

Information in the Neuronal Representation of Individual Stimuli in the Primate Temporal Visual Cortex

EDMUND T. ROLLS

Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford, OX1 3UD, England
edmund.rolls@psy.ox.ac.uk

ALESSANDRO TREVES

Cognitive Neuroscience, S.I.S.S.A., via Beirut 2-4, 34103 Trieste, Italy

MARTIN J. TOVEE

Department of Psychology, University of Newcastle, Ridley Building, Newcastle upon Tyne, NE1 7RU

STEFANO PANZERI

Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford, OX1 3UD, England

Received February 15, 1996; Revised January 22, 1997; Accepted February 20, 1997

Action Editor: L.F. Abbott

Abstract. To analyze the information provided about individual visual stimuli in the responses of single neurons in the primate temporal lobe visual cortex, neuronal responses to a set of 65 visual stimuli were recorded in macaques performing a visual fixation task and analyzed using information theoretical measures. The population of neurons analyzed responded primarily to faces. The stimuli included 23 faces and 42 nonface images of real-world scenes, so that the function of this brain region could be analyzed when it was processing relatively natural scenes.

It was found that for the majority of the neurons significant amounts of information were reflected about which of several of the 23 faces had been seen. Thus the representation was not local, for in a local representation almost all the information available can be obtained when the single stimulus to which the neuron responds best is shown. It is shown that the information available about any one stimulus depended on how different (for example, how many standard deviations) the response to that stimulus was from the average response to all stimuli. This was the case for responses below the average response as well as above.

It is shown that the fraction of information carried by the low firing rates of a cell was large—much larger than that carried by the high firing rates. Part of the reason for this is that the probability distribution of different firing rates is biased toward low values (though with fewer very low values than would be predicted by an exponential distribution). Another factor is that the variability of the response is large at intermediate and high firing rates.

Another finding is that at short sampling intervals (such as 20 ms) the neurons code information efficiently, by effectively acting as binary variables and behaving less noisily than would be expected of a Poisson process.

Keywords: coding, vision, inferior temporal visual cortex, distributed representation, grandmother cell

Introduction

The visual pathways project by a number of cortico-cortical stages from the primary visual cortex until they reach the temporal lobe visual cortical areas (Seltzer and Pandya, 1978; Maunsell and Newsome, 1987; Baizer et al., 1991; Rolls, 1991, 1992a). Neurons with different types of sensitivity to visual stimuli tend to be found in different parts of these temporal cortical areas (Baylis et al., 1987). In some areas neurons respond to stimulus properties such as shape, orientation, texture, and color (Baylis et al., 1987; Tanaka et al., 1991), and in some other areas, especially areas in the cortex in the superior temporal sulcus, up to 20% of the neurons with visual responses have selectivity for faces (Desimone and Gross, 1979; Bruce et al., 1981; Desimone et al., 1984; Gross et al., 1985; Rolls, 1981a, 1981b, 1984, 1992a, 1992b; Perrett et al., 1982; Desimone, 1991). Some of the temporal cortical areas provide a representation of objects and faces that is relatively invariant with respect to retinal position, size, rotation, and even view, and such invariant representations form appropriate inputs to associative neuronal networks in structures to which the temporal cortical areas project, such as the hippocampus and amygdala (see for example, Rolls, 1992a, 1992b, 1992c; Treves and Rolls, 1994). Consistent with this, lesions of the inferior temporal visual cortex impair the ability of monkeys to respond to objects irrespective of changes in size, lighting, and viewing angle (Weiskrantz and Saunders, 1984).

An important question then arises of how the information about objects and faces is represented by the activity of temporal cortical neurons. Important issues are how selective and “information-bearing” (Suga, 1989) the neurons are for different classes of stimulus, such as face versus nonface; how selective or information-bearing the neurons are for individual items within a class; and whether the neurons use “local” or grandmother cell encoding, with strong or even great selectivity of a single neuron for a particular object in the environment (Barlow, 1972), or fully distributed representations in which all the neurons participate (Hinton et al., 1986; Churchland and Sejnowski, 1992), or sparse representations in which encoding by a sparse ensemble is used (Rolls and Treves, 1990, 1998; Treves and Rolls, 1991). Sparse ensemble encoding might have a number of advantages, such as showing continuous generalization as the nature of the input changes, showing graceful degradation if the network is incompletely formed or damaged, and providing for

large numbers of representations to be stored. The sparseness of the encoding is an important parameter in setting the number of memories that can be stored in a network (Rolls and Treves, 1990, 1998; Treves and Rolls, 1991). It is the aim of the work described here to provide a quantitative analysis of evidence on these issues.

In a previous investigation of whether these neurons respond differently to different faces (and so could provide information useful for face identification), we have shown that in many cases (77% of one sample), these neurons are sensitive to differences between faces (as shown by analyses of variance) (Baylis et al., 1985). However, each neuron does not respond only to one face. Instead, each neuron has a different relative response to each of the members of a set of faces. To quantify how finely these neurons were tuned to the faces of particular individuals, a measure derived from information theory, the breadth of tuning metric developed by Smith and Travers (1979), was calculated. This is a coefficient of entropy (H) for each cell which ranges from 0.0 (representing total specificity to one stimulus) to 1.0 (which indicates an equal response to the different stimuli). The breadth of tuning of the majority of the neurons analyzed was in the range 0.8 to 1.0. It was thus clear from this and other quantitative measures of the tuning of these face-responsive neurons that they did not respond only to the face of one individual but that instead typically each neuron responded to a number of faces in the stimulus set (which included five different faces) (Baylis et al., 1985). In a more recent investigation of the fineness of the representation provided by these neurons, the following measure of sparseness was introduced because it can be applied to analytic investigations of the storage properties of neuronal networks (Treves and Rolls, 1991, 1994):

$$a = \left(\sum_{i=1,n} r_i/n \right)^2 / \sum_{i=1,n} (r_i^2/n),$$

where r_i is the firing rate to the i th stimulus in the set of n stimuli. The sparseness ranges from $1/n$, when the cell responds to only one stimulus, to a maximal value of 1.0, attained when the cell responds with the same rate to all stimuli. It was found that the sparseness of the representation of the 65 stimuli of which 23 were faces by each neuron had an average across all neurons of 0.68. This indicates a rather distributed representation (Rolls and Tovee, 1995b). Such evidence shows

that only across a population or ensemble of such cells is information conveyed that would be useful in making different behavioral responses to different faces. The evidence from the firing-rate responses of these neurons about which individual was being seen was significant, as shown by the finding that the number of standard deviations that separated the response to the most effective from that to the least effective face in the set (a measure analogous to detectability, d' , in signal detection theory) was for many neurons greater than 1.0 (Baylis et al., 1985). This was confirmed by the application of an information theoretic analysis to the responses of these neurons, which showed that their responses contained much more information about which of 20 face stimuli had been seen (on average 0.39 bits) than about which (of 20) nonface stimuli had been seen (on average 0.07 bits) (Rolls and Tovee, 1995b).

A further quantitative approach to the issue of neural encoding of information about individual stimuli that has not been taken before but that is taken here is to analyze how much information about each individual stimulus in the set is reflected in the response of the neuron to each member of the set of stimuli. Information theory provides a measure of how much information is provided by the occurrence of events. The more unlikely an event is, the more information is gained by its occurrence. In neural systems, the temporal occurrence of action potentials can be utilized to determine the information reflected in the responses of a cell to a set of stimuli. It is natural to think of information as being reflected by the increase of firing rate of a neuron to a stimulus, but the absence of a neuronal response might also convey information about which stimulus was presented, or information might be reflected in the temporal characteristics of the response of a neuron to different stimuli. The unit of information is the bit. The amount of information is the negative \log_2 of the probability of an event occurring.¹ If the event s_i has probability $P(s_i)$, then the information $I(s_i)$ gained by the occurrence of s_i is

$$I(s_i) = -\log_2 P(s_i).$$

Thus if the probability of an event occurring is 1/16, then four bits of information are provided by the occurrence of that event. If a cell responded at every presentation of one stimulus of a set of 64 stimuli and not at all to the other 63 stimuli, then 6 bits ($-\log_2 1/64$) of information are reflected about which stimulus was shown when the cell fires (and $-\log_2(63/64) \simeq 0.02$ bits when it does not fire). In the present study

information analysis was applied to the responses of individual neurons recorded in the cortex in the anterior part of the superior temporal sulcus to a large set of face and nonface stimuli. The information conveyed about individual visual stimuli and about the stimulus set was calculated to provide a metric for the responses of these temporal visual cortex neurons. The information theoretic analyzes used were developed from those developed by Optican and Richmond (1987), Optican et al. (1991), and Tovee et al. (1993). The further developments in the analysis used here (described in the Methods section) are by Treves and Panzeri (1995) and Panzeri and Treves (1996).

This investigation is one of a series (Rolls, 1992a, 1994, 1995; Hornak et al., 1996) designed to investigate the normal functions of the temporal lobe visual cortical areas and how damage to these brain regions may underlie the perceptual deficits found in patients with disruption of function of these and connected regions.

Methods

The responses of single neurons in the temporal cortical visual areas were measured to a set of 65 visual stimuli in macaques performing a visual fixation task. The stimuli included 23 monkey and human faces and 42 nonface images of real-world scenes, so that the function of this brain region could be analyzed when it was processing natural scenes. The neurons were selected to meet the previously used criteria of face selectivity by responding more than twice as much to the optimal face as to the optimal nonface stimulus in the set (Rolls, 1984, 1992). Earlier investigations on this set of neurons have been described previously (Rolls and Tovee, 1995b), and the neurophysiological methods are described there. The majority of these neurons were recorded in the cortex in the anterior part of the superior temporal sulcus, and the recording sites are shown elsewhere (Rolls and Tovee, 1995b).

Information Analysis

The principles of the information theoretic analysis were similar to those developed by Richmond and Optican (1987) and Optican and Richmond (1987), except that we applied a novel correction procedure for the limited number of trials. Whereas previous correction procedures were basically empirical ad hoc methods, we were recently able to derive analytically

a correction term that significantly improves the reliability of information estimates, as verified with computer simulations and as reported by Treves and Panzeri (1995), Panzeri and Treves (1996) and Golomb et al. (1997). The last paper reports also, for comparison, the results obtained with another procedure designed to allow reliable information estimates from limited samples, developed by John Hertz and collaborators, and it shows that both methods, in particular ours, allow accurate information measurements when the limited data would have made uncorrected measurements grossly inaccurate. A novel aspect of the data analysis described here is that we investigated how much information was available about each stimulus in the set. Because most of the information about which stimulus is shown is made evident by measuring the firing rate of the neuron, and temporal encoding adds relatively little additional information for this population of neurons (Tovee et al., 1993; Tovee and Rolls, 1995), the information theoretic analyzes described here were based on the information available from the firing rate. The period in which the firing rate was measured was the poststimulus period 100 to 600 ms with respect to the onset of the visual stimulus, as most of the information about which stimulus was seen is available in this period (Tovee et al., 1994; Tovee and Rolls, 1995).

Raw Information Measures

If each stimulus s were to evoke its own response r (or its own set of unique responses), then on measuring r one would ascertain s and thus gain $I(s) = -\log_2 P(s)$ bits of information, where $P(s)$ is the probability of occurrence of a particular stimulus or event s . If instead, as happens in general, the same response can sometimes be shared, with different probabilities, by several stimuli, the probabilistic stimulus-response relation will be expressed by a table of probabilities $P(s, r)$, or, equivalently, of conditional probabilities $P(s | r) = P(s, r)/P(r)$. The information gain about a single s on measuring r , can be assessed as follows.

The total amount of information or entropy in a set of spike trains R is

$$H(R) = -\sum_r P(r) \log_2 P(r). \quad (1)$$

We can separate this out into components related to each individual stimulus s by noting that

$$P(r) = \sum_s P(s) P(r | s), \quad (2)$$

so that

$$H(R) = -\sum_s P(s) \sum_r P(r | s) \log_2 P(r). \quad (3)$$

This enables us to identify the total amount of information associated with each stimulus

$$H(s, R) = -\sum_r P(r | s) \log_2 P(r). \quad (4)$$

However, not all of this information is actually about the stimulus. Some of it may concern other completely unrelated things. Therefore, we have to subtract from this total information the amount of information unrelated to the stimulus s . Clearly, any variability in the spiking when the stimulus is held fixed represents information unrelated to the stimulus. The information associated with variations in the spike train for a particular stimulus is just

$$H(R | s) = -\sum_r P(r | s) \log_2 P(r | s). \quad (5)$$

The amount of information actually about the stimulus s is the amount of irrelevant information (5) subtracted from the total information associated with that stimulus (4).

$$I(s, R) = -\sum_r P(r | s) [\log_2 P(r) - \log_2 P(r | s)]. \quad (6)$$

Using $P(s, r) = P(r | s)P(s) = P(s | r)P(r)$ this can be written in different forms as

$$\begin{aligned} I(s, R) &= \sum_r P(r | s) \log_2 \frac{P(s, r)}{P(s)P(r)} \\ &= \sum_r P(r | s) \log_2 \frac{P(r | s)}{P(r)}. \end{aligned} \quad (7)$$

This can be regarded as the difference between the original uncertainty (or *a priori* entropy) and the residual uncertainty after r is known and attains its maximum value $I(s) = -\log_2 P(s)$ only if the probabilistic relation reduces to the deterministic one $P(s | r) = 1$ for $s = s(r)$, and $P(s | r) = 0$ otherwise.

Averaging over different stimuli s in the set of stimuli S one obtains the average information gain about the set of stimuli S present in the spike data R (where R

denotes the set of responses r) as

$$\begin{aligned}
 I(S, R) &= \sum_s P(s) I(s, R) \\
 &= \sum_{s,r} P(s, r) \log_2 \frac{P(s, r)}{P(s)P(r)} \\
 &= \sum_{s,r} P(s)P(r | s) [\log_2 P(s | r) - \log_2 P(s)].
 \end{aligned}
 \tag{8}$$

In the results we show both $I(s, R)$, the information available in the responses of the cell about each individual visual stimulus s , and $I(S, R)$, the average (or mutual) information across all stimuli that is provided about which of the set of stimuli was presented.

In evaluating the information content from the data recorded, the neuronal responses were simply quantified by the number of spikes within a preset time period, 100 to 600 ms poststimulus unless otherwise stated (a unidimensional measure based on a firing-rate measurement). We did not perform a principal component analysis of the time course of the neuronal responses to quantify any information that might be available in the temporal pattern of the spike arrival times because we have performed this previously and have found that most of the information is available in the firing rate, which is close to the first principal component of the neuronal response, (see Tovee et al., 1993, 1994; Tovee and Rolls, 1995; Rolls and Tovee, 1995b). Measured in this way, responses are not continuous but span a discrete set of responses R (although the integers are expressed as firing rates, in spikes/sec). Particularly because of the limited number of trials of data available, R needs to be further quantized into a smaller number of firing rate bins (Optican and Richmond, 1987; Treves and Panzeri, 1995). We perform the quantization as follows. The original data are represented by the number of spikes n^k -recorded in trial k within the prescribed window and are therefore positive integers. Their range is divided into a preselected number D of bins (we usually used $D = 15$; see Optican and Richmond, 1987), with the bin limits selected so that each bin contains the same number of trials (within ± 1). For example, if 100 trials have to be allocated to 15 bins, the first bin extends from zero to midway between the seventh and eighth responses, the second from there to midway between the fourteenth and fifteenth responses, and so on until the last bin, which extends from midway between the ninety-fourth and ninety-fifth responses to plus infinity. A smoothing procedure is applied by convolving

the individual values n with a Gaussian kernel whose width is proportional to the square root of each rate value (the proportionality factor is set such that on average the smoothing widths match the standard deviations σ_n of the values for each stimulus, which appear to scale with roughly the square root of the mean rates). The result, normalized by dividing by the total number of trials, is quantized into the bins defined above, the area within each bin being used as an estimate of the joint probability $P(s, r)$, where r corresponds to one of the D response bins. Summing over all stimuli gives $P(r) = \sum_{s \in S} P(s, r)$.

Information values are in general dependent on the smoothing and binning procedures adopted and, most important, on the number of bins D and on the smoothing widths. The information values reported here are therefore to be considered as measures relative to the present regularization methods. The parameters chosen are a compromise between the need to maintain the originally continuous nature of the data, which would require fine bins and little smoothing, and therefore high D and small widths, and the need to control finite-size distortions in the information estimate, which as discussed by Treves and Panzeri (1995) requires that either D be small (for example, in the absence of smoothing, D should be smaller than the number of trials per stimulus) and/or the widths be large.

Information Estimates Corrected for the Limited Number of Trials

The procedure introduced so far for estimating the probability $P(s, r)$ of a particular response is rather simple. In practice, because of the limited number of trials that can be collected, the various probability tables are not available, and one can at best approximate them with frequency tables $P_N(r | s)$, $P_N(r)$ computed on the basis of a (limited) number of trials N . If N is very large, the frequencies should get close to the underlying probabilities, but for any finite N there will be a discrepancy, which will result in an error in the estimated information gain. Because information quantities depend on probabilities not in a linear but in a greater than linear manner, the error deriving from this *limited sampling* does not cancel out on averaging many measurements; it is, instead, usually biased upward, resulting in an (average) overestimate of the information gain.

The net bias, or average error (usually an overestimating error), can be expressed analytically as a formal

expansion in $1/N$, and the first few terms (in particular, the very first) of this expansion can be evaluated directly (Treves and Panzeri, 1995) in a variety of situations. Simulation experiments have shown that the first term in the expansion is responsible for most of the discrepancy between the raw and correct information measures, whereas successive terms do not in fact correlate with the remainder of the discrepancy. This first term can then be subtracted from the raw estimates, to produce corrected estimates of $I(S, R)$ and $I(s, R)$. This procedure has been shown to improve significantly the reliability of information estimates based on limited data samples, over various alternative empirical remedies that have been proposed and that do not rely on analytical results (Chee-Orts and Optican, 1993; Hertz et al., 1994). The simplest of these alternatives is to subtract from the raw information estimate a correction derived from a *random shuffling* (so-called *bootstrap*) procedure (Optican et al., 1991). We use here the analytically based procedure, referring back to Panzeri and Treves (1996) for a more detailed explanation of how the correction term is computed from the data. (We note that this provides the average correction, so that sometimes the correction could be larger than the computed raw information; in this case, the corrected information should be set to 0.) With respect to the correction based on the shuffling procedure, which we used in previous investigations (Tovee et al., 1993; Tovee and Rolls, 1995), one should note that, while in several cases it yields results for $I(S, R)$ that are very close to those obtained with the present procedure, it cannot be applied to compute the stimulus-specific information $I(s, R)$, simply because the random shuffling mixes responses occurring to different stimuli. This is one case, therefore, in which it was essential to develop a novel procedure to correct for limited sampling.

It should be noted, finally, that an information estimate based on a quantized response set tends to grow with the size D of the set for D small, until it saturates once the width of the D bins becomes negligible with respect to the standard deviation of the responses for each stimulus. We used $D = 15$ (similar to the value of 12 used by Optican and Richmond, 1987) after checking that no marked increase in $I(s, R)$ resulted from using larger D values.

Information Estimates for the 20 ms Response Window

The above procedure was modified for deriving information measures from the response recorded in a

20 ms window, chosen 100 to 120 ms poststimulus onset. The important point is that these cells emitted at most a few spikes within the short 20 ms window (even though their response in that window tended to be higher than in any other 20 ms period). Therefore, the raw response—that is, the number of spikes emitted—already fell into few (typically four to five) natural bins, and there was no need to either regularize the raw responses or to choose the bins (for example, to be equipopulated). Results obtained for the 20 ms window are then independent of any regularization procedure. Note that a regularization similar to that described above, when applied to the 20 ms data, was found to yield reduced information values (by around 10 percent). This is understandable in light of the fact that the regularization distorts more the discrete nature of the raw responses for the 20 ms period. It is believed that it remains meaningful to produce a rough comparison of the results obtained for 20 ms without regularization, with those obtained for the 500 ms window, because the regularization, which is present for the latter, is expected to have relatively minor effects in that case, in which raw responses are already closer to having a continuous distribution.

A further comparison was produced between the information measures derived from the 20 ms window, and the information expected to be contained in the response of cells that fired to each stimulus with exactly the same mean rate as recorded experimentally, but with either a Poisson or a strictly periodic behavior. The important parameter that differentiates these two theoretical examples (the actual data is in most cases intermediate between them) is the variability, within trials and between trials, of the response. For Poisson firing this variability is high because the probability of emitting n spikes in time t , when firing at mean rate r_i , is $P(n) = (r_i t)^n \exp - (r_i t) / n!$, which is a distribution with mean $r_i t$, standard deviation equal to the square root of the mean, and coefficient of variation equal to the inverse square root, and therefore typically somewhat above 1. For periodic firing the variability is minimal because the number of spikes recorded on any trial may be only either the largest integer below $r_i t$ or the smallest integer greater above it, depending on where the window falls with respect to the last spike fired. The information content of these two theoretical examples can be derived, then, from knowing only the means r_i and not the detailed distribution of responses across trials. To obtain a fair comparison with the actual information measures derived from the data, a correction was applied for the average discrepancy between

the experimental mean rates to each stimulus and the true mean rates. Such a correction (which turned out to be generally small, in the 0.02 bits range) was dependent, of course, on the measured standard deviation of the means, and was the equivalent procedure to the finite sampling correction operated for the actual information estimates. It resulted in the theoretical information for Poisson and periodic firing being occasionally negative, as a result, again, of subtracting an average bias from a fluctuating raw estimate. Full details on the derivation of the information measures for the Poisson and periodic cases are reported elsewhere (Treves et al., in preparation).

Sparseness of the Representation

The sparseness a of the representation of a set of visual stimuli provided by these neurons can be defined and was calculated as

$$a = \left(\sum_{i=1,n} (r_i/n) \right)^2 / \sum_{i=1,n} (r_i^2/n),$$

where r_i is the firing rate to the i th stimulus in the set of n stimuli. The sparseness has a maximal value of 1.0. This is a measure of the extent of the tail of the distribution, in this case of the firing rates of the neuron to each stimulus. A low value indicates that there is a long tail to the distribution, equivalent in this case to only a few stimuli with high firing rates. If these neurons were binary (either responding with a high firing rate or not responding), then a value of 0.2 would indicate that 20% of the stimuli produced high firing rates in a neuron, and 80% produced no response. In the more general case of a continuous distribution of firing rates, the sparseness measure a still provides a quantitative measure of the length of the tail of the firing-rate distribution (Treves and Rolls, 1991). This measure of the sparseness of the representation of a set of stimuli by a single neuron has a number of advantages. One is that it is the same measure of sparseness that has proved to be useful and tractable in formal analyzes of the capacity of neural networks that use an approach derived from theoretical physics (see Treves, 1990; Treves and Rolls, 1991, 1994; Rolls and Treves, 1990). A second is that it can be applied to neurons that have continuously variable (graded) firing rates and not just to firing rates with a binary distribution (for example, 0 or 100 spikes/sec) (Treves and Rolls, 1991). A third is that it makes no assumption about the form of the firing-rate

distribution (such as binary, ternary, exponential, and so on), and can be applied to different firing-rate distributions (Treves and Rolls, 1991). Fourth, it makes no assumption about the mean and the variance of the firing rate. Fifth, the measure does not make any assumption about the number of stimuli in the set and can be used with different numbers of test stimuli. Its maximal value is always 1.0, corresponding to the situation when a neuron responds equally to all the stimuli in a set of stimuli. The use of this measure of sparseness in neurophysiological investigations has the advantage that the neurophysiological findings then provide one set of the parameters useful in understanding theoretically how the system operates (Treves and Rolls, 1991, 1994; Rolls and Treves, 1990).

As described in the Results section, a measure of the response sparseness was also calculated, in which the spontaneous firing rate was subtracted from the firing rate (and the response was clipped to zero). This corresponds to the intuition of some neurophysiologists that it is changes from the spontaneous firing rate that are important, and that is why we show it. However, this response sparseness measure a_r does have problems if the neuron decreases its firing rate below the spontaneous rate for some stimuli, in which case it may be more appropriate to calculate the responses as changes from the lowest firing rate to any stimulus. For these reasons, more emphasis is placed here on the sparseness measure a as defined in the previous paragraph based on the absolute firing rates. That measure a also has the advantage that in models of neuronal networks, it is the absolute firing rate of the input to each synapse that must be considered when quantifying measures such as the capacity of the network and the interference between stimuli (Treves and Rolls, 1991; Rolls and Treves, 1990).

Results

It was possible to complete the analyses described here using 65 stimuli (23 face and 42 nonface) for 14 different face-selective neurons recorded in the same monkey. Each stimulus was repeated for about 8 to 10 trials with each cell. To confirm that the 14 cells were a representative sample, we compared results (for a subset of stimuli) with a larger set of 25 cells, recorded in part in a second monkey, and obtained consistent results.

The value for the average information, $I(S, R)$ about a set of 20 face stimuli available in the responses of the

neurons, is shown for each of the 14 neurons in histogram form in Fig. 1a. It is clear that each of these neurons had reasonable amounts of average information about which one of the set of 20 faces had been shown on an individual trial (Fig. 1a). The neurons reflected in their firing rates (in this case in the period 100 to 600 ms poststimulus) on average 0.36 bits of information about which face was presented. In contrast, the same general population of neurons has much less information about which (of 20) nonface stimuli has been seen (on average 0.07 bits) (Rolls and Tovee, 1995b). The mutual information $I(S, R)$ shows the average of all the information contained in the responses to the individual stimuli (corrected for the minor differences in the number of trials). The value for the average information available in the responses of the neurons $I(S, R)$, about which of a set of 65 stimuli (including 23 faces) had been shown, is shown for each of the 14 neurons in histogram form in Fig. 1b.

In order to understand the representation of individual stimuli by individual cells, the information $I(s, R)$ available in the neuronal response about each stimulus in the set of stimuli S was calculated. First, we show in Fig. 2a the firing rate response profile of a single neuron (Am242) to the set of 65 stimuli. The spontaneous firing rate of this neuron was 20 spikes/sec, and the histogram bars indicate the change of firing rate from the spontaneous value, produced by each stimulus. Stimuli that are faces are marked *F* or *P* if they are in profile. *B* refers to images of scenes that included either a small face within the scene, sometimes as part of an image that included a whole person, or other body parts, such as hands (*H*) or legs. The nonface stimuli are unlabeled. The neuron responded best to three of the faces (profile views), had some response to some of the other faces, and had little or no response, and sometimes had a small decrease of firing rate below the spontaneous firing rate, to the nonface stimuli. The sparseness value a for this cell across all 65 stimuli was 0.70, and the response sparseness a_r (based on the evoked responses minus the spontaneous firing of the neuron) was 0.25.

Second, we show for the same neuron the information $I(s, R)$ available in the neuronal response about each of the stimuli in the set of 20 face stimuli (Fig. 2b). The average information $I(S, R)$ about this set of 20 faces for this neuron was 0.55 bits. It is shown in Fig. 2b that 2.2, 2.0, and 1.5 bits of information were present about the three face stimuli to which the neuron had the highest firing-rate responses. The neuron conveyed

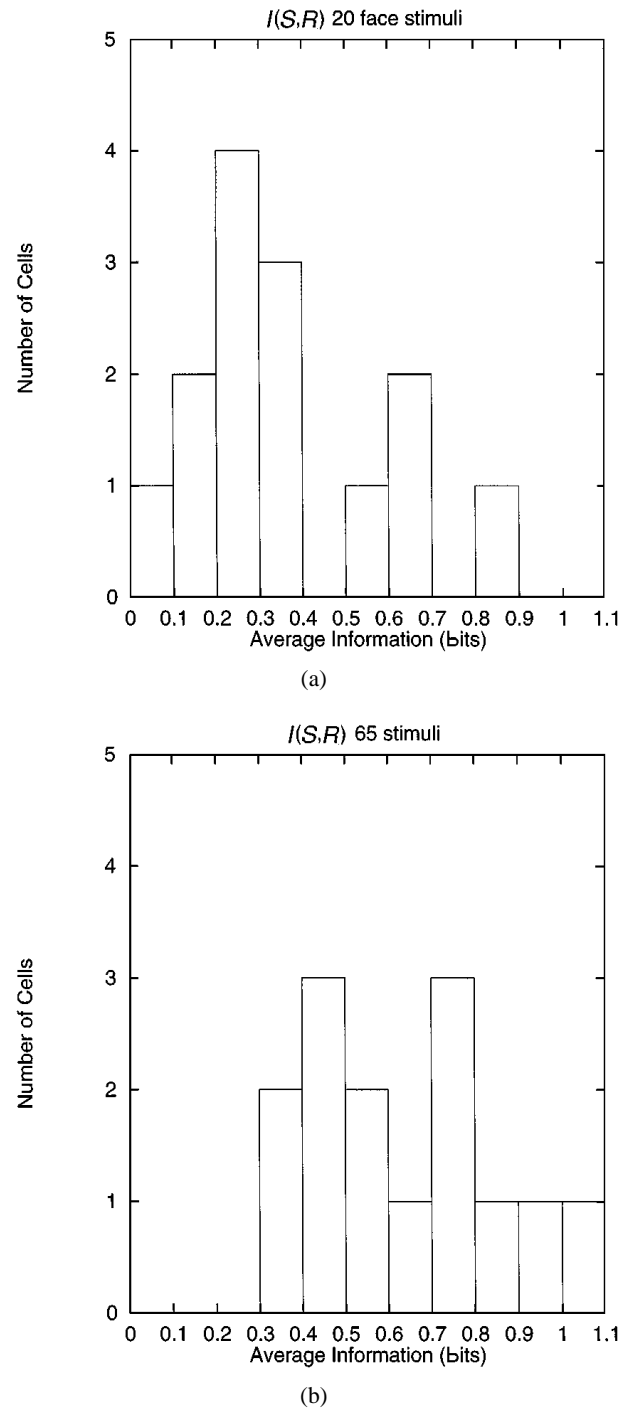


Figure 1. a. Histogram showing the values for the average information $I(S, R)$ about a set of 20 face stimuli available in the responses of each of the 14 neurons. b. The values for the average information about the larger set of 65 face and nonface stimuli.

some but smaller amounts of information about the remaining face stimuli. The average firing rate of this neuron to these 20 face stimuli was 54 spikes/sec. It is clear from Fig. 2b that little information was available from the responses of the neuron to a particular face

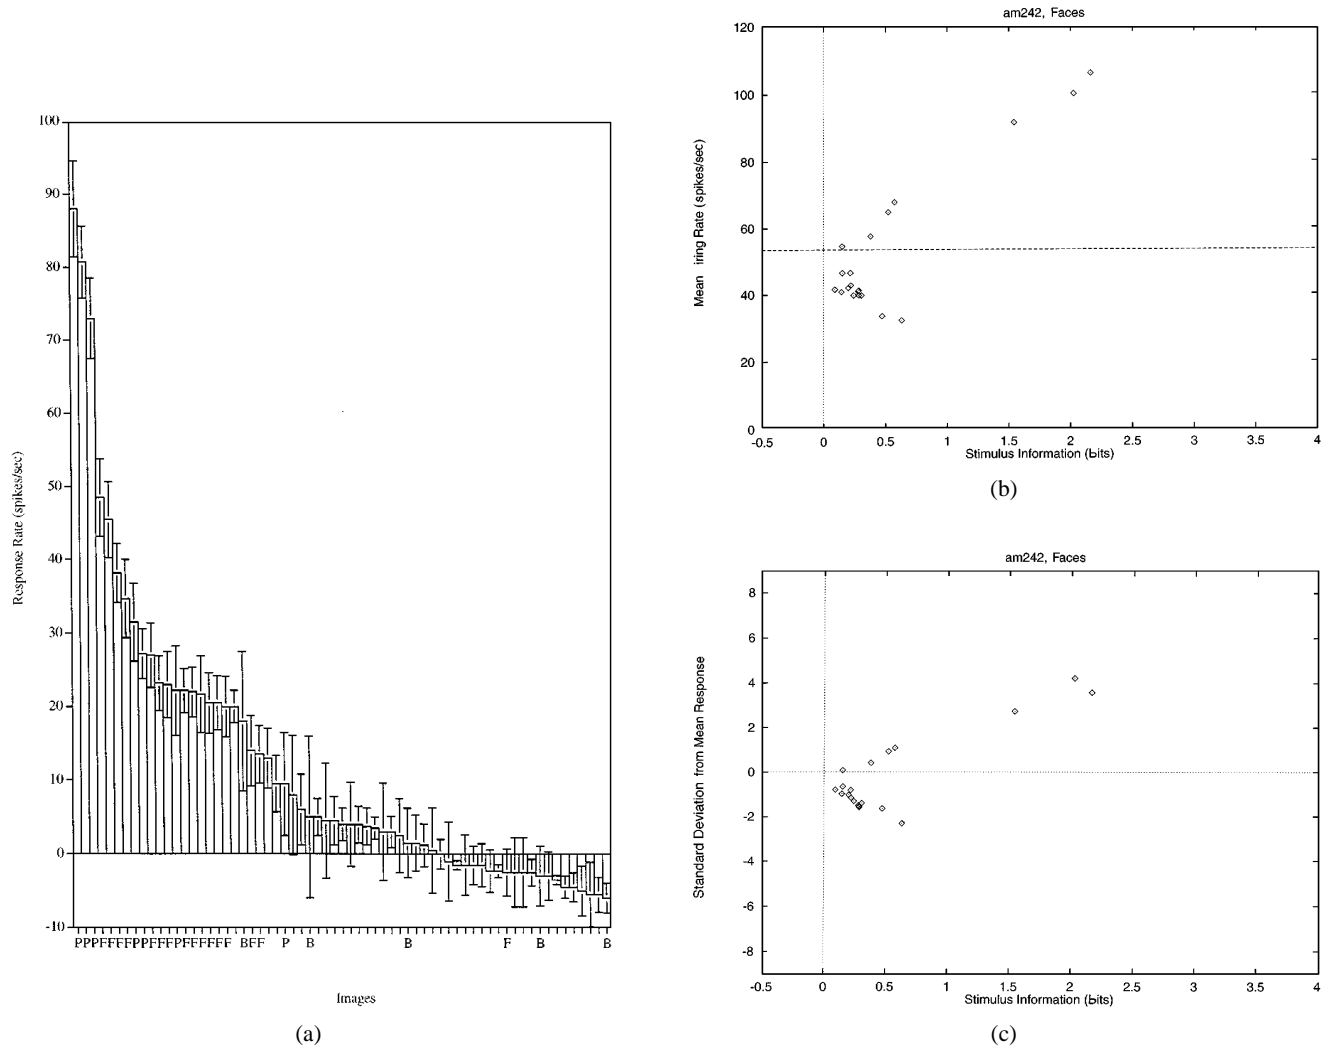


Figure 2. Neuron Am242. a. The distribution of firing rates to all 65 stimuli. The firing rate of the neuron is shown on the ordinate, the spontaneous firing rate of the neuron was 20 spikes/sec, and the histogram bars are drawn to show changes of firing rate from the spontaneous rate (that is, neuronal responses) produced by each stimulus. Stimuli that are faces are marked F or P if they are in profile. B refers to images of scenes that included either a small face within the scene, sometimes as part of an image that included a whole person or other body parts, such as legs. The nonface stimuli are unlabeled. The error bars denote the standard deviation of the mean firing rate to each stimulus. b. The information $I(s, R)$ available in the response of the same neuron about each of the stimuli in the set of 20 face stimuli (abscissa), with the firing rate of the neuron to the corresponding stimulus plotted as a function of this on the ordinate. The horizontal dashed line is the mean firing rate of the cell. c. The number of standard deviations the response to a stimulus was from the average response to all stimuli (see text, z score) plotted as a function of $I(s, R)$, the information available about the corresponding stimulus. d. The information $I(s, R)$ available in the response of the same neuron about each of the stimuli in the set of 23 face and 42 nonface stimuli (abscissa), with the firing rate of the neuron to the corresponding stimulus plotted as a function of this on the ordinate. The horizontal dashed line is the mean firing rate of the cell. The 23 face stimuli in the set are indicated by a diamond, and the 42 nonface stimuli by a cross. e. Histogram showing the frequency of occurrence of different values of $I(s, R)$ for the same neuron. Solid lines-23 face stimuli; dashed line-23 face plus 42 nonface stimuli. f. The probability distribution of firing rates for neuron am242 to the whole set of 65 stimuli. An exponential distribution fitted to this mean rate is shown by the line. The sampling period is 500 ms, starting at 100 ms after the onset of the stimuli. g. The variance of the firing rate of cell am242 as a function of the firing rate, shown for a 20 ms period starting 100 ms after the stimulus was turned on. Each point shows the rate and variance for one stimulus. For comparison, the variance that would be expected from a Poisson process with the same rate is indicated by the line. Where values for different stimuli take the same value, the number of such values is indicated by the number beside the point. h. The variance of the firing rate of cell am242 as a function of the firing rate, shown for a 500 ms period starting 100 ms after the stimulus was turned on. Each point shows the rate and variance for one stimulus. For comparison, the variance that would be expected from a Poisson process with the same rate is indicated by the line.

(Continued on next page.)

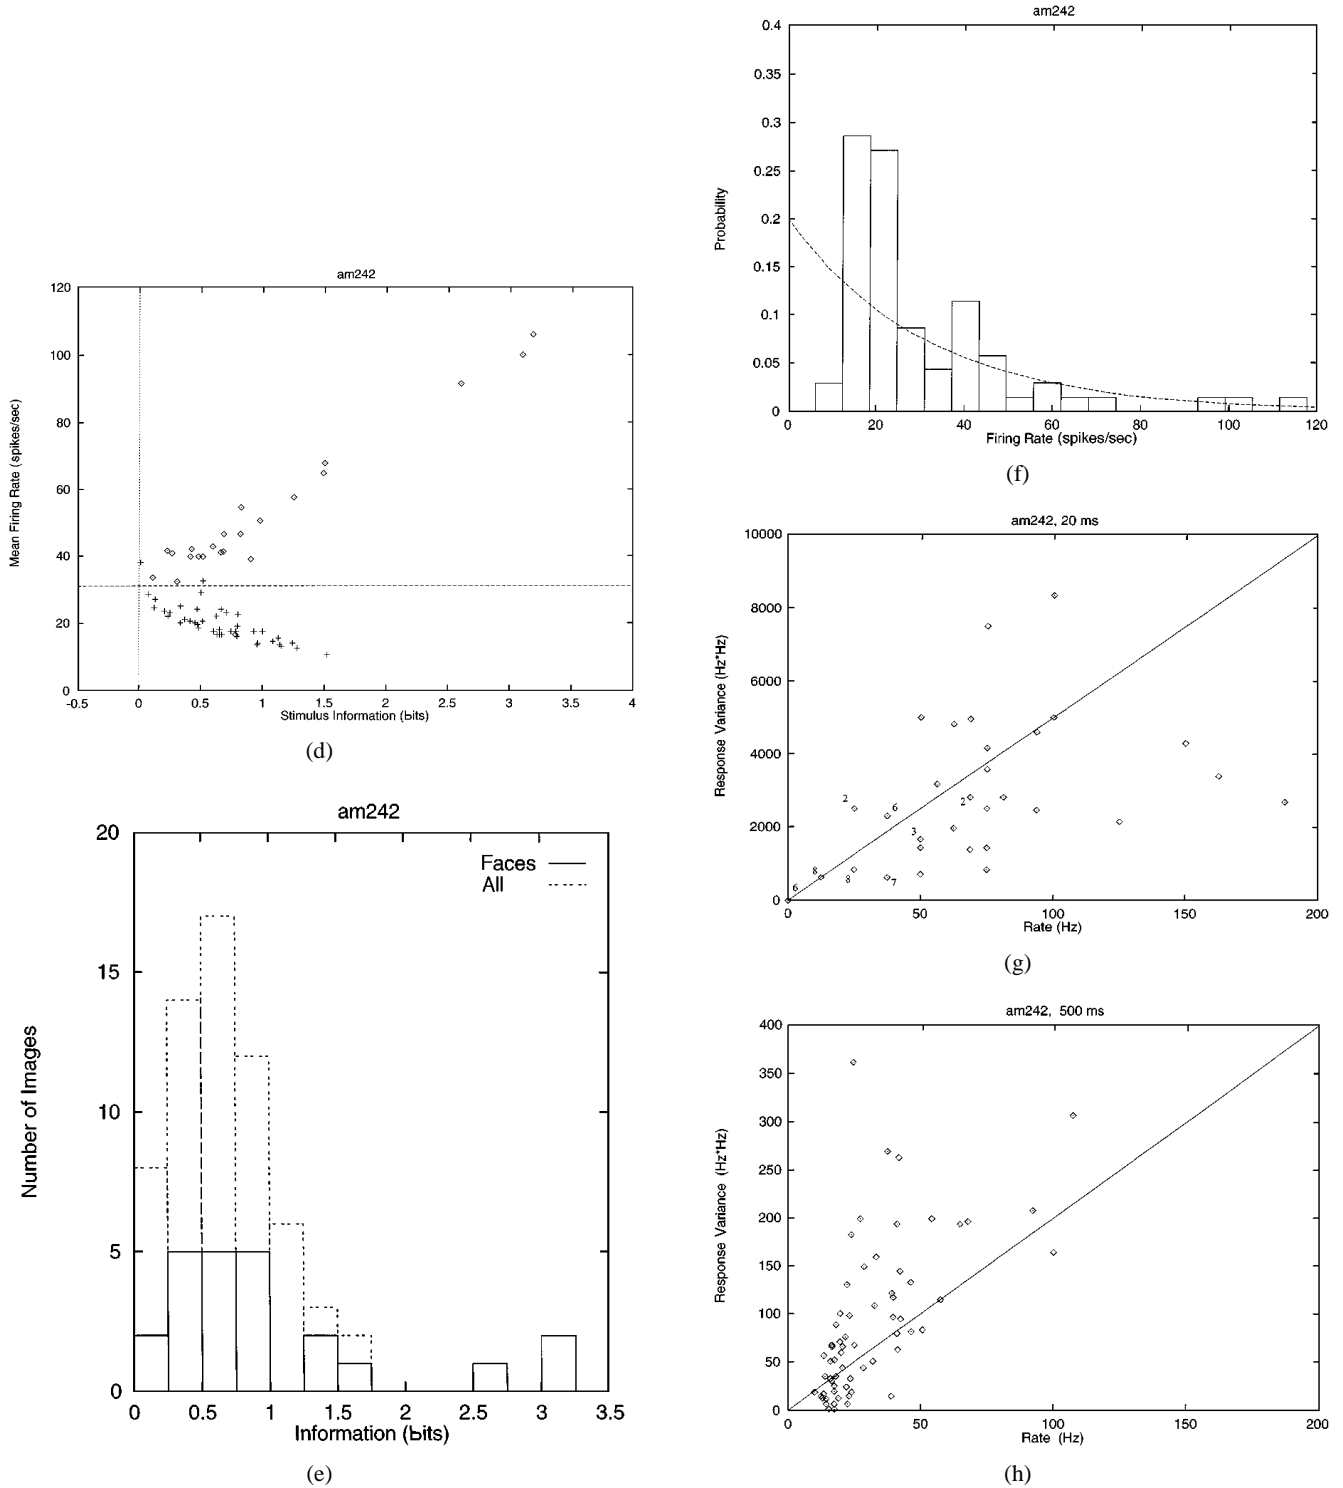


Figure 2. (Continued).

stimulus s if that response was close to the average response of the neuron across all stimuli. At the same time, it is clear from Fig. 2b that information was present depending on how far the firing rate to a particular stimulus s was from the average response of the neuron to the stimuli. Of particular interest, it is evident

that information is present from the neuronal response about which face was shown if that neuronal response was below the average response, as well as when the response was greater than the average response.

One intuitive way to understand the data shown in Fig. 2b is to appreciate that low-probability firing-rate

responses, whether they are greater than or less than the mean response rate, convey much information about which stimulus was seen. This is, of course, close to the definition of information. Given that the firing rates of neurons are always positive and follow an asymmetric distribution about their mean, it is clear that deviations above the mean have a different probability to occur than deviations by the same amount below the mean. One may attempt to capture the relative likelihood of different firing rates above and below the mean by computing a z score obtained by dividing the difference between the mean response to each stimulus and the overall mean response, by the standard deviation of the response to that stimulus. The greater the number of standard deviations (that is, the greater the z score) from the mean response value, the greater the information might be expected to be. We therefore show in Fig. 2c the relation between the z score and $I(s, R)$. The z score was calculated by obtaining the mean and standard deviation of the response of a neuron to a particular stimulus s and dividing the difference of this response from the mean response to all stimuli by the calculated standard deviation for that stimulus. This results in a C-shaped curve in Figs. 2c and 2b, with more information being provided by the cell the further its response to a stimulus is in spikes per second or in z scores either above *or below* the mean response to all stimuli (which was 54 spikes/sec).

As noted in the discussion, this analysis shows what information is *available* in the responses of the cell to the different stimuli. It is a separate issue of whether the neurons that receive inputs from these visual cells can make use of the information potentially available in a response to a stimulus that is below the mean response of the cell to visual stimuli and, if so, what type of neuronal networks would best enable this information to be utilised.

We show in Fig. 2d the information $I(s, R)$ about each stimulus in the set of 65 stimuli for the same neuron, am242. The 23 face stimuli in the set are indicated by a diamond, and the 42 nonface stimuli by a cross. Using this much larger and more varied stimulus set, which is more representative of stimuli in the real world, a C-shaped function again describes the relation between the information conveyed by the cell about a stimulus and its firing rate to that stimulus. In particular, this neuron reflected information about most, but not all, of the faces in the set, those faces that produced a higher firing rate than the overall mean firing rate to all the 65 stimuli, which was 31 spikes/sec. In

addition, it conveyed information about the majority of the 42 nonface stimuli by responding at a rate below the overall mean response of the neuron to the 65 stimuli. This analysis usefully makes the point that the information available in the neuronal responses about which stimulus was shown is relative to (dependent on) the nature and range of stimuli in the test set of stimuli.

To further relate the information $I(s, R)$ about each stimulus in the set of 65 stimuli to the distribution of responses constituting the representation of the set of stimuli, we show for the same neuron, am242, a histogram of the frequency of occurrence of different values for $I(s, R)$ in Fig. 2e. (Each value included in the histogram is the information for one stimulus.) Figure 2e makes it clear that, for this cell, little information was reflected in the responses to most of the stimuli and that for a few, much information was available. This reflects the relatively sharp tuning of this neuron, shown in the firing-rate response profile provided in Fig. 2a.

Another way to plot the data shown in Fig. 2a is to show the probability distribution of particular firing rates when the whole set of 65 stimuli was shown. It is shown in Fig. 2f that neuron am242 is frequently firing with low rates (reflecting the fact that many of the stimuli produced low firing rates) and fires with high rates with a decreasing probability (in that few stimuli elicit high firing rates). To show whether the probability of high rates decreases approximately exponentially (as suggested by Levy and Baxter, 1996; Baddeley, 1996), we also show an exponential curve with the same mean firing rate. The fit for this single cell is not especially good (Kolmogorov-Smirnov goodness of fit, significantly different from exponential at $p < 0.01$). However, this issue is considered again for the whole population of cells in Fig. 8a.

The variance of the responses of neuron am242 as a function of its firing rate is shown in Figs. 2g and 2h. With 500 ms epochs, the variance is in many cases greater than would be expected from a Poisson process. This could be due to the fact that these cortical neurons tend to have high rates early in their response period (for example, from 100 to 200 ms poststimulus), while after that the rate tends to drop to approximately half of its peak rate (see, for example, Tovee et al., 1993, Fig. 3). This could tend to increase the estimated variance with respect to the mean rate estimated over 500 ms.

A similar analysis for a cell (am240) with a more distributed representation is shown in Fig. 3. The response profile of the cell to the 65 different visual stimuli is

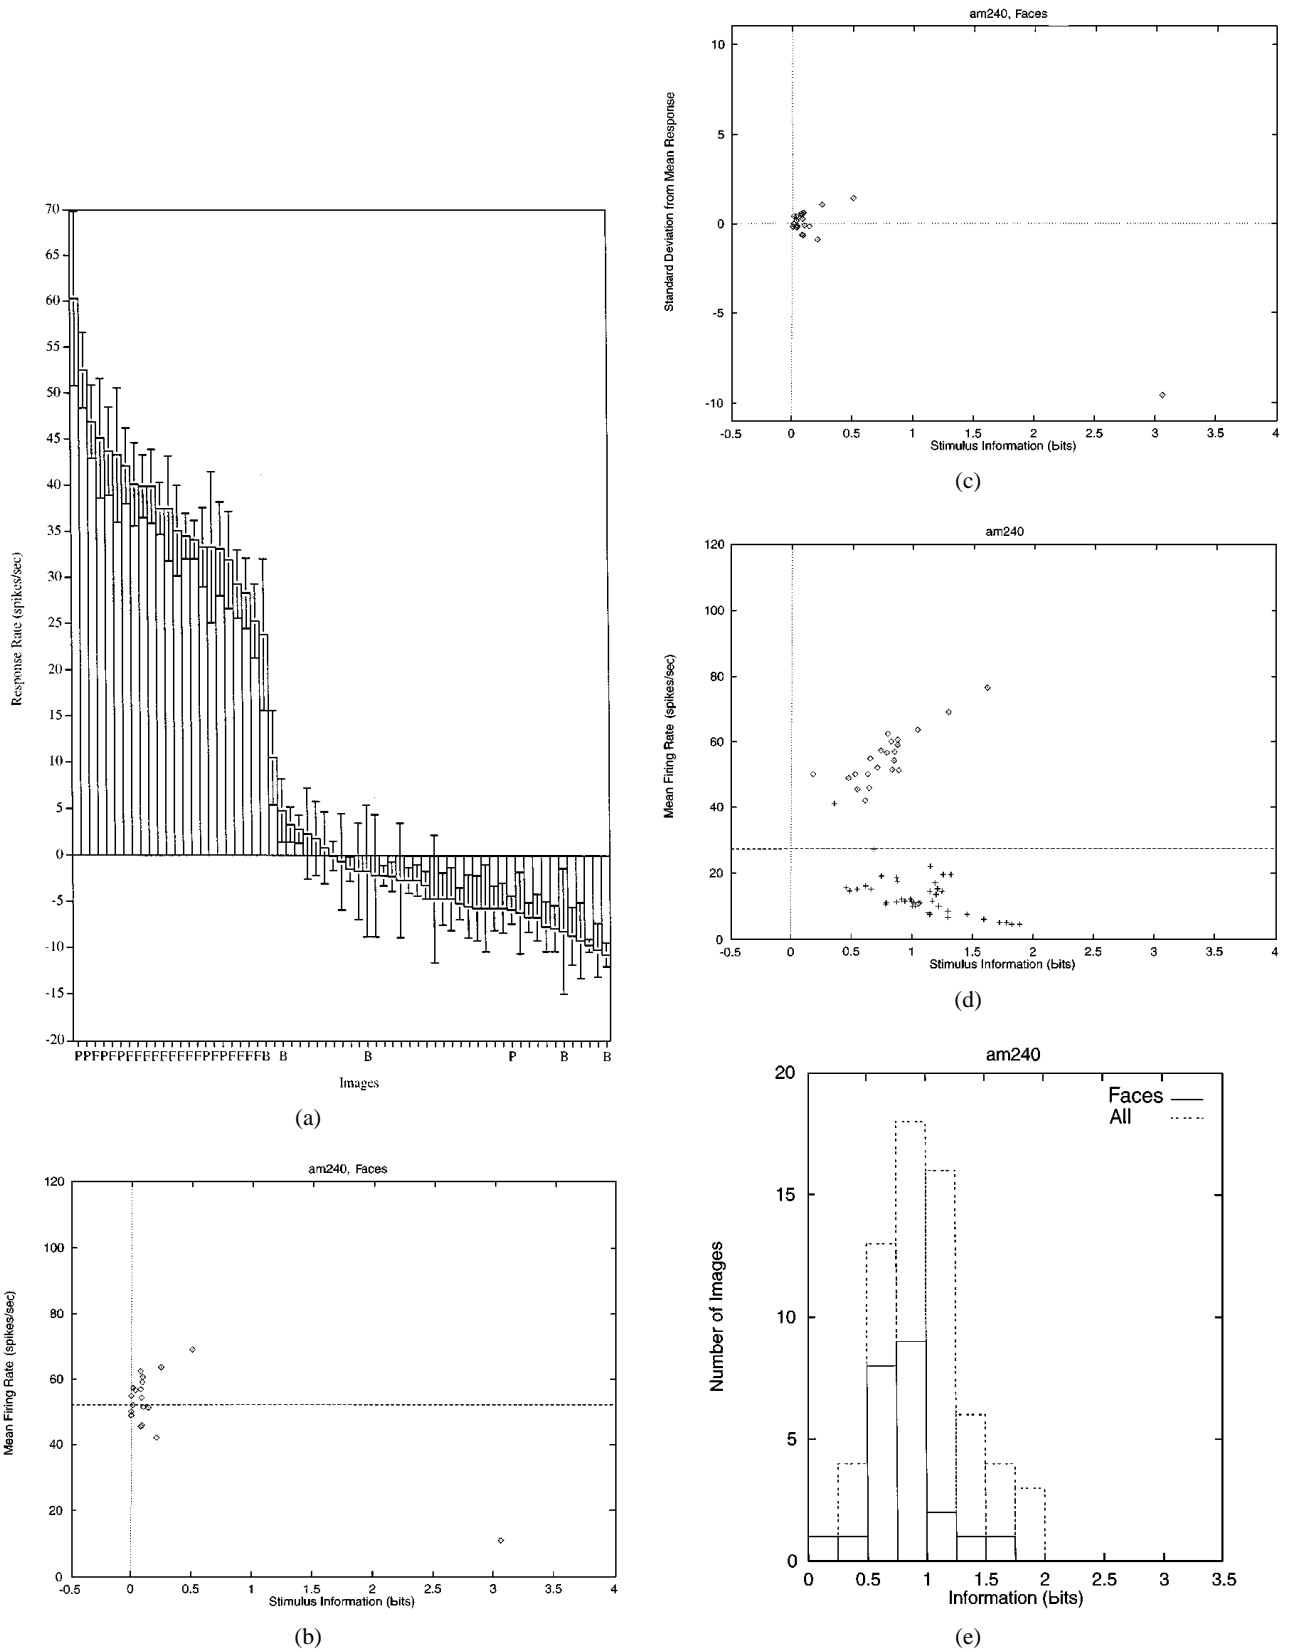


Figure 3. a-e. Similar graphs to those shown in Figs. 2a-e for another neuron, am240.

shown in Fig. 3a. The sparseness value a for this cell across all 65 stimuli was 0.62, and the response sparseness a_r was 0.35. The cell responded well to many of the faces in the set but not to all of them and had little response to the nonface stimuli. The average information $I(S, R)$ about this set of 65 stimuli for this neuron was 0.93 bits.

For the same neuron (am240) the information $I(s, R)$ available in the neuronal response about each of the stimuli in the set of 20 face stimuli is shown in Fig. 3b. The average information $I(S, R)$ about this set of 20 faces for this neuron was 0.24 bits. Because this neuron had high firing-rate responses to 19 of the face stimuli in this set of 20 but did not respond to one of the faces in the set of 20, the amount of information available about the one outlier was high, 3.1 bits. The next-highest information provided when one of the other faces was shown was 0.5 bits, and to the next stimulus, 0.25 bits. This emphasizes that information about which of a set of stimuli has been shown can be available in a lack of or in a small firing-rate response to one or a few of the stimuli in the set being considered, if the firing rate response to most of the stimuli in the set is large. The average firing-rate response of this neuron to these 20 face stimuli was 52 spikes/sec. It is clear also from Fig. 3b that little information was available from the responses of the neuron to a particular face stimulus if that response was close to the average response of the neuron across all stimuli.


Again, the implication of the data shown in Fig. 3b is that low-probability firing-rate responses, whether they are greater than or less than the mean response rate to the particular set of stimuli considered, convey much information about which stimulus was seen. This is supported by the analysis shown in Fig. 3c of the relation between the z score and $I(s, R)$ for the same set of 20 faces. The z score analysis shown in Fig. 3c indicates that the greater the number of z scores the neuronal response is from the mean response of the cell to all stimuli (and therefore the less probable is the response), the greater is the information provided about that stimulus, even when the response to a stimulus is below the average response to all stimuli. This results in a C-shaped curve in Figs. 3c and 3b, with more information being provided by the cell the further its response to a stimulus is in spikes per second or in z scores either above or below the mean response to all 20 face stimuli (which was 52 spikes/sec).

We show in Fig. 3d the information $I(s, R)$ about each stimulus in the set of 65 stimuli for the same


neuron, am240. The 23 face stimuli in the set are indicated by a diamond, and the 42 nonface stimuli by a cross. It is clear that relative to this much larger and more varied stimulus set, which is more representative of stimuli in the real world, this neuron reflected information about most (all but one) of the faces in the set with responses that were greater than the overall mean response to all the 65 stimuli, which was 27 spikes/sec. Because this neuron responded well to the majority of the faces and had little response to the majority of the nonface stimuli, the overall mean response of the neuron lay between the response to most of the faces and the response to most of the nonfaces. The effect of this is that, as the graph shows, a considerable amount of information was available about which of the nonfaces in the set had been shown, as is evident by the lower part of the C-shaped curve in Fig. 3d. This analysis again usefully makes the point that the information available in the neuronal responses about which stimulus was shown is relative to (dependent on) the nature and range of stimuli in the test set of stimuli. The average information $I(S, R)$ about this set of 65 stimuli for this neuron was 0.93 bits.

For the same neuron, am240, a histogram of the frequency of occurrence of different values for $I(s, R)$ is shown in Fig. 3e, for the set of 65 stimuli. For this cell, there are few outliers in the distribution—that is, there are no stimuli for which very large amounts of information become available when that stimulus is shown, and moderate amounts of the information available were reflected in the responses to most of the stimuli. This reflects the fact that some information is available about most of the face stimuli because the response to most of them is above the mean response of the neuron to all stimuli; and that some information is available about most of the nonface stimuli because the response to most of them is below the mean response to all of the stimuli (cf. Figs. 3a and d).


To illustrate further how these information analyzes can be interpreted with neurons with different types of tuning, we show in Fig. 4 similar analyzes for another cell (am229) that had intermediate tuning between the types illustrated in Figs. 2 and 3. The response profile of the cell to the 65 different visual stimuli shown in Fig. 4a indicates that the cell had some tuning within the class faces, with large responses to some but not to other faces, and had little response to the nonface stimuli. The sparseness value a for this cell across all 65 stimuli was 0.60, and the response sparseness a_r was 0.30. The average information




(a)




(b)



(c)



(d)



(e)

Figure 4. a–e. Similar graphs to those shown in Figs. 2a–e for another neuron, am229. f. The information $I(s, R)$ available in the response of the same neuron, when firing rates (plotted on the ordinate and used to calculate information on the abscissa) are measured over 20 ms rather than 500 ms. The same set of 20 stimuli is used (diamonds). The stars indicate the amounts of information that would have been provided by a theoretical cell firing with Poisson statistics and with the same distribution of mean rates, and the x 's those of a theoretical cell firing with exact periodicity (see text).

(Continued on next page.)




Figure 4. (Continued).

$I(S, R)$ about this set of 65 stimuli for this neuron was 0.79 bits.

For the same neuron (am229) we show the information $I(s, R)$ available in the neuronal response about each of the stimuli in the set of 20 face stimuli in Fig. 4b. The average information $I(S, R)$ about this set of 20 faces for this neuron was 0.61 bits. Because this neuron had a wide distribution of firing rates to the different members of this set of 20 faces, the information available about each of the stimuli in the set of 20 faces had a wide distribution, with considerable amounts of information available both for stimuli with responses above the mean response for faces of 70 spikes/sec, and for stimuli with responses below the mean response to faces. It is clear also from Fig. 4b that little information was available from the responses of the neuron to a particular face stimulus if that response was close to the average response of the neuron across all stimuli.

Again, the z score analysis, shown in Fig. 4c, indicates that in most cases the greater the number of z scores the neuronal response is from the mean response of the cell to all stimuli (and therefore the less probable is the response), the greater is the information provided about that stimulus, even when the response to a stimulus is below the average response to all stimuli (which was 70 spikes/sec).

We show in Fig. 4d the information $I(s, R)$ about each stimulus in the set of 65 stimuli for the same neuron, am229. The 23 face stimuli in the set are indicated by a diamond, and the 42 nonface stimuli by a cross. It is seen that for this cell there was a somewhat greater degree of mixing between the responses to faces and to nonfaces. Still, the information versus rate curve has the characteristic C shape. The overall mean firing rate

to all the 65 stimuli, was 38 spikes/sec. The average information $I(S, R)$ about this set of 65 stimuli for this neuron was 0.79 bits.

For the same neuron, am229, a histogram of the frequency of occurrence of different values for $I(s, R)$ is shown in Fig. 4e, for the set of 65 stimuli. For this cell, the stimuli that produced information-rich responses were both faces and nonfaces, and high values of information (1.5 to 2.0 bits) were conveyed by high rates for the former and by low rates for the latter.

Figure 4f shows, for the same neuron, am229, the information $I(s, R)$ calculated from the firing rates extracted over the 20 ms period from 100 to 120 ms after stimulus onset. Note that the information values are plotted on the same (horizontal) scale as in Fig. 4b, while the vertical (firing rate) scale is expanded to accommodate the higher mean rates recorded in these 20 ms at the peak of the responses than in the whole 500 ms period. The actual information values measured (denoted with diamonds) follow a similar C-shaped curve as those obtained from responses over 500 ms, but the values for individual stimuli are more scattered, reflecting a noisier relation to the underlying C curve, and also the curve itself is on the whole closer to the vertical axis, reflecting the smaller amount of information available from the short 20 ms period than from the long 500 ms period. Averaging across stimuli, the mean rate in the 20 ms period was for this cell 93 spikes/sec (compared with 70 spikes/sec over 500 ms), and the average information in the rates was 0.20 bits (compared with 0.61 bits over 500 ms). The stars in Fig. 4f denote the amount of information available in the response to each individual stimulus, of a theoretical neuron that responded with the same distribution of mean rates to each stimulus as cell am229, but with a Poisson distribution of spikes in the 20 ms window, for each trial, as explained under Methods. The corresponding values follow quite closely another C-shaped curve, which is even closer to the vertical axis, reflecting the fact that Poisson firing, with its high variability, would convey less information than the actual firing mode observed. Averaging across stimuli the theoretical Poisson cell would convey only 0.12 bits of information. Finally, the x 's in Fig. 4f denote the amount of information that would be carried by a neuron firing strictly periodically, again with firing rates set equal to the mean rates observed for cell am229. Information values are in this case again scattered around an underlying C-shaped curve, with a degree of scattering similar to that of the actual information values but

with a curve extending farther away from the vertical axis, reflecting higher values, on average. In fact, this theoretical cell, firing exactly periodically with no variability, would convey on average as much as 0.80 bits in 20 ms. In summary the real neuron am229, in firing with a variability intermediate between (though closer to) that of a Poisson cell and that of a periodically firing cell, conveys intermediate amounts of information about individual stimuli between the two extreme cases of Poisson (low information) and periodic (high information). It is important, however, that the C-shape underlying the information-versus-rate values across stimuli is the same in all three cases. The fact that in all three cases the curve appears to touch the vertical axis close to the overall mean rate is related to the 20 ms period being fairly short, as commented on below. The typical variability of the responses of these cells at different firing rates is shown (for neuron am242) in Figs. 2g and 2h. For 20 ms epochs, the variance is less than would be expected from a Poisson process. The high amounts of information at short epochs are related to this low variability. The way in which the variance depends on the firing rate and the sampling interval shown in Figs. 2g and 2h is typical for these cells.

The distribution of the maximum values of information about any one stimulus in the set of 20 faces available in the responses of the 14 neurons is shown in Fig. 5a and in Table 1. The mean of the maximum values of $I(s, R)$ was 1.8 bits. For comparison, the average information $I(S, R)$ conveyed by this particular set of neurons about the set of 20 faces was 0.36 bits. For the set of 65 stimuli the maximum values of information about any one stimulus are shown in Fig. 5b and in Table 2. The mean of the maximum values of $I(s, R)$ was 2.5 bits. For comparison, the average information $I(S, R)$ conveyed by this particular set of neurons about the set of 65 stimuli was 0.62 bits.

To relate these information measures further to the sparseness of the representation provided by these neurons, histograms showing the sparseness values for each of the neurons are shown in Fig. 6. In Fig. 6a, the sparseness a , and in Fig. 6b, the response sparseness a_r , is shown calculated across the face stimuli in the set. In Fig. 6c, the sparseness a , and in Fig. 6d, the response sparseness a_r , is shown calculated across 25 nonface stimuli in the set. (For the same set of neurons, across the 65 stimuli in the set, the mean sparseness was 0.68, and the mean response sparseness was 0.30). It is clear, in line with the data shown in Figs. 1 to 4, that the sparseness values for these neurons are quite high,

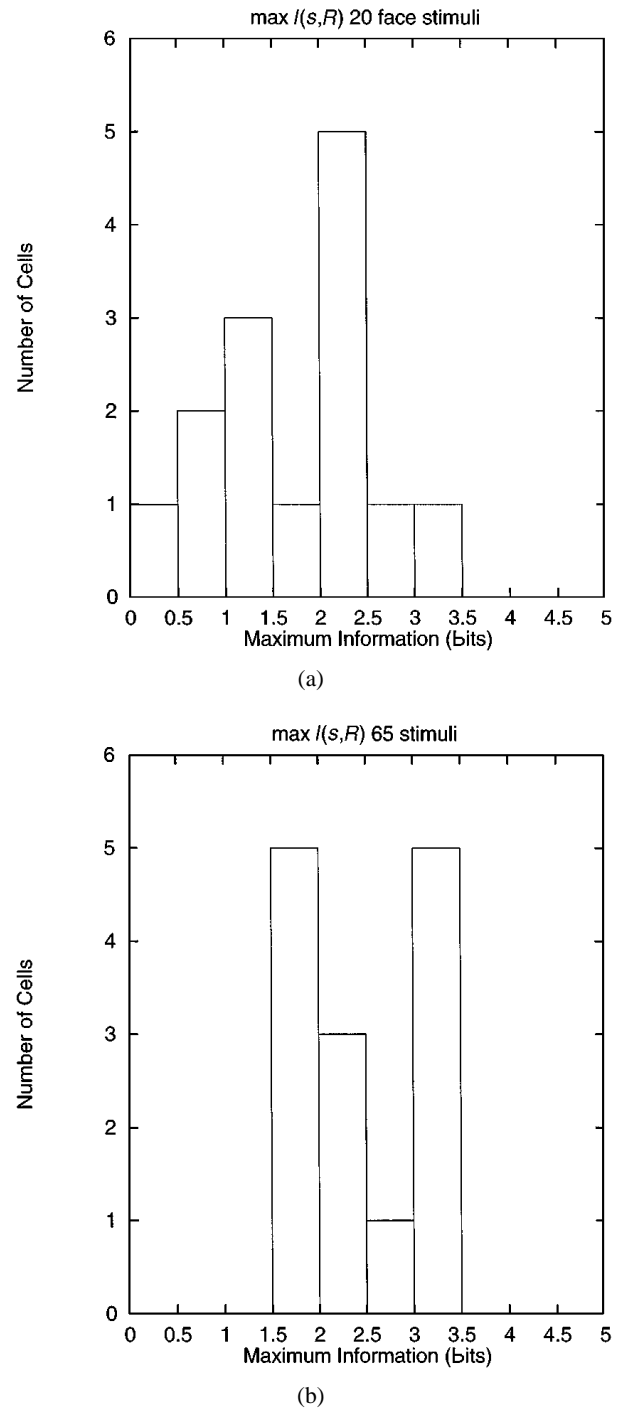


Figure 5. a. The distribution of the maximum values of information about any one stimulus available in the responses of the 14 neurons, for the 20 face stimuli. b. The distribution of the maximum values of information about any one stimulus available in the responses of the 14 neurons, for the 65 face and nonface stimuli.)

reflecting relatively distributed encoding. If response sparseness values are considered, they are seen to be lower, indicating, again in line with Figs. 1 to 4, that relative to the spontaneous firing rates, the representation,

Table 1. For each cell analyzed, the average information $I(S, R)$, the average firing rate (spikes/sec), the maximum information to any one stimulus $I_{\max}(s, R)$ and the associated firing rate r_s , the sparseness a of the representation of the stimuli by the neuron, and the response sparseness a_r (see text). The stimulus set was 20 face stimuli. The standard deviations of the sparseness measures were calculated by propagating the standard error of the mean firing rates with the normal procedure.

Cell number	$I(S, R)$	r_{ave}	$I_{\max}(s, R)$	r_s	a (\pm s.d.)	a_r (\pm s.d.)
243	0.601	45	2.36	135	0.67 ± 0.02	0.55 ± 0.02
242	0.552	54	2.16	106	0.86 ± 0.01	0.66 ± 0.03
241	0.343	88	1.33	46	0.92 ± 0.01	0.89 ± 0.01
240	0.244	52	3.06	11	0.95 ± 0.01	0.92 ± 0.01
236	0.399	40	2.48	116	0.73 ± 0.02	0.19 ± 0.04
235	0.245	20	1.82	41	0.90 ± 0.01	0.66 ± 0.04
234	0.801	72	2.41	21	0.86 ± 0.01	0.59 ± 0.04
233	0.149	47	1.28	82	0.94 ± 0.02	0.56 ± 0.09
232	0.089	33	0.35	47	0.97 ± 0.01	0.76 ± 0.05
231	0.295	15	1.22	1	0.71 ± 0.02	0.52 ± 0.04
229	0.608	70	2.25	13	0.86 ± 0.01	0.75 ± 0.02
227	0.180	52	0.98	23	0.92 ± 0.02	0.57 ± 0.07
222	0.318	30	0.87	49	0.81 ± 0.03	0.15 ± 0.07
164	0.242	29	2.57	56	0.86 ± 0.03	0.70 ± 0.04

Table 2. For each cell analyzed, the average information, the average firing rate, the maximum information to any one stimulus and the associated firing rate, the sparseness of the representation of the stimuli by the neuron, and the response sparseness. The stimulus set was 65 face and nonface stimuli (23 faces and 42 nonfaces).

Cell number	$I(S, R)$	r_{ave}	$I_{\max}(s, R)$	r_s	a (\pm s.d.)	a_r (\pm s.d.)
243	0.607	30	3.1	135	0.58 ± 0.03	0.39 ± 0.03
242	0.835	31	3.2	106	0.70 ± 0.01	0.25 ± 0.02
241	1.019	47	2.3	8	0.64 ± 0.01	0.45 ± 0.01
240	0.926	27	1.9	5	0.62 ± 0.01	0.35 ± 0.02
236	0.515	26	3.4	116	0.81 ± 0.02	0.08 ± 0.04
235	0.343	20	3.2	57	0.84 ± 0.01	0.53 ± 0.03
234	0.748	65	2.1	143	0.79 ± 0.01	0.39 ± 0.03
233	0.588	27	2.6	82	0.73 ± 0.02	0.19 ± 0.04
232	0.406	21	1.7	6	0.75 ± 0.02	0.25 ± 0.05
231	0.429	8	1.9	36	0.53 ± 0.03	0.22 ± 0.05
229	0.791	38	1.8	118	0.60 ± 0.02	0.30 ± 0.03
227	0.407	50	2.4	115	0.80 ± 0.02	0.34 ± 0.04
222	0.780	12	1.8	49	0.39 ± 0.02	0.05 ± 0.03
164	0.318	18	3.1	56	0.73 ± 0.03	0.33 ± 0.04

although distributed, is much sparser than fully distributed.

Finally, the analysis of the information available in the 20 ms window 100 to 120 ms post stimulus onset

was repeated for all 14 cells. The resulting information values, averaged across stimuli, are shown in Fig. 7, and compared with the average information available in 500 ms, as in the analyzes above, and with the

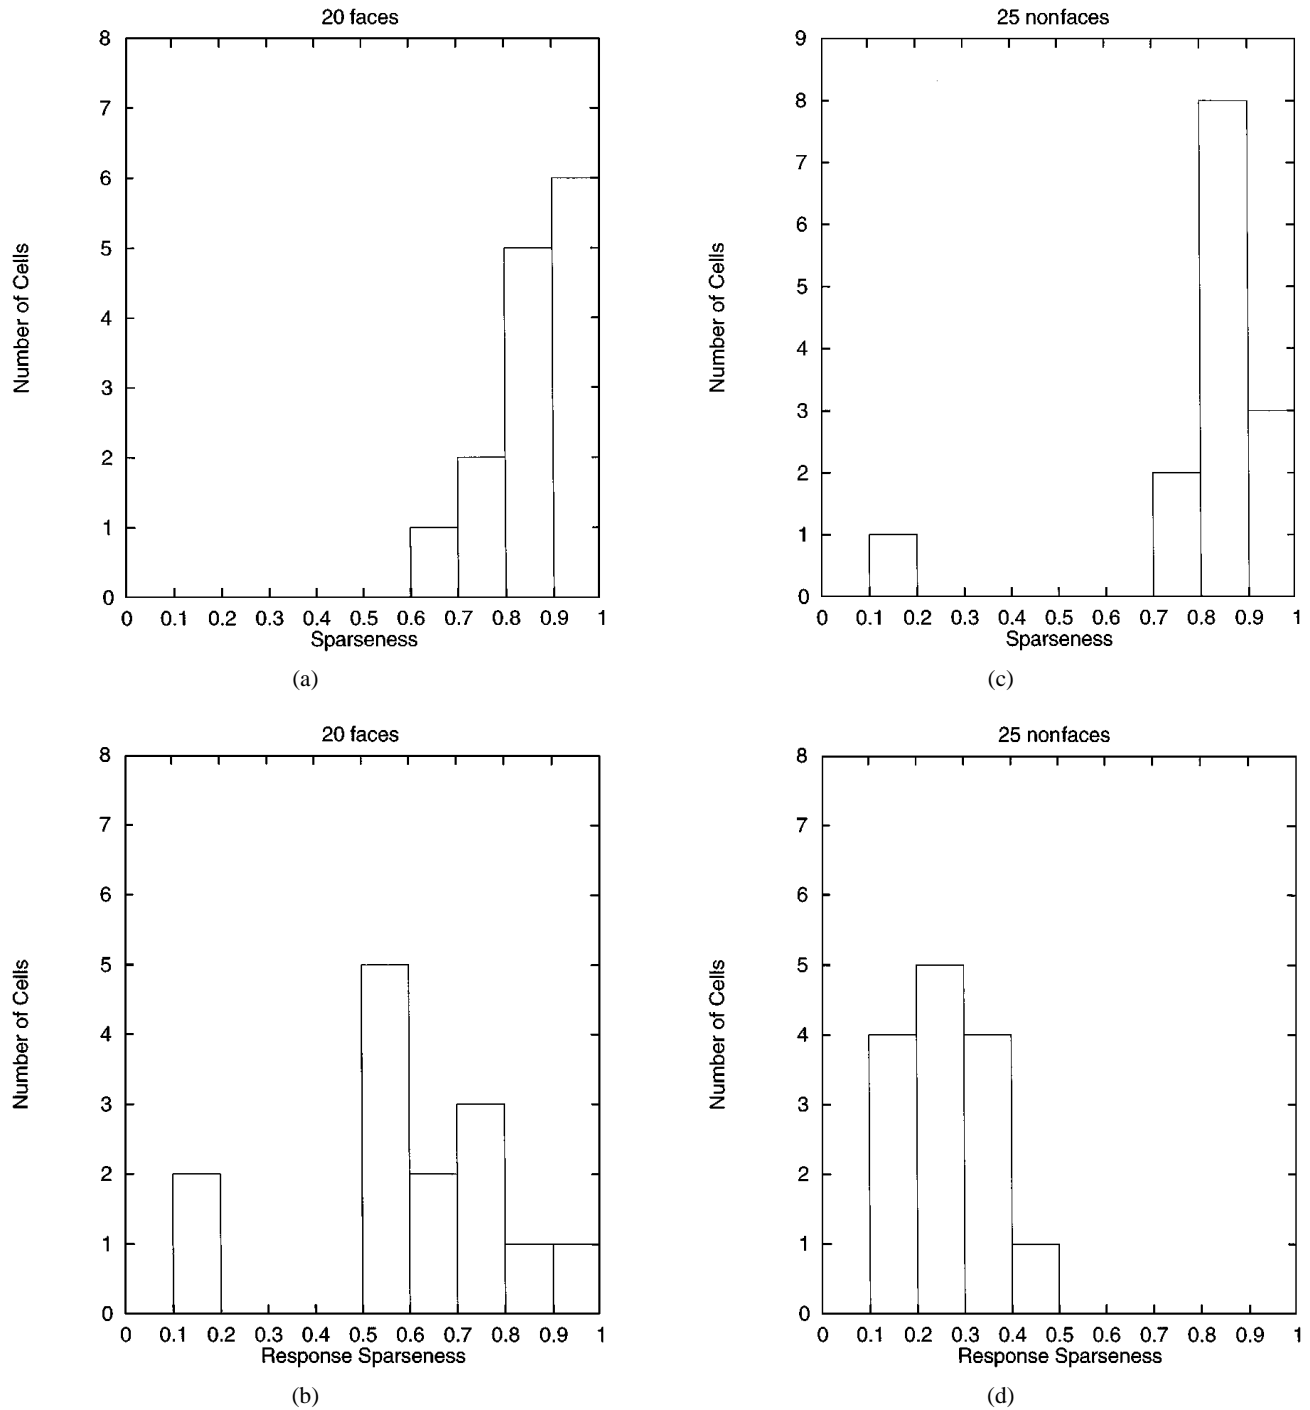


Figure 6. a. The distribution of sparseness values a for the population of cells, calculated across the face stimuli in the set. The mean sparseness was 0.86 ± 0.09 (\pm s.d.). b. As Fig. 6a, but response sparseness values, a_r , are shown. The mean response sparseness was 0.61 ± 0.22 (\pm s.d.). c. The distribution of sparseness values a for the population of cells, calculated across 25 nonface stimuli in the set. The mean sparseness was 0.79 ± 0.11 (\pm s.d.). d. As Fig. 6c, but response sparseness values a_r are shown. The mean response sparseness was 0.22 ± 0.24 (\pm s.d.).

information that would be provided in 20 ms by cells with Poissonian firing and by cells with periodic firing, and with the same distribution of mean rates to each stimulus as the real cells, as explained for cell am229. One can see from Fig. 7 that the values obtained for

the 20 ms period (indicated by the full line) are almost always, except for three cells, in between the values corresponding to the two theoretical cases of Poissonian firing (indicated with dashed impulses) and periodic firing (denoted with dashed boxes). For 12 out of


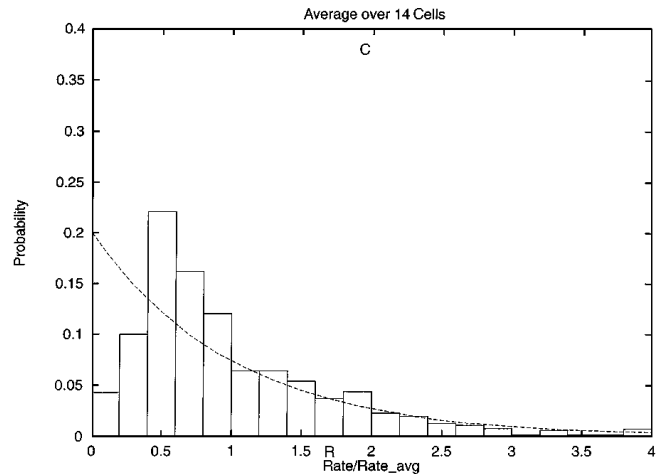


Figure 7. The average information $I(S, R)$ about the 20 face stimuli contained in the firing rates of each of the 14 neurons when rates are measured over 20 ms (solid line), compared to the corresponding values when rates are measured over the standard 500 ms period (diamonds). The values for 20 ms are contrasted with the information that would be provided by theoretical cells with Poisson firing (dashed impulses) and by theoretical cells with periodic firing (dashed histogram).

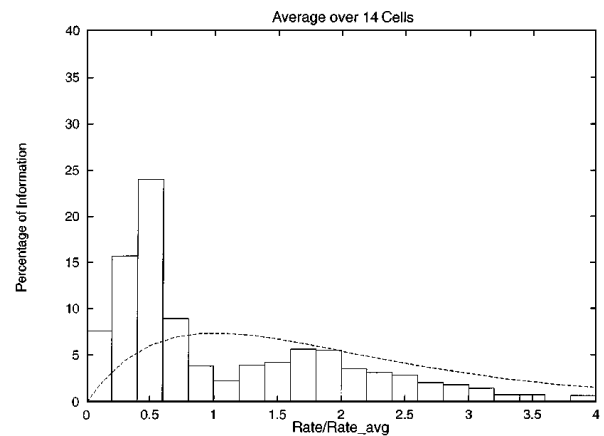
14 cells the values for 20 ms are below, but represent a large fraction of, the values for 500 ms (indicated with diamonds). For two cells the information from the rates measured over 500 ms is actually below that from the 20 ms rates; obviously, in those cases some information in the longer period must be present in a form different from the firing rate and could be extracted by performing, for example, principal component analysis of the firing records (Optican and Richmond, 1987). In fact, for cell am233 (labeled cell 8 in Fig. 7), the mutual information calculated over 500 ms rises from the 0.15 bits carried by the firing rate alone to 0.26 bits when considering the first 3 principal components of the spike train; and for cell am227 (cell 12 in Fig. 7) from 0.18 bits to 0.33 bits, respectively. For the other cells, the first three principal components bring about only minor increases over the information contained in the firing rate, as reported previously (Tovee and Rolls, 1995).

It remains interesting that the firing rate alone appears to convey in a period as short as 20 ms a large fraction of the information it conveys in 500 ms. The grand averages across all 14 cells are 0.362 bits for the information in 500 ms and 0.193 bits (53%) for the information in 20 ms, while the theoretical values are 0.124 bits for Poissonian cells and 0.560 bits for periodic cells.

The probability distribution of firing rates for one neuron to the whole set of 65 stimuli was shown in Fig. 2f. We calculated the average across the whole



(a)



(b)

Figure 8. a. The average across the whole population of 14 cells of the probability distribution of different firing rates. In order to enable the cells to be combined, the firing rate of each cell was normalized to an average rate of 1. An exponential distribution fitted to the mean rate is shown by the line. b. The percentage of the information carried by the different levels of firing rate for this population of 14 cells. The average rate for each cell was normalized to 1 so that the cells could be combined. The firing rate measure on the abscissa is the firing rate expressed as a fraction of the mean rate for the cell (which was set to 1). The line shows what would be expected if the cells had an exponential distribution of their firing rates, and encoded information noiselessly.

population of 14 cells of the probability distribution of different firing rates. In order to enable the cells to be combined, we normalized the firing rate of each cell to a mean value of 1. The resulting probability distribution (calculated over all individual trials to all 65 stimuli for all 14 cells) is shown in Fig. 8a. We also, as in Fig. 2f, show an exponential distribution fitted to this mean rate. While the exponential distribution may provide a reasonable fit at high rates for this set of stimuli and cells, there are too few very low rates to fit

the exponential well at this part of the curve and also too many rates just below the mean rate (see further, Discussion, paragraph 5). (This may be related to the fact that cortical neurons typically have a significant spontaneous firing rate, which averaged 23 spikes/sec for this population of neurons, so that very low rates may be unusual.)

It is possible to show how much of the information is carried by the different levels of firing rate, given the mean firing rates elicited by each stimulus and the corresponding $I(s, R)$ values that have been the subject of this article. The result is shown in Fig. 8b. It is of considerable interest that much of the information was available from the firing rates that were below the mean, related to the fact that low firing rates were very common (see Discussion). The mode of this distribution is at 0.5 with respect to the mean firing rate across all stimuli. We were also able to calculate the percentage of information that would be carried by different firing rates if the firing rates were exponentially distributed and with the additional assumption implicit in that approach (Levy and Baxter, 1996; Baddeley, 1996) of noise-free coding. This comparison is interesting because (1) the exponential distribution approximately fits the tails of the experimental firing rate distributions and (2) the exponential distribution maximizes, in the noiseless coding case, the information transmitted with the constraint of fixed mean firing rate (and thus fixed metabolic cost) (Levy and Baxter, 1996; see also Rieke et al., 1996). It is shown in Fig. 8b that the exponential distribution has most of the information available in the responses centred on the mean response. In contrast, the real responses of these inferior temporal cortex neurons have most of the information about this large set of natural images available in firing rates well below the mean and show in fact a dip in the information available in those rates that are near the mean rate for each cell.

Discussion

The analyzes described here elucidate a quantitative approach to analyzing the representation provided by neurons. First, the average information about the set of stimuli $I(S, R)$ was found to be 0.36 bits on average for the population of neurons. Second, it was shown that the information available in the neuronal responses was not equal to all stimuli but that there was usually much more information about some stimuli than about others. The maximum information provided by the responses of these neurons to one of the stimuli in the

set (of 20 faces) was for the majority of neurons in the range 0.5 to 2.5 bits and averaged 1.8 bits (see Fig. 5a and Table 1). The maximum information provided by the responses of these neurons to one of the stimuli in the set of 65 face and nonface stimuli was in the range 1.5 to 3.5 bits, and averaged 2.5 bits (see Fig. 5b and Table 2). Third, it was shown that the amount of information $I(s, R)$ about a stimulus s is related to how far the response to that stimulus is from the average response to all stimuli (see Figs. 2b, 3b, and 4b). This does not mean that the information measure can just be replaced by the number of spikes per second that the neuronal response to a stimulus is different from the average response to all stimuli, for the information measure takes into account the variability of the response. (This can be made fully evident by noting that if the response were very variable—that is, noisy—to each stimulus, then little information would be available on any one trial about which stimulus was presented, even if the firing rate was very different from the mean firing rate.) Fourth, it was found in particular that information is available if the neuronal response to a stimulus is below the mean response to all stimuli, as made clear in Figs. 2b, 3b, and 4b. Fifth, and in line with the preceding statement about variability in the neuronal response, it was shown in Figs. 2c, 3c, and 4c that the number of standard deviations the response to a stimulus was above or below the mean response was also closely related to the information $I(s, R)$ available about that stimulus. While such a z score takes into account the variability in the response to each stimulus, it does not reflect the distribution of mean firing rates across stimuli, nor, more important, does it explain why the information about those stimuli that elicited mean responses close to the overall mean rate should always be close to zero. In fact, if different stimuli elicited mean responses uniformly distributed in a given range, and with zero or uniform variability across trials, each response would be roughly equiprobable. In this theoretical case, the information conveyed about each stimulus would be roughly equal across stimuli and would not show, as suggested instead by the z score analysis, a lower value at rates close to the overall mean. In another theoretical but slightly more realistic case, different stimuli could be taken to produce responses that are exponentially distributed. If the variability of each response across trials was, again, minimal or small and uniform, the distribution of information carried by different rates would show a maximum, not a minimum, at rates close to the mean rate (see Fig. 8b)! In such a

case, lower rates carry less information (than rates close to the mean) because too many stimuli elicit responses crowded at low rates, while high rates are individually more informative than intermediate rates, but on the whole carry less information just because they occur very infrequently. Our observations, summarized by the C-shaped curve that was similar across all cells and by the cumulative trend reported in Fig. 8b, are clearly not consistent with the assumption of uniform variability or of a noiseless encoding.

One approach to understand the C-shaped curve evident in Figs. 2b, 3b, and 4b is to note that it follows very closely the C-shaped curve that describes the instantaneous information transmission rate, as a function of firing rate, associated with the distribution of responses to the stimuli in the set. It is possible to show that the time derivative of the information in the response to each stimulus (calculated assuming no prior knowledge of the stimulus set) is in general

$$dI(s, R)/dt = r_s \log_2(r_s/r_{\text{mean}}) + (r_{\text{mean}} - r_s) * \log_2(e) \quad (9)$$

bits per second (Panzeri et al., 1996a). This time derivative means calculating the information available in the rates measured over a short period dt , dividing by dt , and then taking dt to zero; when averaged over different stimuli with the corresponding probabilities and divided by the overall mean rate, it yields the mean information carried by each spike, as used, for example, by Skaggs et al. (1993) (see also Bialek, 1991). The curve $dI(s, R)/dt$ is plotted in Fig. 9, and can be seen




Figure 9. The time derivative of the information in the response to each stimulus (information rate in bits/sec) shown in relation to the firing rate of a neuron. The mean firing rate response of the neuron to all the stimuli was set to 50 Hz for this calculated graph (see text).

to describe rather well the data plotted in Figs. 2b, 3b, and 4b, apart from the appropriate rescalings. In particular, the data indicate very small amounts of information about stimuli when the mean rate to that stimulus is close to the overall average response, in agreement with the curve of Fig. 9, which touches the vertical axis when the rate equals the average rate. The instantaneous information rate depends solely on the probability of emitting, in a very small time window, the first spike of a train, and in fact once averaged across stimuli and divided by the overall mean rate r_{mean} it is equal to the mutual information per spike. Its being zero at the mean rate r_{mean} is related to the fact that a binary distribution of rates is optimal for short time windows (Panzeri et al., 1996a). This is because, intuitively, the emission of a spike within the window tends to signal stimuli that elicit high rates, whereas the failure to emit a spike can be associated with stimuli that elicit no firing; and stimuli that elicit intermediate responses are the most difficult to predict from the observed response. The observation, then, is that the information carried by a 500 ms firing rate response about each stimulus is of essentially the same nature, in its dependence on firing rate, as the information carried by rates extracted from much shorter time periods, although it may attain higher values. Although the distribution of firing rates over 500 ms is continuous and its tail can be approximately exponential (see Fig. 8a), the information conveyed by those rates appears to still reflect the binary nature of the spikes, whereby intermediate probabilities of producing a spike provide less information than either high or low probabilities. This is apparent from Fig. 8b, which shows a significant depression right at r_{mean} , where instead the information transmitted by a continuous, exponential distribution of noiseless rates would have its maximum. This point is made in further detail by calculating the information provided by the firing rates measured over a period as short as 20 ms. As expected (since 20 ms is much shorter than 500 ms), the information about each stimulus is related to the mean rate to that stimulus through, again, the same C-shaped curve that corresponds to the time derivative of the information. One may wonder, then, to what extent measuring firing rates over longer periods (intuitively, with higher precision) can result in higher information values. Data on this issue is provided, in Fig. 7, only for averages across stimuli because information values for individual stimuli are too noisy for a detailed comparison. Averaging further across cells, the firing rate measured over 20 ms taken at the peak

of the response conveys over 50% of the information it conveys over 500 ms. Therefore, waiting 25 times longer produces only a 100% gain in information. To convey this relatively large amount of information in a short period, the variability in the firing of each cell must be somewhat constrained. This point is explored here by calculating the information conveyed, in the same manner, by theoretical cells that fired with high variability (the Poisson cells) or with no variability (the periodic cells) and used the same distribution of mean firing rates to represent the set of stimuli as the actual cells. It is shown in Fig. 7 that different (actual) cells convey information at different intermediate levels between the two theoretical sample situations. On average, across stimuli and cells, Poisson firing would have resulted in about half the information (63%), and periodic firing in nearly three times the information (288%). Analyzing individual cells in more detail, we noted that the majority of them, 11 out of 14, provided in 20 ms more information than they would have with Poisson firing *and* tended to display significantly less variability than Poisson. Three cells had similar variability and provided similar amount of information to the theoretical Poisson mode. Finally, the one cell that in 20 ms conveyed with its actual firing significantly less information than its theoretical Poisson counterpart, was also the one cell (am164) that happened to fire *more* irregularly than Poisson.

The actual results obtained provide quantitative information theoretic evidence for the first time on the issue of how specifically neurons are tuned to individual stimuli in the higher parts of the visual system. If a neuron were to respond to one of the faces in the set of 20 faces and to none of the other faces, and to do this consistently without noise, then the amount of information gained when the neuron fired would be 4.32 bits, (that is, $-\log_2(1/20)$ bits). In another possible type of encoding, the neuron might respond at the same rate to half the faces in the set and not to the other half of the faces in the set—that is, might show fully distributed encoding. In this case, the information available when any one of the faces in the set was shown would be 1 bit (that is, $-\log_2(1/2)$ bits). In fact, the encoding was intermediate between these fully localist and fully distributed forms, with individual neurons providing reasonably high information about a number of stimuli, but in a partly distributed way; moreover, the encoding was different from both these simple examples in that it was based on a graded and noisy response. Within the localist-distributed scale, the emphasis was toward

distributed rather than localist, as shown also by the sparseness measures illustrated in Fig. 6.

It is interesting to relate the information about each stimulus $I(s, R)$ to the average information provided on any one trial about the stimulus set $I(S, R)$. If for the 20 equiprobable face stimuli the neurons were localist, then the information when the effective stimulus was shown would be $-\log_2(1/20) = 4.32$ bits, and when one of the other stimuli was shown would be $-\log_2(19/20) = 0.074$ bits. On average, the information on any one trial would be (given the equiprobable stimuli) $(1 \times 4.32 \text{ bits} + 19 \times 0.074 \text{ bits})/20 = 0.29$ bits. If the neuron showed fully distributed encoding, then the average information on any one trial would be 1 bit. (These calculations again for illustration assume no noise, and binary firing rates.) The actual state found was between these, with $I(S, R)$ (for 20 faces) having an average of 0.36 bits.

The analysis shown in Fig. 8a indicates that the firing rates of this population of inferior temporal cortex cells approximately fit an exponential distribution (as suggested by Levy and Baxter, 1996, and Baddeley, 1996) at high rates but that there are fewer very low rates and more moderate rates (that is, close to 0.5 of the mean) than would be predicted by an exponential distribution. Similar, at least partly exponential distributions of firing rate have also been found from a similar group of neurons when the monkey looks at a continuous video film of his normal environment (Baddeley et al., 1998). The ways in which these firing-rate distributions could be produced are the subject of a simple model using Gaussian input distributions of the activation received by a neuron and a threshold linear activation function (Panzeri et al., 1996b). The analysis shown in Fig. 8b shows that with the calculation of stimulus-specific information introduced in this article, one can actually quantify the distribution of the fraction of information carried by different firing rates taking into account also the real variability of responses (which is naturally included in the definition of mutual or stimulus-specific information). The interesting result is that most of the information about the large set of natural images is available for this set of neurons in firing rates that are well below the mean. For comparison, an exponential distribution of firing rates would have most of the information available from the firing rates centred close to the mean (see Fig. 8b). The fact that the exponential distribution would be only optimal for encoding information when there is a constraint on maintaining a low mean firing rate for metabolic efficiency (Levy and

Baxter, 1996) only applies with noise-free encoding is considered elsewhere (Treves et al., in preparation). The evidence that firing rate distributions, at least for the population of neurons described here, are not exponentially distributed, and the fact that we were able to take into account with the method described here the actual variability of the firing rates to stimuli are probably both factors that contribute to the actual distribution of information available from the neuronal responses as a function of firing rate (Fig. 8b). This actual distribution shows the interesting results that taking into account the actual firing-rate distribution, and the trial-to-trial variability in the neuronal response, the mode of the distribution of the information available from different rates is well below the mean response, centered close to half the mean response for this population of neurons. Further, at the mean response itself the distribution shows, instead, a depression, consistent with the behavior of the instantaneous information rate that is strictly zero at that point. This suggests that even when considering relatively long time windows, the distributions that optimize *instantaneous* information transmission, which happen to be binary distributions, would be more efficient than the exponential distribution that, according to Shannon's noiseless coding theorem (Levy and Baxter, 1996), would optimize information transmission in the absence of noise (both subject to the constraint of fixed mean rate). The fact that the observed firing-rate distributions at long sampling intervals are continuous and are not binary suggests that optimal information transmission does not explain the firing rate distributions. Instead, we suggest the hypothesis that, for large sets of ecological stimuli, the firing-rate distributions (which can be close to but do not usually perfectly fit an exponential distribution; see Fig. 8a) reflect the firing that would be produced by Gaussian inputs to a neuron that has a threshold linear activation function, in the presence of Gaussian thermal noise (Panzeri et al., 1996b; see further Treves et al., in preparation).

It has been argued elsewhere (Rolls and Tovee, 1995b) that the rather distributed encoding found in this population of neurons allows a relatively large amount of information about a set of stimuli to be provided by such a population—provided, of course, that they do not have the same profile of responsiveness to the set of stimuli. This can be realized considering again the simple case of binary noiseless encoding; in that case, the ideal representation for *discrimination* would be when each cell fires to 50% of the stimuli (a different 50%

for each cell), for the maximum information suitable for comparing fine differences between different stimuli would then be made available across the population. However, a representation as distributed as this would not be appropriate for a memory system, in which the aim is to store a large number of memories. In an associative memory with neurons with continuously variable firing rates, such as the autoassociative memory believed to be implemented in the hippocampus (Rolls, 1989), we have shown (Treves, 1990; Treves and Rolls, 1991, 1994) that the maximum number p_{\max} of firing patterns that can be (individually) retrieved is proportional to the number C of (associatively) modifiable synapses per cell, by a factor that increases roughly with the inverse of the sparseness a of the neuronal representation. Approximately,

$$p_{\max} \simeq \frac{C}{a \ln(1/a)} k,$$

where k is a factor that depends weakly on the detailed structure of the rate distribution, on the connectivity pattern, and so forth. Thus, in a memory system, one parameter limiting the number of memories that can be stored and retrieved correctly is the number of modifiable synapses received by each neuron (in the order of 5,000 to 15,000), and another is the sparseness of the representation a . (Similar considerations apply to pattern associators (Rolls and Treves, 1990) such as are suggested to be implemented in the amygdala and orbitofrontal cortex, and in cortico-cortical backprojections (Rolls, 1989, 1992a, 1992c; Treves and Rolls, 1994)). Thus, in memory systems optimal performance requires a sparse representation, and this is consistent with values for sparseness in the hippocampus that are in the order of 0.01 to 0.04 (Leonard and McNaughton, 1990; Barnes et al., 1990; Jung and McNaughton, 1993; Rolls and O'Mara, 1993; O'Mara et al., 1994; Rolls et al., 1989; Cahusac et al., 1989; see Treves and Rolls, 1994). It is therefore proposed that these fundamentally different constraints account for the different sparsenesses of representations found in the high-order sensory cortices, such as the temporal cortical areas described here, and in memory systems such as the hippocampus. In the sensory cortex, a relatively distributed representation may be used in order to optimize discriminative ability. In memory systems, much more sparse representations may be used in order to maximize the number of memories that can be stored. (We note that many of the cells of the type analyzed here have other properties that make them

suitable for discrimination between faces, including invariance with respect to size, spatial frequency, translation, and even in some cases view (see Rolls, 1992a, 1994; Rolls and Tovee, 1995a.)

This discussion on the sparseness of representation provided by these face-selective cells runs counter to the possibility that they are very specifically tuned and provide a cardinal or “grandmother cell” type of very sparse representation (Barlow, 1972). Instead, the data presented in this article indicate that they are very selective, in that they respond rather selectively to stimuli within the class faces and respond less to stimuli that are not faces. However, within the class for which they are selective in their responses (faces), the representation is very distributed, implying great discriminative capacity, including the representation of small differences between faces presented simultaneously. It will be of interest in future studies to determine the sparseness of the representation of nonface stimuli provided by other neurons in the temporal cortical visual areas.

Acknowledgments

This research was supported by Medical Research Council Grant PG8513790 to Professor E.T. Rolls, by the Human Frontier Science Program, and by the European Community Human Capital and Mobility program. We are very grateful for helpful comments and assistance provided by Professor L.F. Abbott, W. Bialek, Dr. R. Baddeley, and D. Foster.

Note

1. Logs to the base 2 are used throughout this article.

References

- Baddeley R (1996) An efficient code in V1? *Nature* 381:560–561.
- Baddeley, R.J., Abbot, L.F., Booth, M.J.A., Sengpiel, F., Freeman, T., Wakeman, E.A., and Rolls, E.T. (1998) Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society B*, in press.
- Baizer JS, Ungerleider LG, Desimone R (1991) Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *J. Neurosci.* 11:168–190.
- Barlow HB (1972) Single units and sensation: A neuron doctrine for perceptual psychology? *Perception* 1:371–394.
- Barnes CA, McNaughton BL, Mizumori SJ, Lim LH (1990) Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Prog. Brain Res.* 83:287–300.
- Baylis GC, Rolls ET, Leonard CM (1985) Selectivity between faces in the responses of a population of neurons in the cortex in the superior temporal sulcus of the monkey. *Brain Res.* 342:91–102.
- Baylis GC, Rolls ET, Leonard CM (1987) Functional subdivisions of temporal lobe neocortex. *J. Neurosci.* 7:330–342.
- Bialek W (1991) Optimal signal processing in the nervous system. In: W Bialek, ed. *Princeton Lectures on Biophysics*. World Scientific, London.
- Bruce C, Desimone R, Gross CG (1981) Visual properties of neurons in a polysensory area in superior temporal sulcus of the macaque. *J. Neurophysiol* 46:369–384.
- Churchland PS, Sejnowski TJ (1992) *The Computational Brain*. MIT Press, Cambridge, MA.
- Desimone R (1991) Face-selective cells in the temporal cortex of monkeys. *J. Cog. Neurosci.* 3:1–8.
- Desimone R, Gross CG (1979) Visual areas in the temporal lobe of the macaque. *Brain Res.* 178:363–380.
- Desimone R, Albright TD, Gross CG, Bruce C (1984) Stimulus-selective properties of inferior temporal neurons in the macaque. *J. Neurosci.* 4:2051–2062.
- Fukunaga K (1972) *Introduction to Statistical Pattern Recognition*. Academic Press, New York.
- Golomb D, Hertz J, Panzeri S, Richmond BH, Treves A (1997) How well can we estimate the information carried in neuronal responses from limited samples? *Neural Computation*, 9:649–665.
- Gross CG (1992) Representation of stimuli in inferior temporal cortex. *Phil. Trans. Roy. Soc. Lond. B* 335:3–10.
- Gross CG, Desimone R, Albright TD, Schwartz EL (1985) Inferior temporal cortex and pattern recognition. *Exp. Brain Res. Suppl.* 11:179–201.
- Hinton GE, McClelland JL, Rumelhart DE (1986) Distributed representations. In: DE Rumelhart, JL McClelland, eds. *Parallel Distributed Processing*, MIT Press, Cambridge, MA. pp. 77–109.
- Hornak J, Rolls ET, Wade D (1996) Face and voice expression identification and their association with emotional and behavioural changes in patients with ventral frontal lobe damage. *Neuropsychologia* 34:247–261.
- Jung MW, McNaughton BL (1993) Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus* 3:165–182.
- Leonard BW, McNaughton BL (1990) Spatial representation in the rat: Conceptual, behavioral and neurophysiological perspectives. In: RP Kesner, DS Olton, eds. *Neurobiology of Comparative Cognition*. Erlbaum, Hillsdale, NJ.
- Levy WB, Baxter RA (1996) Energy efficient neural codes. *Neural Computation* 8:531–543.
- Macrae AW (1971) On calculating unbiased information measures. *Psychol. Bull.* 75:270–277.
- Maunsell JHR, Newsome WT (1987) Visual processing in monkey extrastriate cortex. *Ann. Rev. Neurosci.* 10:363–401.
- Optican L, Richmond BJ (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J. Neurophysiol.* 57:132–146.
- Optican LM, Gawne TJ, Richmond BJ, Joseph PJ (1991) Unbiased measures of transmitted information and channel capacity from multivariate neuronal data. *Biol. Cyber.* 65:305–310.
- Panzeri S, Treves A (1996) Analytical estimates of limited sampling biases in different information measures. *Network* 7:87–107.

- Panzeri S, Biella G, Rolls ET, Skaggs WE, Treves A (1996a) Speed, noise, information and the graded nature of neuronal response. *Network* 7:365–370.
- Panzeri S, Booth M, Wakeman EA, Rolls ET, Treves A (1996b) Do firing rate distributions reflect anything beyond just chance? *Society for Neuroscience Abstracts* 22:445.5.
- Perrett DI, Rolls ET, Caan W (1982) Visual neurons responsive to faces in the monkey temporal cortex. *Exp. Brain Res.* 47:329–342.
- Richmond BJ, Optican L (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. II. Quantification of response waveform. *J. Neurophysiol.* 57:147–161.
- 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.
- Rieke F, Warland D, De Ruyter Van Steveninck RR, Bialek W (1996) *Spikes: Exploring the Neural Code*. MIT Press, Cambridge, MA.
- Rolls ET (1981a) Processing beyond the inferior temporal visual cortex related to feeding, learning, and striatal function. In: Y Katsuki et al., eds. *Brain Mechanisms of Sensation*. Wiley, New York, pp. 241–269.
- Rolls ET (1981b) Responses of amygdaloid neurons in the primate. In: Y Ben-Ari, eds. *The Amygdaloid Complex*. Elsevier, Amsterdam, pp. 383–393.
- Rolls ET (1984) Neurons in the cortex of the temporal lobe and in the amygdala of the monkey with responses selective for faces. *Human Neurobiol.* 3:209–222.
- Rolls ET (1989) Functions of neuronal networks in the hippocampus and neocortex in memory. In: JH Byrne, WO Berry, eds. *Neural Models of Plasticity: Experimental and Theoretical Approaches*. Academic Press, San Diego, pp. 240–265.
- Rolls ET (1991) Neural organisation of higher visual functions. *Curr. Opinion in Neurobiol.* 1:274–278.
- Rolls ET (1992a) Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. *Phil. Trans. Roy. Soc.* 335:11–21.
- Rolls ET (1992b) The processing of face information in the primate temporal lobe. In: V Bruce, M Burton, eds. *Processing Images of Faces*. Ablex, Norwood NJ, pp. 41–68.
- Rolls ET (1992c) Neurophysiology and functions of the primate amygdala. In: JP Aggleton, ed. *The Amygdala*, Wiley-Liss, New York, pp. 143–165.
- Rolls ET (1994) Brain mechanisms for invariant visual recognition and learning. *Behavioural Processes* 33:113–138.
- Rolls ET (1995) Learning mechanisms in the temporal lobe visual cortex. *Behav. Brain Res.* 66:177–185.
- Rolls ET, Baylis GC (1986) Size and contrast have only small effects on the responses to faces of neurons in the cortex of the superior temporal sulcus of the monkey. *Exp. Brain Res.* 65:38–48.
- Rolls ET, Miyashita Y, Cahusac PMB, Kesner RP, Niki H, Feigenbaum J, Bach L (1989) Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *J. Neurosci.* 9:1835–1845.
- Rolls ET, O'Mara S (1993) Neurophysiological and theoretical analysis of how the hippocampus functions in memory. In: T Ono, LR Squire, ME Raichle, DI Perrett, M Fukuda, eds. *Brain Mechanisms of Perception and Memory: From Neuron to Behavior*, Oxford University Press, New York, pp. 276–300.
- Rolls ET, Treves A (1990) The relative advantages of sparse versus distributed encoding for associative neuronal networks in the brain. *Network* 1:407–421.
- Rolls ET, Tovee MJ (1995a) The responses of single neurons in the temporal visual cortical areas of the macaque when more than one stimulus is present in the visual field. *Exp. Brain Res.* 103:409–420.
- Rolls ET, Tovee MJ (1995b) The sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J. Neurophysiol.* 73:713–726.
- Rolls, E.T., Treves, A. (1998) *Neural Networks and Brain Function*. Oxford University Press, Oxford.
- Seltzer B, Pandya DN (1978) Afferent cortical connections and architectonics of the superior temporal sulcus and surrounding cortex in the rhesus monkey. *Brain Res.* 149:1–24.
- Skaggs WE, McNaughton BL, Gothard KM, Markus EJ (1993) An information-theoretic approach to deciphering the hippocampal code. In: SJ Hanson, JD Cowan, CL Giles, eds. *Advances in Neural Information Processing Systems 5*. Morgan Kaufman, San Mateo, pp. 1030–1037.
- Smith DV, Travers JB (1979) A metric for the breadth of tuning of gustatory neurons. *Chem. Senses* 4:215–229.
- Suga N (1989) Principles of auditory information-processing derived from neuroethology. *J. Exp. Biol.* 146:277–286.
- Tanaka K, Saito H-A, Fukada Y, Moriya M (1991) Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *J. Neurophysiol.* 66:170–189.
- 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.
- Tovee MJ, Rolls ET, Azzopardi P (1994) Translation invariance and the responses of neurons in the temporal visual cortical areas of primates. *J. Neurophysiol.* 72:1049–1060.
- Tovee MJ, Rolls ET (1995) Information encoding in short firing rate epochs by single neurons in the primate temporal visual cortex. *Visual Cognition* 2:35–58.
- Treves A (1990) Graded-response neurons and information encodings in autoassociative memories. *Physical Review A* 42:2418–2430.
- Treves A, Panzeri S (1995) The upward bias in measures of information derived from limited data samples. *Neural Computation* 7:399–407.
- Treves A, Rolls ET (1991) What determines the capacity of autoassociative memories in the brain? *Network* 2:371–397.
- Treves A, Rolls ET (1994) A computational analysis of the role of the hippocampus in memory. *Hippocampus* 4:374–391.
- Weiskrantz L, Saunders RC (1984) Impairments of visual object transforms in monkeys. *Brain* 107:1033–1072.