Finally, it is possible that V1 neurons transmit relatively more information about the natural stimuli used in our experiments. This is consistent with the observation of Reich et al. (2000) that information transmission rates are stimulus-dependent. If this stimulus dependence underlies the difference in information rates, then it suggests that V1 neurons are optimized for representing natural visual stimuli.

Comparison with information transmission in H1 neurons

Because information is measured using a rather abstract scale (bits), it is difficult to appreciate the meaning of the information-transmission rates observed in a single sensory system. One way to better understand our results is to compare them with information processing in a different nervous system. The blowfly (*Calliphora vicina*) is an interesting case in this regard. From the standpoint of information theory, the most studied cells in the blowfly are the wide-field, velocity-sensitive H1 neurons. The blowfly has only two H1 neurons and must rely on this pair to provide crucial flight control information. If H1 neurons are highly optimized for representing wide-field motion, their information-transmission properties should reflect this fact.

Previous studies have reported that H1 neurons transmit ~ 1 bit/spike and operate at an efficiency of $\sim 50\%$ (de Ruyter van Steveninck et al., 1997; Strong et al., 1998). For V1 neurons provided with $4 \times$ CRF natural-vision movies, the average information transmission is 1.2 bits/spike and the average efficiency is $\sim 27\%$. These values are probably lower bounds on the true information-transmission capacities of V1 neurons; we make several conservative assumptions in our analysis while simultaneously removing potential inflationary biases (see Materials and Methods). Given this situation, it is impressive that some V1

neurons are more efficient than the average H1 neuron (i.e. >3 bits/spike and/or efficiency values approaching 50%).

The neocortex is a comparatively recent evolutionary innovation. Despite this, during natural vision the information-transmission properties of some V1 neurons are roughly comparable with the transmission properties of H1 neurons responding to wide-field motion. This also supports the idea that V1 neurons are optimized to process the information in natural scenes.

Coding density in the visual system

Coding density is a fundamental property of any neural representation: do a few neurons encode the important information or is information distributed across most of the available neurons? Each stage of sensory processing offers an opportunity to alter the sparseness of the stimulus representation. At each stage, the representation of sensory input may be refined so that the most informative components are easily accessible to higher areas. Thus, it is important to identify the factors that influence coding density at each processing stage and to determine whether they increase or decrease sparseness. Unfortunately, this matter has received little experimental attention.

Previous researchers have suggested that the spatiotemporal tuning properties of the retina (Srinivasan et al., 1982) and the lateral geniculate nucleus (Dan et al., 1996) minimize the encoding of redundant visual input. Beyond V1, there have been only two studies of coding density (Young and Yamane, 1992; Rolls and Tovee, 1995). The representation of the visual world in the inferotemporal cortex appears to be at least as sparse as the representation we find in V1. However, direct comparisons are difficult because of differences in stimulation and analysis. Future studies using common stimulus sets and sparseness metrics while recording from neurons in multiple cortical visual areas will allow rigorous comparisons and may lead to a deeper understanding of how information is represented during visual processing.

What does the nCRF do?

Previous studies of the nCRF have focused on suppressive modulation and have suggested that the nCRF is critical for contrast-gain control (Geisler and Albrecht, 1992; Heeger, 1992; Wilson and Humanski, 1993). Others have noted that appropriate nCRF stimulation can actually enhance responses, suggesting that it plays a role in representing specific features such as curvature, extended contours, corners, and texture boundaries (Gilbert and Wiesel, 1990; Knierim and Van Essen, 1992; Wilson and Richards, 1992; Sillito et al., 1995; Fitzpatrick, 2000). Our results demonstrate that natural nCRF stimulation can both facilitate and suppress responses. However, suppression predominates, consistent with contrast-gain control models. In addition, the nCRF also increases selectivity and optimizes filtering, which may allow for more efficient processing of the sparse components of natural scenes. The finding that the nCRF of V1 neurons promotes sparse coding appears consistent with previous studies of nCRF function and adds a new dimension to our understanding of the nCRF.

Full-field stimulation is common during normal vision. In this regime, V1 neurons represent the visual world with a relatively sparse code and operate at their peak efficiency. However, some neurophysiological experiments confine stimuli to the CRF and avoid nCRF stimulation. Under these conditions V1 neurons evidently operate below their true potential and transmit less information with lower efficiency than they would if both the CRF and nCRF were stimulated. Our results suggest that during nat-

ural vision, the CRF and nCRF act together as a single unit optimized for processing natural scenes.

REFERENCES

- Abbott LF, Varela JA, Sen K, Nelson SB (1997) Synaptic depression and cortical gain control. *Science* 275:220–224.
- Allman J, Miezin F, McGuinness E (1985) Stimulus-specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons on visual neurons. *Annu Rev Neurosci* 8:407–430.
- Bair W (1999) Spike timing in the mammalian visual system. *Curr Opin Neurobiol* 9:447–453.
- Barlow HB (1961) Possible principles underlying the transformation of sensory messages. In: *Sensory communication* (Rosenbluth WA, ed), pp 217–234. Cambridge, MA: Cambridge UP.
- Barlow HB (1989) Unsupervised learning. *Neural Comput* 1:295–311.
- Barlow HB (2001) Redundancy reduction revisited. *Network* 12:241–253.
- Bell AJ, Sejnowski TJ (1997) The “independent components” of natural scenes are edge filters. *Vision Res* 37:3327–3338.
- Bishop PO, Coombs JS, Henry GH (1973) Receptive fields of simple cells in the cat striate cortex. *J Physiol (Lond)* 231:31–60.
- Blakemore C, Tobin EA (1972) Lateral inhibition between orientation detectors in the cat's visual cortex. *Exp Brain Res* 15:439–440.
- Borst A, Theunissen FE (1999) Information theory and neural coding. *Nat Neurosci* 2:947–957.
- Burman DD, Seagraves MA (1994) Primate frontal eye field activity during natural scanning eye movements. *J Neurophysiol* 71:1266–1271.
- Chance FS, Nelson SB, Abbott LF (1998) Synaptic depression and the temporal response characteristics of V1 cells. *J Neurosci* 18:4785–4799.
- Connor CE, Preddie DC, Gallant JL, Van Essen DC (1997) Spatial attention effects in macaque area V4. *J Neurosci* 17:3201–3214.
- Cover TM, Thomas JA (1991) *Elements of information theory*. New York: Wiley.
- Dan Y, Atick JJ, Reid RC (1996) Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J Neurosci* 16:3351–3362.
- 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.
- Efron B, Tibshirani RJ (1993) *An introduction to the bootstrap*. London: Chapman and Hall.
- Field DJ (1987) Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am A* 4:2379–2394.
- Field DJ (1993) Scale-invariance and self-similar “wavelet” transforms: an analysis of natural scenes and mammalian visual systems. In: *Wavelets, fractals, and Fourier transforms* (Farge M, Hunt J, Vascillicos C, eds), pp 151–193. Oxford: Oxford UP.
- Fitzpatrick D (2000) Seeing beyond the receptive field in primary visual cortex. *Curr Opin Neurobiol* 10:438–443.
- Foldiak P, Young MP (1995) Sparse coding in the primate cortex. In: *The handbook of brain theory and neural networks* (Arbib MA, ed), pp 895–898. Cambridge, MA: Cambridge UP.
- Gallant JL, Connor CE, Van Essen DC (1998) Neural activity in areas V1, V2 and V4 during free viewing of natural scenes compared to controlled viewing. *NeuroReport* 9:85–90.
- Geisler WS, Albrecht DG (1992) Cortical neurons: Isolation of contrast gain control. *Vision Res* 8:1409–1410.
- Gilbert CD, Wiesel TN (1990) The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Res* 30:1689–1701.
- Heeger DJ (1992) Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9:181–197.
- Hirsch JA, Gilbert CD (1991) Synaptic physiology of horizontal connections in the cat's visual cortex. *J Neurosci* 11:1800–1809.
- Hubel DH, Wiesel TN (1965) Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat. *J Neurophysiol* 28:229–289.
- Jones BH (1970) Responses of single neurons in cat visual cortex to a simple and more complex stimulus. *Am J Physiol* 218:1102–1107.
- Kapadia MK, Westheimer G, Gilbert CD (2000) Spatial distribution of contextual interactions in primary visual perception. *J Neurophysiol* 84:2048–2062.
- Keating CF, Keating EG (1982) Visual scan patterns of rhesus monkeys viewing faces. *Perception* 11:211–219.
- Knierim J, Van Essen DC (1992) Neural responses to static texture patterns in area V1 of the alert macaque monkey. *J Neurophysiol* 67:961–980.
- Krieger G, Zetsche C (1996) Nonlinear image operators for the evaluation of local intrinsic dimensionality. *IEEE Trans Image Process* 5:1026–1041.
- Laughlin SB, de Ruyter van Steveninck RR, Anderson JC (1998) The metabolic cost of neural information. *Nat Neurosci* 1:36–41.
- Levitt JB, Lund JS (1997) Contrast dependence of contextual effects in primate visual cortex. *Nature* 387:73–76.
- Ma S (1981) Calculation of entropy from data of motion. *J Stat Phys* 26:221–240.
- Manly BFJ (1991) *Randomization and Monte-Carlo methods in biology*. New York: Chapman and Hall.
- Meister M (1996) Multineuronal codes in retinal signaling. *Proc Natl Acad Sci USA* 93:609–614.
- Nelson SB (1991) Temporal interactions in the cat visual system. I. Orientation-selective suppression in the visual cortex. *J Neurosci* 11:344–356.
- Olshausen BA, Field DJ (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381:607–609.
- Olshausen BA, Field DJ (1997) Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Res* 23:3311–3325.
- Olshausen BA, Field DJ (2000) Vision and the coding of natural images. *Am Sci* 88:238–245.
- Optican LM, Richmond BJ (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J Neurophysiol* 57:162–178.
- Rao RPN, Ballard DH (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci* 2:79–87.
- Reich DS, Mechler F, Victor JD (2000) Formal and attribute-specific information in primary visual cortex. *J Neurophysiol* 85:305–318.
- Richmond BJ, Optican LM, Podell M, Spitzer H (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. I. Response characteristics. *J Neurophysiol* 57:132–146.
- Rolls ET, Tovee MJ (1995) Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J Neurophysiol* 73:713–726.
- Shannon CE, Weaver W (1949) *The mathematical theory of communication*. Chicago: University of Illinois.
- Sibson NR, Dhankhar A, Mason GF, Rothman DL, Behar KL, Shulman RG (1998) Stoichiometric coupling of brain glucose metabolism and glutamatergic neuronal activity. *Proc Natl Acad Sci USA* 95:316–321.
- Sillito AM, Grieve KL, Jones HE, Cudeiro J, Davis J (1995) Visual cortical mechanisms detecting local focal orientation discontinuities. *Nature* 378:492–496.
- Simoncelli EP, Olshausen BA (2001) Statistical properties of natural images. *Annu Rev Neurosci* 25:1193–1216.
- Snodderly DM, Gur M (1995) Organization of striate cortex of alert, trained monkeys (*Macaca fascicularis*): ongoing activity, stimulus selectivity, and widths of receptive field activating regions. *J Neurophysiol* 74:2100–2125.
- Srinivasan MV, Laughlin SB, Dubs A (1982) Predictive coding: a fresh view of inhibition in the retina. *Proc R Soc Lond B Biol Sci* 216:427–459.
- 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.
- Tovee MJ, Rolls ET, Treves A, Bellis RP (1993) Information encoding and the responses of single neurons in the primate temporal visual cortex. *J Neurophysiol* 70:640–654.
- Treves A, Panzeri S (1995) The upward bias in measures of information derived from limited data samples. *Neural Comput* 7:399–407.
- van Hateren JH, Ruderman DL (1998) Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proc R Soc Lond B Biol Sci* 265:2315–2320.
- Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287:1273–1276.
- Wilson HR, Humanski R (1993) Spatial frequency adaptation and gain control. *Vision Res* 33:1133–1149.
- Wilson HR, Richards WA (1992) Curvature and separation discrimination at texture boundaries. *J Opt Soc Am A* 9:1653–1662.
- Young MP, Yamane S (1992) Sparse population coding of faces in the inferotemporal cortex. *Science* 256:1327–1331.
- Zetsche C, Barth E (1990) Fundamental limits of linear filters in the visual processing of two-dimensional signals. *Vision Res* 30:1111–1117.