JOURNAL OF NEUROPHYSIOLOGY
Vol. 75, No. 5, May 1996. *Printed in U.S.A.*

# Representation of Olfactory Information in the Primate Orbitofrontal Cortex

EDMUND T. ROLLS, HUGO D. CRITCHLEY, AND ALESSANDRO TREVES
*University of Oxford, Department of Experimental Psychology, Oxford OX1 3UD, United Kingdom*

## SUMMARY AND CONCLUSIONS

*1*. To analyze the information represented about individual odor stimuli in the responses of single olfactory neurons in the primate orbitofrontal area, neuronal responses were measured to a set of seven to nine odorants in macaques performing an olfactory discrimination task. The population of neurons analyzed had responses that were significantly differential to the odorants.

*2*. Information theoretic analyses were applied to the responses of the neurons, and information measures were calculated from the firing rate of the responses and from the principal components of the responses. The information reflected by the firing rate of the response accounted for the majority of the information present (86%) when compared with the information derived from the first three principal components of the spike train. This indicated that temporal encoding had a very minor role in the encoding of olfactory information by single orbitofrontal olfactory cells.

*3*. The average information about which odorant was presented, averaged across the 38 neurons, was 0.09 bits, a figure that is low when compared with the information values previously published for the responses of temporal lobe face-selective neurons.

*4*. Application of information theoretic analyses to the responses of these neurons showed how much information about which stimulus was delivered was present in the responses of individual neurons. It was found that for the majority of the neurons significant amounts of information were reflected about one or two of the odorants presented.

*5*. For each neuron, the information reflected in the responses of that neuron about the reinforcement value and the information about the identity of the odorants were calculated. It is shown that many neurons carry information about which of the odorants was presented; in addition, some neurons reflect information only about the taste association of the stimuli and not about odorant identity.

*6*. Measurements of the sparseness of the representation indicated that a broadly distributed representation of the identity of odorants was present in this population of neurons.

## INTRODUCTION

The primate orbitofrontal cortex forms the ventral aspect of the frontal lobe. The primary olfactory (pyriform) cortex projects into area 13a in the caudal orbitofrontal cortex, from which there are onward projections to an extensive part of the orbitofrontal cortex (Barbas 1993; Morecraft et al. 1992; Price et al. 1991). The orbitofrontal cortex also contains, more laterally, the secondary taste cortex, which also has extensive onward connections to other parts of the orbitofrontal cortex (Baylis et al. 1994; Rolls 1995; Rolls et al. 1990).

The orbitofrontal cortex in primates is known to be important for the identification and discrimination of odors. In humans it has been shown that pathology involving the right orbitofrontal cortex affects the ability of subjects to identify odors (Jones-Gotman and Zatorre 1988; Zatorre and Jones-Gotman 1991). This finding has been supported by positron emission tomography evidence of activation of the right orbitofrontal cortex during the inhalation of odors in normal subjects (Zatorre et al. 1992). In monkeys, lesions to the lateral orbitofrontal cortex impair the ability to recognize food from nonfood by associated odors (Tanabe et al. 1975). Consistent with this, single-neuron recordings from the orbitofrontal cortex showed the presence of olfactory neurons, at least some of which were selective in their responses to the odorants used (Tanabe et al. 1974, 1975a,b; Yarita et al. 1980). Rolls and Baylis (1994) confirmed the presence of odor-responsive neurons in the primate orbitofrontal cortex, and went on to show that for some orbitofrontal neurons information from the taste and olfactory modalities converges onto single neurons. Such neurons may exhibit cross-modal correspondence of responses in gustatory and olfactory modalities for the representation of foods, and thereby show selectivity to particular food odors.

In more recent work it has been shown that the responses of at least some olfactory neurons in the orbitofrontal cortex are influenced by the reward value and taste association of the olfactory stimuli (Critchley and Rolls 1996a,b; Rolls et al. 1996). In particular, it has been shown first that some of these olfactory neurons that respond to food-related odors have this responsiveness modulated by hunger, and respond much less to that odor when the monkey is satiated with that food (Critchley and Rolls 1996b). Second, it has been shown that in an automated olfactory discrimination task, the responses of 35% of the neurons reflected whether the odor was associated with a food taste (sucrose) or with the taste of saline (Critchley and Rolls 1996a). The reinforcement association of the stimuli was also shown to influence the representation of the odorants in multidimensional analysis of the neuronal responses. Third, it was shown in the reversal of an olfactory discrimination task that the responses of some orbitofrontal olfactory neurons (68%) depend on whether the odor was currently associated with the taste of sucrose or with saline (Rolls et al. 1996).

In recent years there has been much interest in the olfactory system from computational neuroscientists (e.g., Davis and Eichenbaum 1991). The neuroanatomy of the olfactory bulb and the pyriform cortex have provided computational modelers with specific examples of parallel distributed processing (Kauer et al. 1991) and architecture of a type well suited for associative memory (Haberly and Bower 1989) and hierarchical clustering (Ambros-Ingerson et al. 1990;

Granger et al. 1989). A question that arises and that is important in understanding the operation of the olfactory system with its networks of neurons is how the information about odorants is represented by the activity of cortical olfactory neurons. The investigation described here takes a quantitative information theoretic approach to this issue.

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(base 2) of the probability of an event occurring. [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 4 bits of information are provided by the occurrence of that event. If a cell responded at every presentation of 1 stimulus of a set of 8 stimuli, and not at all to the other 7 stimuli, then 3 bits of information are reflected about the stimulus set when the cell fires.] In the present study information analysis was applied to the responses of individual olfactory neurons to the odorants in the olfactory discrimination task. The information conveyed about individual odorants, and about the stimulus set, was calculated to provide a metric for the responses of olfactory neurons. The information theoretic analyses 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 are described in METHODS, by Treves and Panzeri (1995), or by Panzeri and Treves (1996).

The selectivity of olfactory neurons to the class of odor stimuli is a related major aspect of how information is represented in the brain. Neurons may use either "local" or grandmother cell encoding, with strong or even great selectivity of the neuron for a particular environmental stimulus (Barlow 1972); or fully distributed representations in which all the neurons participate (Hinton et al. 1986); or sparse representations in which encoding by a sparse ensemble is used (Rolls and Treves 1990; Treves and Rolls 1991). Sparse ensemble encoding has some advantages of distributed representations such as generalization as the nature of the input changes, and graceful degradation if the network is incompletely formed or damaged. At the same time, sparse representations allow large numbers of representations to be stored and retrieved in associative neural networks (Rolls and Treves 1990; Treves and Rolls 1991), and in general, the sparseness of the representation has implications for how the network operates. Because of the importance of the sparseness of the representation for understanding processing in the orbitofrontal olfactory cortex, it was measured in this investigation. The sparseness measure used was one that can be directly applied to analyses of the storage capacity of networks of neurons (Treves and Rolls 1991, 1994). The olfactory neurons studied consisted of 38 olfactory neurons that responded differentially to different odorants in an olfactory discrimination task.

## METHODS

### *Recordings*

Recordings were made from single neurons in the orbitofrontal cortex, which included both medial and lateral areas in which olfactory responses have previously been described. A few neurons were also recorded in the adjacent olfactory regions (e.g., pyriform cortex, insula, and ventral striatum). The subjects were two male rhesus macaques (*Macaca mulatta*) weighing 2.5–3.5 kg. Neurophysiological methods were the same as described previously (Rolls and Baylis 1994; Rolls et al. 1976, 1990; Scott et al. 1986; Yaxley et al. 1990). All procedures, including preparative and subsequent ones, were carried out in accordance with the "Guidelines for the Use of Animals in Neuroscience Research" of the Society for Neuroscience, and were licensed under the UK Animals (Scientific Procedures) Act 1986. The monkey was fed on return to its home cage and was allowed access to water ad libitum. Glass-coated tungsten microelectrodes were constructed in the manner of Merrill and Ainsworth (1972) without the platinum plating. A computer (IBM 486 DX) collected spike arrival times and displayed on-line summary statistics or a peristimulus time histogram and rastergram. To ensure that recordings were made from single cells, the interspike interval was continuously monitored to make sure that intervals of <2 ms were not seen, and also the waveform of the recorded action potential was continuously monitored using an analog delay line.

### *Localization of recordings*

X-radiography was used to determine the position of the microelectrode after each recording track relative to permanent reference electrodes and to the anterior sphenoidal process. This is a bony landmark whose position is relatively invariant to brain structures (Aggleton and Passingham 1981). Microlesions (60–100 $\mu$A, 100 s) made through the tip of the recording electrode during the final tracks were used to mark the locations of typical units. These microlesions together with the associated X-radiographs allowed the positions of all cells to be reconstructed in the 50-$\mu$m brain sections with the methods described in Feigenbaum and Rolls (1991).

### *Stimuli*

Suprathreshold concentrations of stimuli were determined such that each odorant could be easily identified and discriminated in conditions of delivery identical to those employed in the experiment by the authors and colleagues, and were approximately equally intense. The odorants were selected as representative of differing putative odor classes (Amoore 1977) and were pure chemicals. They are listed in Table 1.

### *Reward associations*

The monkeys were able to acquire the task rapidly and work for the reward with nearly 100% accuracy. Eugenol was associated with the delivery of saline (and therefore was the S odorant) for the first monkey. The rewarded odorants in these experiments were phenylethanol, butyric acid, naphthalene, caprylic acid, citral, and amyl acetate. Eight cells from this monkey are included in this study. Of these, five neurons responded differentially only between

TABLE 1. *Odorants used in the olfactory discrimination task*

| Odorant | Quality | Concentration of Solution | Abbreviation |
|---|---|---|---|
| Eugenol | Cloves | 0.2 M* | eu |
| Hexylamine | Rotten fish | 0.2 M* | hx |
| Phenylethanol | Floral | 0.9 M* | pe |
| Butyric acid | Putrid | 0.05 M | bu |
| Naphthalene | Moth balls | 1.0 M* | na |
| Caprylic acid | Goaty/burnt plastic | 100% | cp |
| Citral | Citrus/boiled sweets | 0.1 M* | ct |
| Amyl acetate | Pear drops | 100% | aa |
| Vanillin | Vanilla | 1.0 M* | vn |

* Diluted in propylene glycol.

eugenol and the remaining odorants [1-way analysis of variance (ANOVA) with post hoc Newman-Keuls analysis]. To determine whether this was directly due to the reward association or due to some quality of eugenol as an odorant (such as trigeminal stimulation), the saline-associated odorant was changed to hexylamine in the task for the second monkey. Eugenol was a rewarded odor for the 30 cells collected from this monkey. At a later stage (for 20 cells in this monkey) a second odorant (vanillin) was also associated with saline, chosen because of its reported lack of trigeminal stimulation at discriminable concentrations. (Thus for 8 of the cells, there were 6 rewarded odors and 1 saline-associated odor; for 20 cells, there were 7 rewarded odors and 2 saline-associated odors; and for 10 cells, there were 7 rewarded odors and 1 saline-associated odor.)

### *Stimulus presentation and recording of neuronal activity*

The activity of single neurons in the orbitofrontal cortex and surrounding areas was recorded during the performance of an olfactory discrimination task. In the task, the delivery of one of eight different odors indicated that the monkey could lick to obtain a taste of sucrose. If one of two other odors was delivered from the olfactometer, the monkey had to refrain from licking, otherwise the monkey received a taste of saline. One reason for measuring responses in this task was that each odor had to be sniffed by the monkey as part of the discrimination, thus providing consistent olfactory stimulation from trial to trial. This was facilitated by signaling to the monkey with a 500-ms tone that the odor delivery was about to start, and allowing the monkey just long enough to make two licks for sucrose if the odor was sniffed immediately as its 1-s delivery began. The 38 of the 1,696 neurons (2.2%) that responded differently to the different odors in the task are considered here. The recording methods and neuronal population analyzed are described further elsewhere (Critchley and Rolls 1996a). Some of the analyses described here were performed on the responses to the seven odors (or, for a few cells, 6 odors) associated in the task with sucrose, so that differing reinforcement associations did not influence the measurements, which thus reflected odor quality independently of different associations with taste. Other analyses, as made clear in the RESULTS section, were performed on these odors and also on the one or two saline-associated odors. Where shown, the firing rates of the neurons are calculated from a 500-ms time window beginning 100 ms after onset of odorant delivery. This same time window was used in the principal component analysis of the spike train, from which informational measures were derived.

### *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) and Panzeri and Treves (1996). 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 we found that most of the information about which stimulus was presented is made evident by measuring the firing rate of the neuron, and temporal encoding adds relatively little additional information for this population of neurons (see Fig. 1), the information theoretic analyses described here and used for all data apart from those shown in Fig. 1 were based on the information available from the firing rate. The period in which this was measured was the poststimulus period of 100–600 ms with respect to the onset of the olfactory stimulus. (The method used for the information calculation shown in Fig. 1 is summarized below and described fully by Tovee and Rolls 1995 and Tovee et al. 1993).

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

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

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

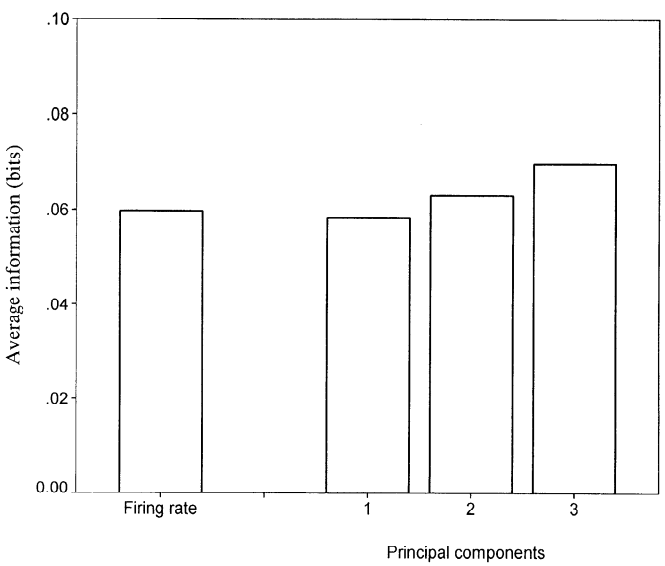


FIG. 1. Average information $I(S,R)$ about the set of 7–9 stimuli available in the responses of the neurons, averaged across 38 neurons. The information measures were calculated using the bootstrap correction described by Tovee et al. (1993) and derived from the firing rate, the 1st principal component of the response, the 1st 2 principal components of the response and the 1st 3 principal components of the response.

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

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

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

However, not all of this information is actually "about" the stimulus. Some of it may concern 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 at fixed 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 (*Eq. 5*) subtracted from the total information associated with that stimulus (*Eq. 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

$$I(s,R) = \sum_r P(r|s) \log_2 \frac{P(s,r)}{P(s)P(r)} = \sum_r P(r|s) \log_2 \frac{P(s|r)}{P(s)} \quad (7)$$

(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

$$I(S,R) = \sum_s P(s)I(s,R) = \sum_{s,r} P(s,r) \log_2 \frac{P(s,r)}{P(s)P(r)}$$

$$= \sum_{s,r} P(s)P(r|s) [\log_2 P(s|r) - \log_2 P(s)] \quad (8)$$

In the results we show both $I(s_i,R)$, the information available in the responses of the cell about each individual stimulus $s_i$, and $I(S,R)$, the average 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, 100–600 ms poststimulus unless otherwise stated (a unidimensional measure based on a firing rate measurement). (We also performed 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. For visual neurons, we 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 Rolls and Tovee 1995; Tovee and Rolls 1995; Tovee et al. 1993, 1994.) Although the set of stimuli can be discrete (it is in the present experiment), $R$ is generally a continuum, e.g., the firing rate in spikes per second. Because in practice one has to evaluate the expression for $I$ by performing a sum rather than an integral, $R$ needs to be quantized. 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 $D = 15$, cf. Optican and Richmond 1987), with the bin limits selected so that each bin contains the same number of trials (within $\pm 1$). For example, if 100 trials have to be allocated to 15 bins, the 1st bin extends from zero to midway between the lowest 7th and 8th response, the 2nd from there to midway between the lowest 14th and 15th response, and so on until the last bin, which extends from midway between the 94th and 95th response 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 $\sigma_n$ 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 $P(r) = \Sigma_{s \subset S} P(s,r)$.

Information values are in general dependent on the smoothing and binning procedures adopted, and most importantly 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-sized distortions in the information estimate, which as discussed by Treves and Panzeri (1995) requires that either $D$ be small (e.g., 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, as described by Tovee et al. (1993) and Treves and Panzeri (1995).

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 1st) 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. 1992). 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, and refer to Panzeri and Treves (1996) for a more detailed explanation

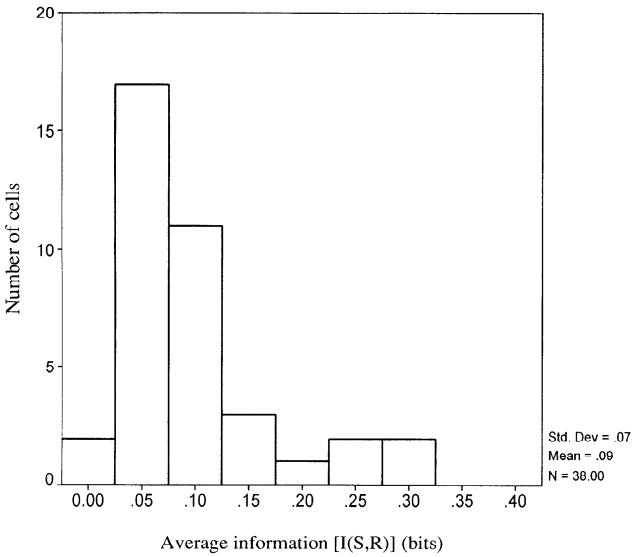


FIG. 2. Histogram showing the values of the average information, $I(S,R)$, about the set of 8–9 odor stimuli available in the responses of each of the 38 neurons.

of how the correction term is computed from the data. With respect to the correction based on the shuffling procedure, which we used in previous investigations (Tovee and Rolls 1995; Tovee et al. 1993), one should note that, although 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, which is described by Panzeri and Treves (1996).

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.

CALCULATION OF INFORMATION BASED ON A PRINCIPAL COMPONENT ANALYSIS. An additional aim of the analysis was to investigate whether information was available about which stimulus was presented in the temporal pattern of arrival of the spikes. For example, one odor might produce a sudden onset of the neuronal response, and another odor a slower onset, and information might be available from this. To analyze the possibility of temporal encoding by these olfactory neurons, the principal components of the time series of the smoothed spike trains to each stimulus were used. A related aim was to investigate whether information of comparable magnitude was present in the firing rate response of the neuron, and for this a single measure of the response (in the poststimulus period 100–600 ms) was used. To perform this calculation, we performed the analysis using the same methods described fully by Tovee et al. (1993; Tovee and Rolls 1995), and we do not redescribe those methods here. The information calculated in this way is shown in Fig. 1 only.

### *Sparseness of the representation*

The sparseness, $a$, of the representation of a set of (odor) stimuli provided by these neurons can be defined and was calculated as

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

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 analyses of the capacity of neural networks that use an approach derived from theoretical physics (see Rolls and Treves 1990; Treves 1990; Treves and Rolls 1991). 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 (e.g., 0 or 100 spikes/s) (Treves and Rolls 1991). A third is that it makes no assumption about the form of the firing rate distribution (e.g., binary, ternary, exponential, etc.), 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 in which 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 (Rolls and Treves 1990; Treves and Rolls 1991) how the system operates.

As described in RESULTS, 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 0). 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 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 on the basis of 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 (Rolls and Treves 1990; Treves and Rolls 1991).

## RESULTS

Thirty-eight olfactory neurons in the orbitofrontal cortex with significant differential responses to odorants presented during the olfactory discrimination task (shown by 1-way ANOVAs with post hoc Newman-Keuls analysis) are the source of the data set analyzed here. Recordings from 1,696 orbitofrontal neurons were performed to obtain this set of olfactory neurons. The responses of these neurons have been described elsewhere (Critchley and Rolls 1996a). Sixty percent of the 38 olfactory neurons with differential responses in the task showed differential selectivity for the stimuli based on the odor quality, and not on the taste reward associ-

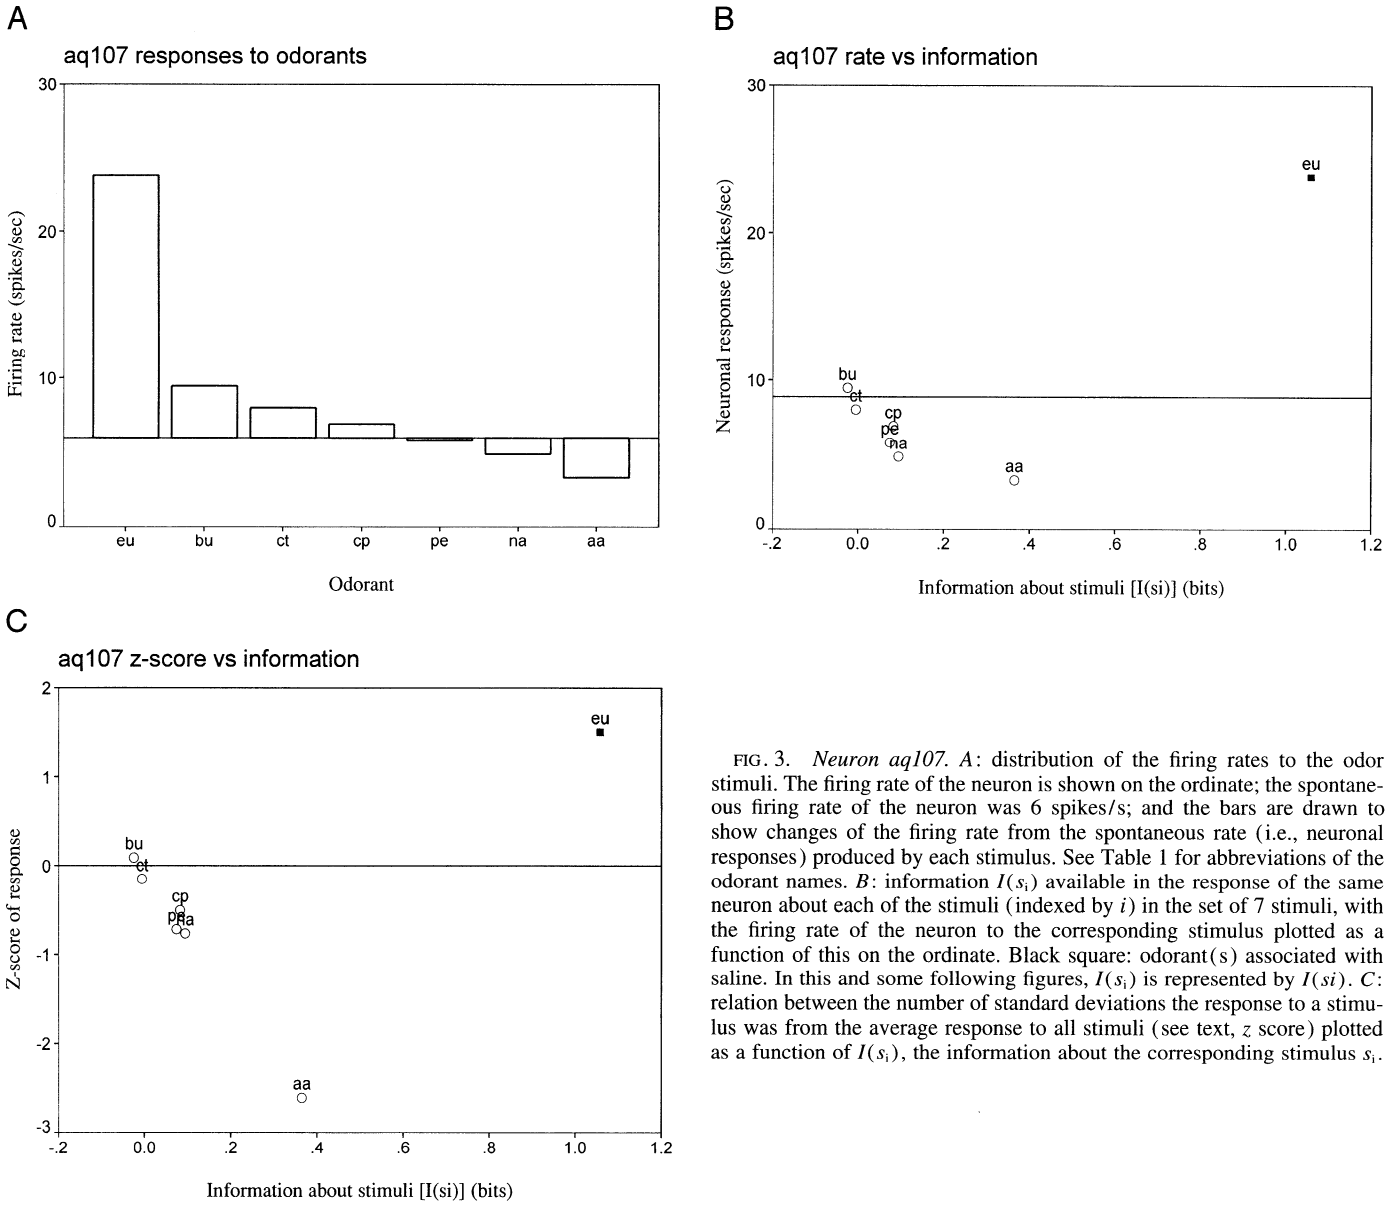


FIG. 3. *Neuron aq107. A*: distribution of the firing rates to the odor stimuli. The firing rate of the neuron is shown on the ordinate; the spontaneous firing rate of the neuron was 6 spikes/s; and the bars are drawn to show changes of the firing rate from the spontaneous rate (i.e., neuronal responses) produced by each stimulus. See Table 1 for abbreviations of the odorant names. *B*: information $I(s_i)$ available in the response of the same neuron about each of the stimuli (indexed by $i$) in the set of 7 stimuli, with the firing rate of the neuron to the corresponding stimulus plotted as a function of this on the ordinate. Black square: odorant(s) associated with saline. In this and some following figures, $I(s_i)$ is represented by $I(si)$. *C*: relation between 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 $I(s_i)$, the information about the corresponding stimulus $s_i$.

ation of the odor. They responded, for example, to some but not to others of the seven olfactory stimuli that were all associated with the delivery of sweet taste. Forty percent of the 38 olfactory neurons responded on the basis of the taste reward association of the odorants. Such neurons responded either to all the rewarded stimuli and to none of the saline-associated stimuli, or vice versa.

The average information $I(S,R)$ about the set of olfactory stimuli available in the responses of the neurons, averaged across all 38 neurons, is shown in Fig. 1. The neurons reflected in their firing rates on average 0.06 bits of information about which odorant was presented. With respect to the temporal encoding analysis, it was found that the majority of information was carried by the first principal component ($\sim$0.06 bits on average), with the second and third principal components (which reflect different temporal time courses to the first principal component) together adding less than a further 0.01 bits of information about the stimulus set. This demonstrates that little information was carried by the temporal firing pattern of neurons and that the major part of the information was carried in the firing rates of the neurons. The average information calculated with this method is a little lower than that calculated with the analytically based correction procedure of Treves and Panzeri (1995) used in the remainder of this paper, because the method used for Fig. 1 used a bootstrap correction procedure, which with the limited number of trials available is quite conservative in the information estimate produced (Tovee et al. 1993; Treves and Panzeri 1995).

A histogram showing the values of $I(S,R)$ (the average information in the responses of a cell about the stimulus set) for each cell is provided in Fig. 2. In this and the subsequent figures, the information measure is derived from the firing rate of the response. Most of the neurons had values for $I(S,R)$ of $<0.1$ bits, with the average across the population of neurons being 0.09 bits, as shown in Fig. 1. A few neurons carried $>0.1$ bits of information, yet the majority of neurons carried little information about the odorant set as a whole.

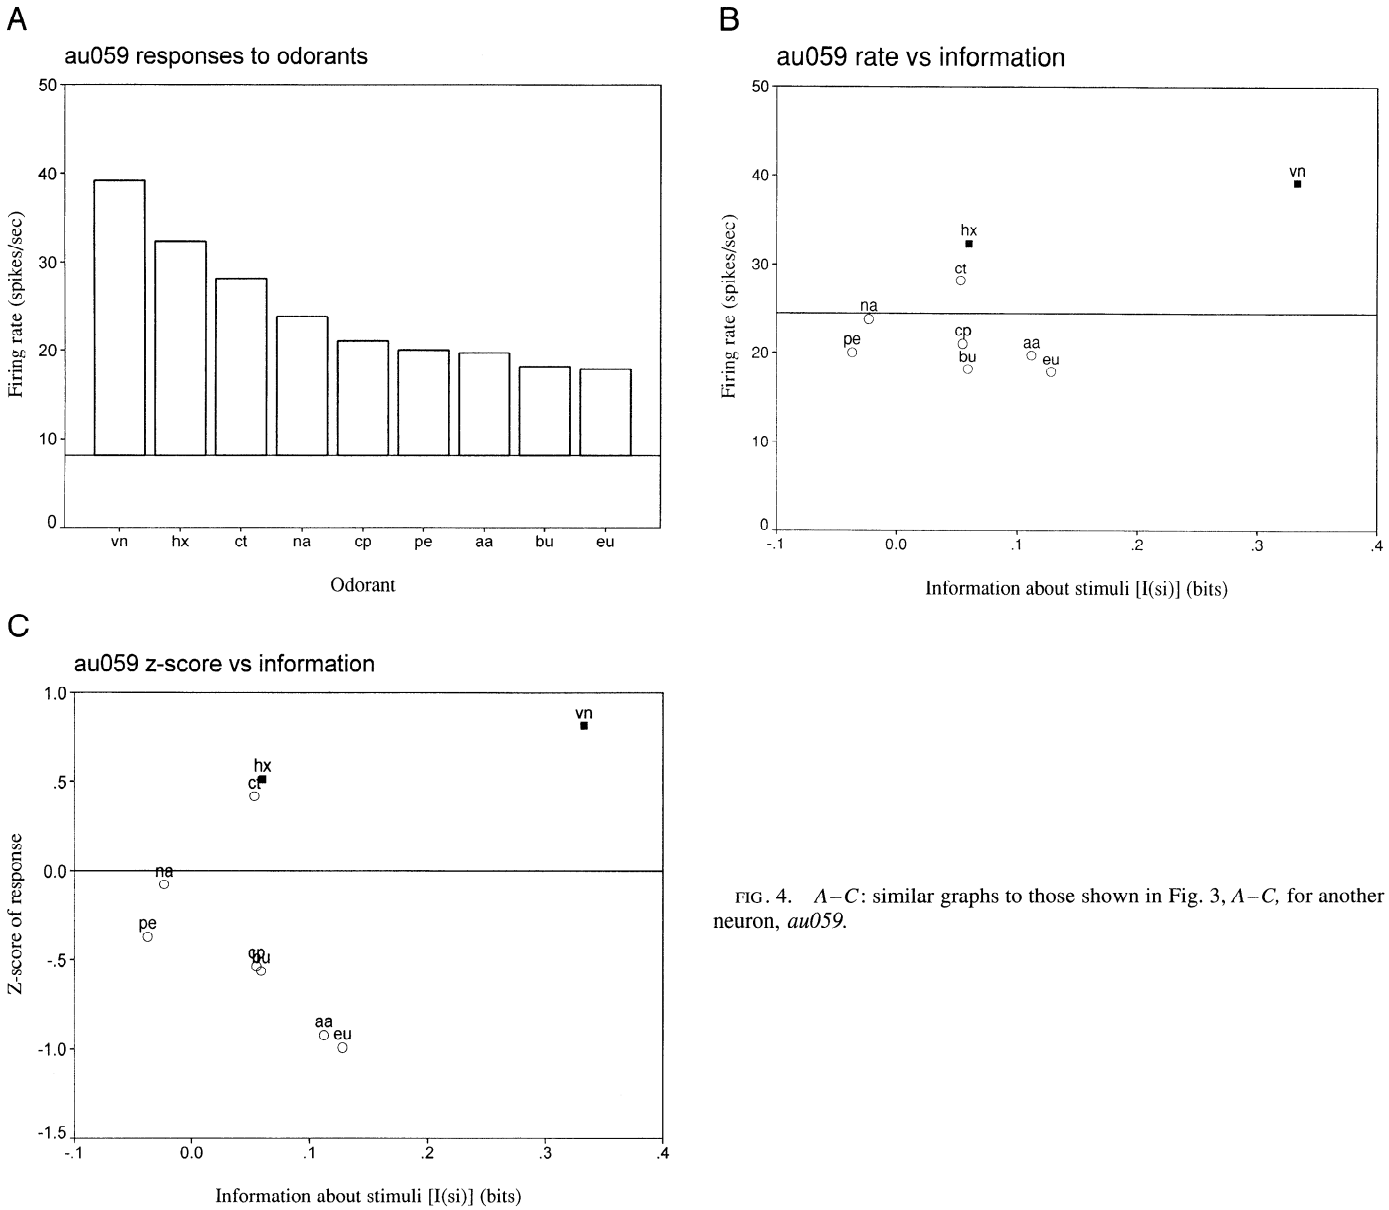


FIG. 4. *A–C*: similar graphs to those shown in Fig. 3, *A–C*, for another neuron, *au059*.

This is despite evidence that some neurons were very selective in their evoked firing rates to one or a few of the odorants, and that all the neurons showed a significant differential response in one-way ANOVAs and post hoc Newman-Keuls analysis. The neuronal information measure described here is equivalent to the average of all the information contained in the responses to the individual stimuli (corrected for the minor differences in the number of trials). If many of the odorants evoke a similar neuronal response, then the average information from the neuronal response about which odor was delivered is low.

To understand the representation of individual stimuli by individual cells, the information $I(s_i)$ available in the neuronal response about each of the stimuli (indexed by $i$) in the set of stimuli $S$ was calculated. First we show in Fig. 3*A* the response profile of a single olfactory neuron to the odor stimuli. The cell responded strongly to one odorant, eugenol, which was negatively reinforced, but did not respond well to the other odorants, which were rewarded with a sweet taste. The sparseness of the representation was 0.66, and the average information $I(S,R)$ provided by the cell across all stimuli in the set $S$ was 0.25 bits. In Fig. 3*B* we show the information $I(s_i)$ available in the neuronal response about each of the stimuli in the set of stimuli, plotted along the abscissa. It is shown that considerable information, ~1.1 bits, could be gained from the neuronal response when one of the stimuli, eugenol, was presented, and that very little information (on average 0.01 bits of information) could be gained when any of the other stimuli were presented. However, 0.37 bits of information were reflected in the neuronal responses to amyl acetate.

Figure 3*B* also shows the relation between $I(s_i)$ and the firing rate to the $i$th stimulus. The odorant eugenol, which was unique among the set in eliciting a high firing rate from the cell, has a high information value of 1.1 bits. The neuron conveyed very little information about the remaining odorants, which are largely indistinguishable from each other by firing rates. Given that the information $I(s_i)$ about the $i$th

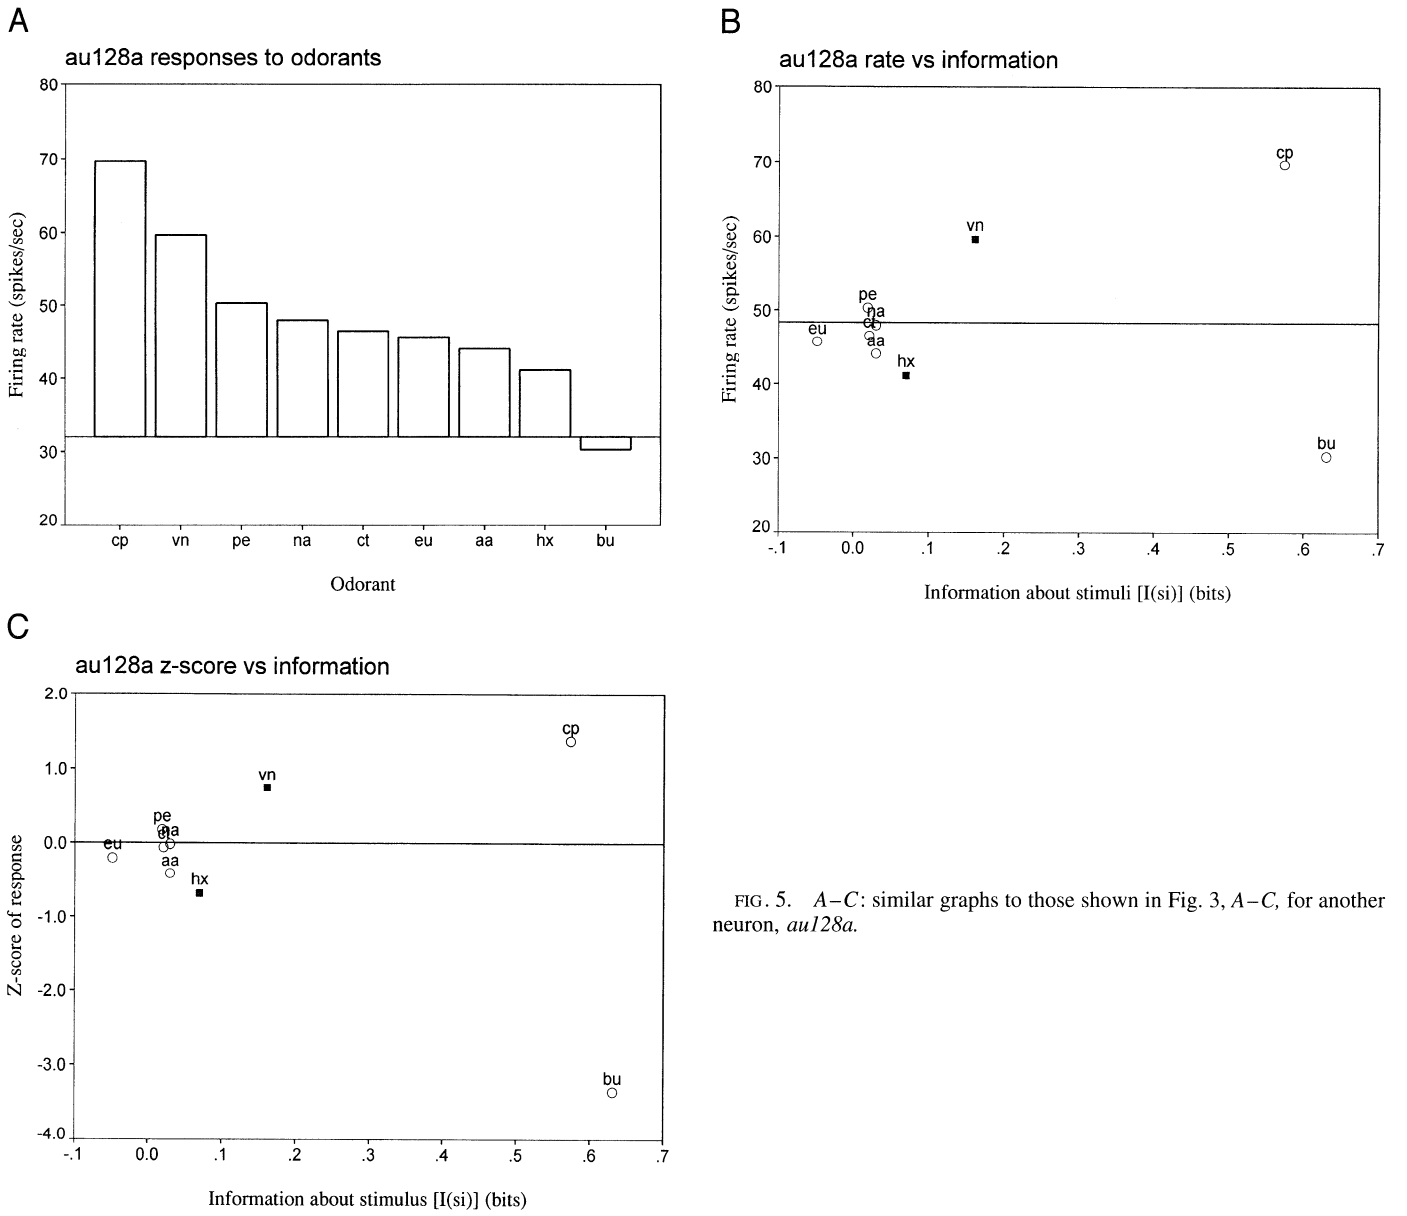


FIG. 5. *A–C*: similar graphs to those shown in Fig. 3, *A–C*, for another neuron, *au128a*.

stimulus from the neuronal response is related to the probability that that particular response will occur, it is natural that the information should be closely related to the number of standard deviations that the response to a particular stimulus is from the mean neuronal response. The greater the number of standard deviations (i.e., the greater the $z$ score) from the mean response value, the greater the information might be expected to be. Indeed, given that the firing rates of neurons tend often to follow a Poisson distribution (apart from the refractory period), for which the variance is proportional to the mean, a more proportional relationship might be expected between the information and the $z$ score than between the information and the firing rate. We therefore show in Fig. 3*C* the relation between the $z$ score and $I(s_i)$. The $z$ score was calculated by calculating the mean and standard deviation of the response of a neuron to a particular stimulus, and dividing the difference of this response from the mean response to all stimuli by the calculated standard deviation for that stimulus. (To the degree to which the firing is Poisson-like, the standard deviation for a stimulus below the mean will be smaller than for a stimulus above the mean, and this will be reflected in the calculated $z$ scores for stimuli below and above the mean.) It can be seen that most information is present on trials on which stimuli are presented that produce responses that are far from the mean response in terms of $z$ score (eugenol and amylacetate). The $z$ scores underline the point that the variability in the firing rate to a particular stimulus as well as the mean response to that stimulus may be important. We will see soon that the $z$ score is useful for cells that display a more distributed representation of the stimuli.

A similar analysis for a cell with a more distributed representation is shown in Fig. 4. The response profile of the cell to the different odorants is shown in Fig. 4*A*. The sparseness of the representation was 0.93, and the average information $I(S,R)$ provided by the cell across all stimuli in the set $S$ was 0.09 bits. The cell responded best to the negatively reinforced odorants, in this case hexylamine and vanillin,

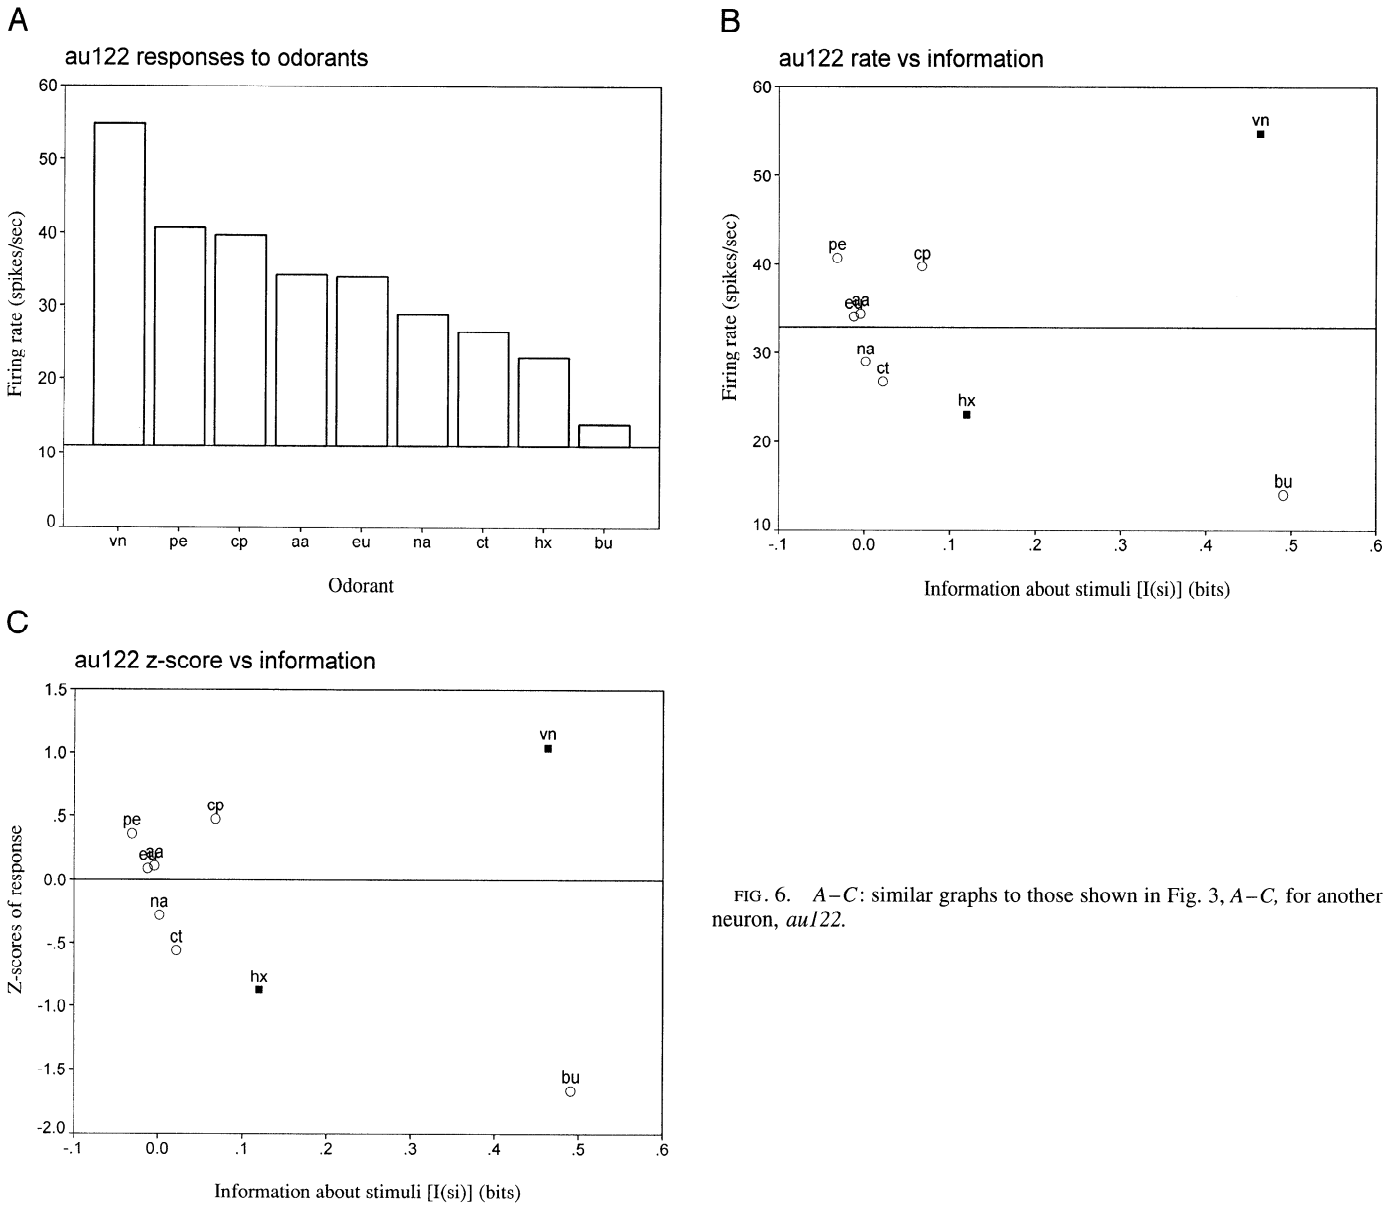


FIG. 6. *A–C*: similar graphs to those shown in Fig. 3, *A–C*, for another neuron, *au122*.

yet had moderate responses to most of the rewarded odorants. The information conveyed about each odorant is plotted against the neuronal firing rate in Fig. 4*B*. The responses to the odorant vanillin carry the most information, almost 0.35 bits, and the responses to the remaining odorants convey less information. (The small amounts of "negative information" arise because of the correction procedure that subtracts what is an estimated correction from the raw information.) Associated with the more distributed representation provided by the responses of the cell shown in Fig. 4, the information about the most effective stimulus, vanillin, does not reach the same level as that of the most effective stimulus, eugenol, for the cell shown in Fig. 3. The greater amount of information reflected by the response to eugenol of the cell shown in Fig. 3 is related to the fact that this was a low probability response. Another interesting aspect of the data for the cell shown in Fig. 4 is that hexylamine, the second most effective stimulus in terms of the firing rate elicited, conveys almost the same amount of information as the least effective stimuli for the cell. This emphasizes the important point that information could be carried by a response that is below the mean response to odors for the cell, again reflecting the fact that a low firing rate response is a relatively improbable response for the cell. Indeed, this point is brought out explicitly by the $z$ score analysis shown in Fig. 4*C*, which 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. This results in a C-shaped curve in Fig. 4*C*, with more information being provided by the cell the further its response to a stimulus is in $z$ scores either above or below the mean response to all stimuli (which was 24.4 spikes/s). The result is clearer with the $z$ scores than with the firing rate shown in Fig. 4*B*, because the $z$ score takes into account the variability of the response. Because of the Poisson-like firing of neurons, and the way in which the standard deviation varies with the mean, the effect of transforming the firing rate to $z$ scores

away from the mean is to clarify the way in which responses to a stimulus that are below the mean response of the cell to all stimuli can convey considerable information. 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 olfactory 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 olfactory stimuli, and if so, of what type of neuronal networks would best enable this information to be utilized.

Figure 5 shows the response profile (Fig. 5*A*), firing rate versus information $I(s_i)$ (Fig. 5*B*), and $z$ score versus information $I(s_i)$ (Fig. 5*C*) of a neuron maximally responsive to the rewarded odorant caprylic acid but much less responsive to the rewarded odorant butyric acid. [The sparseness of the representation was 0.97, and the average information $I(S,R)$ provided by the cell across all stimuli in the set $S$ was 0.16 bits.] The caprylic acid responses carry >0.57 bits of information. The responses to the butyric acid convey 0.63 bits of information. Figure 5*C* shows that both the caprylic acid and the butyric acid elicit neuronal responses that have a low probability with respect to the mean overall firing rate of the neuron (47.7 spikes/s) to the set of odor stimuli $S$. The information $I(s_i)$ contained in the responses to butyric acid illustrates the phenomenon described above in which low firing rates, in this case not different from the spontaneous rate, may carry information if they are improbable or unique.

Figure 6 illustrates the responses of a similar neuron, responsive to the odorant vanillin, and again not responsive to butyric acid. [The sparseness of the representation was 0.90, and the average information $I(S,R)$ provided by the cell across all stimuli in the set $S$ was 0.13 bits.] The vanillin odorant responses carry 0.46 bits of information, and the very small responses (relative to the spontaneous firing rate of the neuron) to butyric acid carry 0.49 bits of information. The curves of firing rate against information (Fig. 6*B*) and of $z$ scores against information (Fig. 6*C*) illustrate the C-shaped curve of the relationship between information and the neuronal response. The overall mean firing rate across all stimuli was 32.8 spikes/s.

It has been shown previously that the reward value of the odorants may strongly influence the responses of some (35%) of these particular neurons in the olfactory discrimination task (Critchley and Rolls 1996a). To determine the degree to which information is reflected about the reinforcement value of the odorants, information theoretic analysis was applied to the responses to stimuli that were grouped as rewarded odorants and saline-associated odorants. In a similar manner, to quantify the amount of information concerning stimulus quality independently of reward value, information analysis was performed on the responses of each cell to only the reward-associated odorants. (There were 6–7 reward-associated odorants and 1–2 saline-associated odorants when each neuron was tested.) Table 2 shows the average information conveyed by the responses of each neuron about the reinforcement value of odorants and about the identity of the odorants. (The information about identity was calculated from the 6 or 7 stimuli that were all equally associated with reward, so that differences in their firing reflected the identity of the odor stimulus presented.) Where the information analysis produced a negative value, this is denoted as 0. It can be seen that the average information about the odorant identity is low (cf. Tovee et al. 1993 for visual cells). The most information that any one cell conveyed about odorant quality was 0.30 bits and the average information $I(S,R)$ over the 38 neurons about which of the rewarded odorants was delivered is 0.06 bits. This indicates that although the cells often respond robustly to the presence of odorants, neurons in this population do not individually express much information about odorant quality. Instead they are likely to contribute to a very distributed representation of information about which odorant is present. The information conveyed by the neurons about the reward value of the odorants averaged 0.04 bits, with 5 cells conveying >0.1 bits of information about the taste reinforcement association. Figure 7 shows for each cell the information about the reinforcement association of the odor (ordinate) as a function of the information about which odor was delivered (with reinforcement constant). It can be seen that at least some neurons appear to reflect primarily taste quality, and others the reinforcement association. The larger amount of

TABLE 2. *Information provided by each cell about odor identity versus reward value*

| Cell | Information About Reward Value | Information About Odor Identity |
|---|---|---|
| *aq105* | 0.231 | 0.072 |
| *aq107* | 0.199 | 0.035 |
| *aq078b* | 0.196 | 0.196 |
| *au037b* | 0.122 | 0 |
| *au011* | 0.106 | 0 |
| *au114* | 0.074 | 0.010 |
| *au059* | 0.069 | 0.022 |
| *au079* | 0.063 | 0.139 |
| *aq079* | 0.058 | 0.047 |
| *au037a* | 0.051 | 0.011 |
| *au161* | 0.050 | 0.023 |
| *au070* | 0.049 | 0.019 |
| *aq112* | 0.044 | 0.006 |
| *au031* | 0.042 | 0 |
| *au097* | 0.042 | 0 |
| *au032a* | 0.041 | 0.082 |
| *au032b* | 0.040 | 0.126 |
| *au030* | 0.032 | 0.41 |
| *au086* | 0.042 | 0 |
| *au033* | 0.020 | 0.080 |
| *au083* | 0.019 | 0.009 |
| *au150* | 0.018 | 0 |
| *au010a* | 0.016 | 0.029 |
| *au119* | 0.016 | 0.065 |
| *au034* | 0.015 | 0.032 |
| *aq092* | 0.008 | 0.025 |
| *au005* | 0.003 | 0.300 |
| *au071* | 0.002 | 0 |
| *au122* | 0.002 | 0.082 |
| *au028* | 0.001 | 0.016 |
| *aq078a* | 0.001 | 0.033 |
| *aq113b* | 0.000 | 0.113 |
| *au089* | 0 | 0.102 |
| *au159* | 0 | 0.049 |
| *au111* | 0 | 0.121 |
| *au142* | 0 | 0.070 |
| *au128a* | 0 | 0.167 |
| *au158* | 0 | 0.061 |

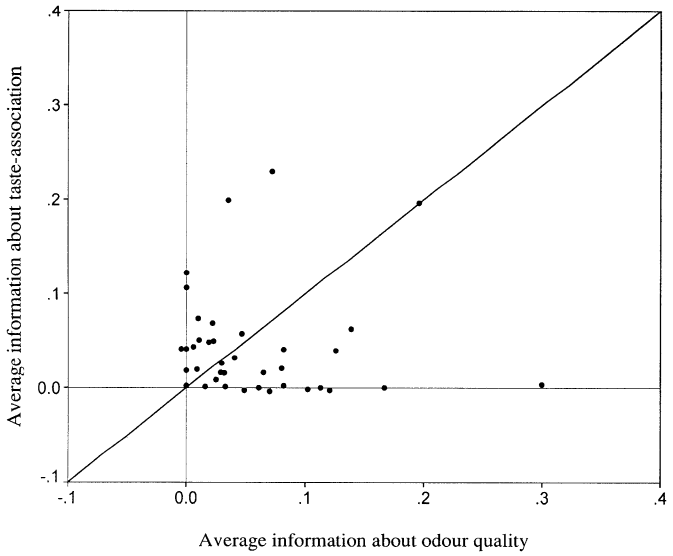


FIG. 7. Average information for each cell about the reinforcement association of the odor (ordinate), plotted as a function of the information about which odor was delivered.

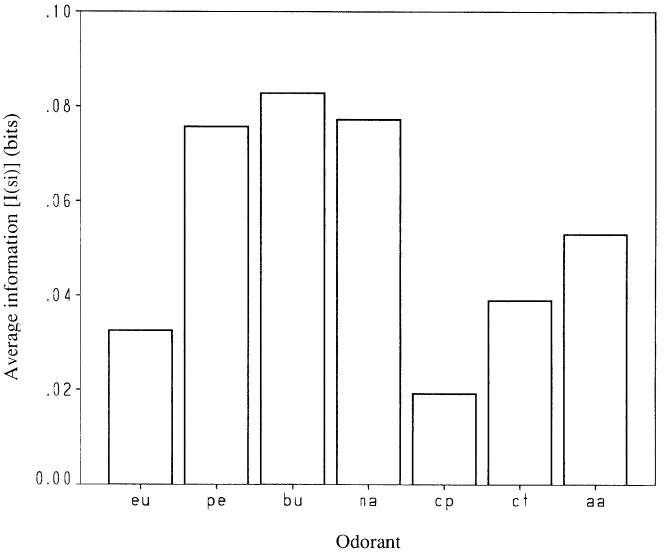


FIG. 9. Values of the information about any 1 stimulus, $I(s_i)$, averaged across 30 neurons for eugenol and across 38 neurons for the remaining odorants.

information carried by some neurons about the reinforcement value of the stimuli adds to further evidence described elsewhere (Critchley and Rolls 1996a) that the responses of some olfactory neurons in the primate orbitofrontal cortex are strongly influenced by the taste association of odorants.

To further show the extent to which these neurons can contribute to the discrimination of odorants independently of reward association, the information about each of the rewarded odorants was calculated for each neuron. Figure 8 shows the mean responses over the neurons to the individual rewarded odorants. Each bar represents the mean of 38 neurons, except for eugenol, which is derived from the 30 neurons in which eugenol was rewarded. It can be seen that the overall representation of each odorant is approximately equal in terms of the firing rate of the neurons, each of the odorants evoking between 16 and 20 spikes/s averaged across the population of cells. Figure 9 shows the corresponding bar plot for the averages of the information $I(s_i)$ held about each odorant $i$. Overall, high levels of information are not conveyed about which odor was presented by the responses to the set of odorants that were a set in that they were all equally associated with reward. Figure 9 shows that the average information about each of the odorants conveyed by this population of neurons was between 0.02 and 0.08 bits. The least amount of information, on average 0.02 bits, was reflected in the responses of the cells to caprylic acid. This is an indication that the response to caprylic acid was consistently more similar than the other odorants to the average responses of the neurons. For only a very few neurons (e.g., Fig. 5) was the response to caprylic acid clearly different from the mean to other odorants.

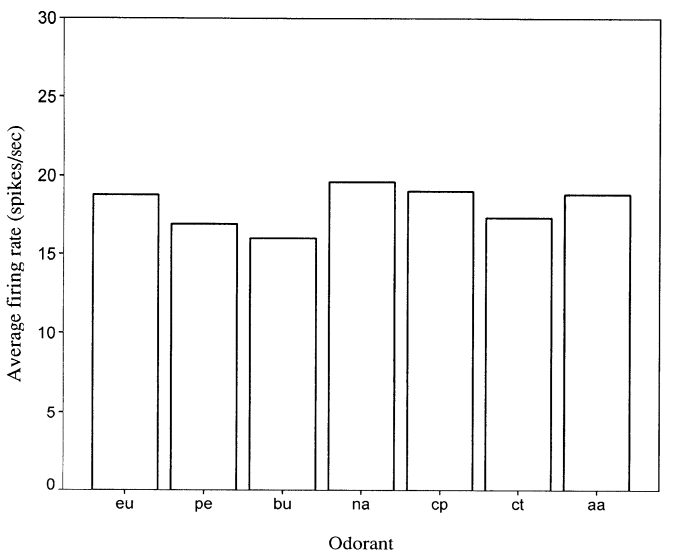


FIG. 8. Mean evoked firing rate to each of the odorants (of equal reinforcement value), averaged across 30 neurons for eugenol and across 38 neurons for the remaining odorants.

The data above indicate that the encoding of information about odorants in this population of neurons is likely to be achieved by very distributed encoding. To quantify this, the measure $a$ of the sparseness of the representation was calculated. The sparseness measure indicates the length of the tail of the distribution of neuronal responses to the stimuli, such that low values indicate high selectivity to one or a few of the stimuli in the set, and a value of 1.0, if the neurons had binary firing rates (e.g., firing or not), would indicate equal responses to all the stimuli. Sparseness was calculated in the manner described in the METHODS section, and was performed from the raw firing rates of the neurons ($a$), and from the responses of the neurons with the spontaneous subtracted with clipping at zero response $a_r$.

Figure 10 shows the sparsenesses $a$ of the neuronal responses to the rewarded odorants. It can be seen that all the neurons had high values for $a$, that is the coding was very distributed. The mean value for $a$ was 0.936, indicating a very distributed encoding of odorant quality in these cells. The response sparseness $a_r$ (based on the evoked responses minus the spontaneous firing of the neuron) is shown in Fig. 11. The mean of the response sparseness was 0.78.

The sites at which these neurons were recorded are shown

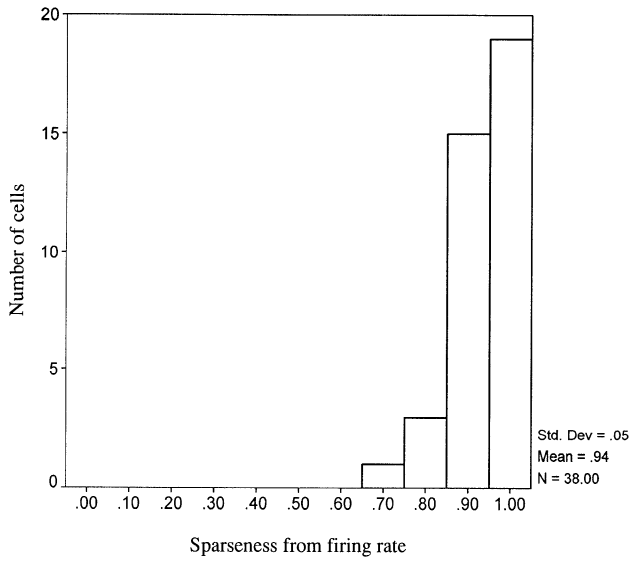


FIG. 10. Distribution of sparseness values $a$ for the population of cells, calculated across the set of odorants of equal reinforcement value. The mean sparseness was 0.94 ± 0.05 (SD).

in Fig. 12. The majority of the neurons were in the orbitofrontal cortex, with four in the neighboring regions of the pyriform cortex, ventral striatum, and insula.

## DISCUSSION

The results described above provide a quantitative analysis of the information represented by olfactory neurons in the orbitofrontal cortex, and of how it is represented. Some of the main conclusions are as follows. First, the average information about the set of stimuli, $I(S,R)$, is low on average for the population of neurons. This implies that a large ensemble of neurons is needed to represent the set of odors sufficiently to allow identification of which odor was presented. Studies are in progress to determine what the minimal size of this ensemble would need to be to enable identification of this set of eight odors. A comparison with the average information provided by neurons in other modalities is given below. The relatively low information about odors represented by this population of neurons may reflect the situation that primates are not especially good at identifying which particular odor in a large set was presented, that is that olfaction is not a modality in which many different stimuli can be clearly and separately represented. Another possibility is that there are neurons in other parts of the olfactory system that are very good at representing the stimulus space of odors, and that the orbitofrontal neurons analyzed here are more involved in coding odors in relation not to their stimulus quality but to their association with events in other modalities, such as whether the odor is associated with a good taste. Certainly some of the neurons in this region are influenced by taste reward association, and because of this, such neurons would become less effective at representing the stimulus quality. Second, the tuning of most of the orbitofrontal olfactory neurons is broad, with high values for the sparseness measure $a$. Broad tuning (i.e., distributed encoding) of this type may be particularly effective when the system is noisy, that is, when there is relatively little information about the stimulus (Linsker 1992). Third, the information about each individual stimulus, $I(s_i)$, could be shown to be simply related to the type of tuning shown by the neuron. If the neuron responded mainly to one odor in the set, then high information about that stimulus was provided by the firing of the neuron to that stimulus, and little information was provided about the other stimuli. This is local encoding. If, on the other hand, the neuron used very distributed encoding, then information about several individual stimuli could be provided by the rate of firing of the neuron, and the amount of information could be shown to be simply related to the probability of the neuron responding differently to a particular stimulus relative to the overall mean response of the neuron to all stimuli, as reflected by the $z$ score measurements. Of particular interest was that for such neurons, significant information was potentially available if the neuron responded much less to one or some of the odors than its average response across all the odors.

### *Comparison with information conveyed by neurons in other modalities*

Information theoretic analysis has been applied to the responses of temporal lobe visual neurons responsive to faces (Rolls and Tovee 1995; Tovee et al. 1993) and to orbitofrontal taste-responsive neurons (Critchley and Rolls, unpublished data). These studies can be compared with each other and with the results presented here for olfactory neurons, because they use a similar method for the calculation of raw information measures and a correction procedure to correct for the low number of trials to each stimulus. For the face-responsive neurons in the superior temporal sulcus temporal lobe region, Tovee et al. (1993) and Rolls and Tovee (1995) showed that the average information $I(S,R)$ conveyed by these neurons was between 0.3 and 0.6 bits (information measures derived from the firing rate or from the 1st principal component were shown to be largely equivalent). Face-responsive units held 0.4 bits of information about a set of 20 faces (Rolls and Tovee 1995). For taste-responsive neu-

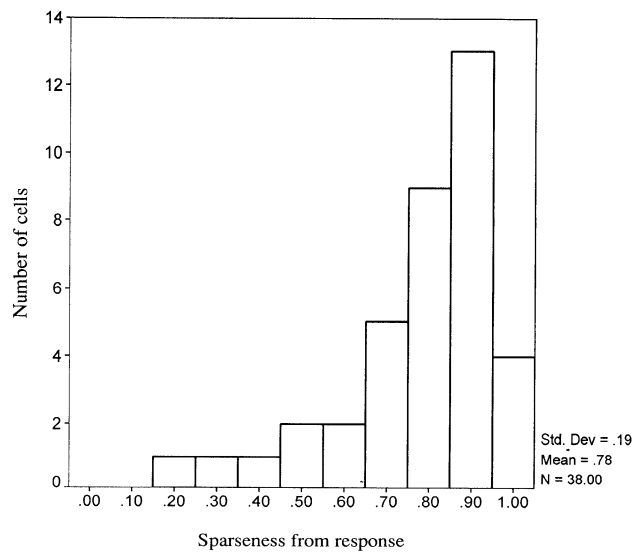


FIG. 11. As Fig. 10, but response sparseness values $a_r$ are shown. The mean response sparseness was 0.78 ± 0.19 (SD).

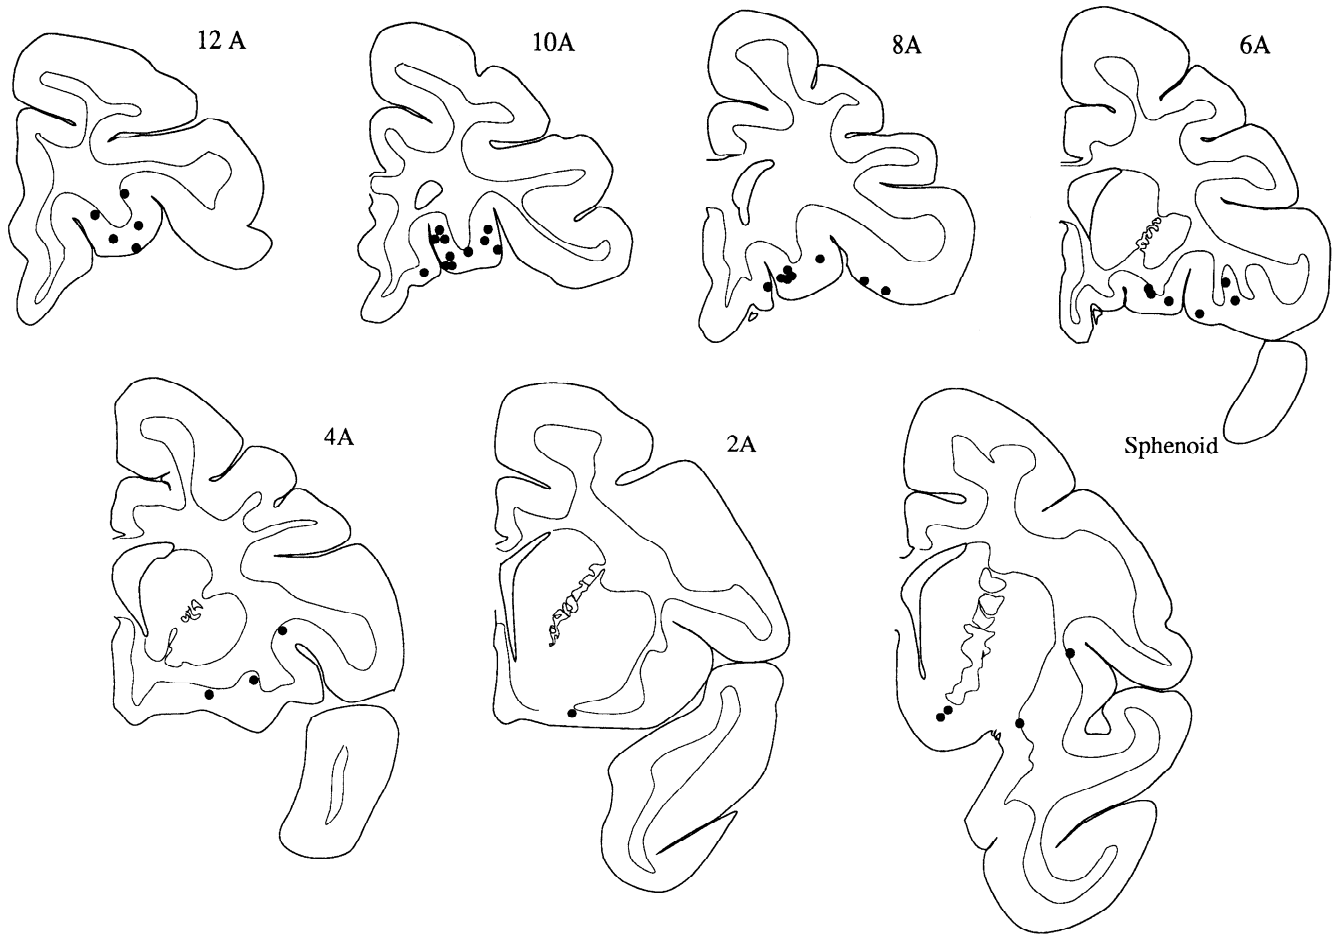


FIG. 12. Sites at which the neurons were recorded. The majority of the neurons were in the orbitofrontal cortex, with 4 in the caudally adjacent olfactory areas.

rons also in the primate orbitofrontal cortex encountered in this study, 0.43 bits of information were conveyed in the neuronal responses to the stimuli, glucose, HCl, NaCl, quinine, monosodium glutamate, and distilled water. In the present study, 0.09 bits of information were contained within the responses of orbitofrontal olfactory cells to the odorants in the olfactory discrimination task. The following factors may contribute to the low information and high sparseness measures of the orbitofrontal olfactory neuron. First, there is generally much more variation in the responses of the olfactory neurons to the odorants. (The more noisy the neuronal response, the less information it will reflect about the stimulus.) This is likely to be the result of different latencies in the sniffs made by the monkey in order to sample the odorants. Because these sniff responses need to be integrated with the normal respiratory cycle of the monkey, a degree of variation is unavoidable. Attempts to reduce this were made; a cue tone was sounded before the onset of odorant deliveries, and the monkey did generally start to sniff just as the tone ended. Also, the presence of saline-associated odorants required the monkey to be attentive to and sample each odorant before making a lick response within the 1-s duration of the odorant presentations. This provided some temporal restriction of the sniffs. The relatively fixed onset latencies of some of the neurons recorded provided evidence that the odor sampling in the task was adequate, although for the reasons stated not as good as can be obtained with visual neurons when a monkey is performing a visual fixation task (Rolls and Tovee 1995; Tovee et al. 1993).

*Information contained within the spike train*

The average information in the responses of the olfactory neurons to the stimulus set was 0.06 bits (calculated with the bootstrap method used for Fig. 1). Analysis of the contribution of the firing rate and first three principal components of the response spike train, showed that the first principal component conveyed the major part of the information about odorants (see Fig. 1). The first principal component is highly correlated with the firing rate of the neurons. Higher principal components do not contribute much information. This indicates that at least for orbitofrontal olfactory neurons, temporal encoding does not have a large role in the representation of odors. This parallels similar findings for face-responsive visual neurons in the temporal lobe visual cortex by Tovee et al. (1993) and Rolls and Tovee (1995).

The analyses shown in Figs. 3, *B* and *C*, 4, *B* and *C*, 5, *B* and *C*, and 6, *B* and *C*, show that the information reflected in the neuronal response when a particular stimulus was delivered was related to how far the firing rate to that stimulus was from the mean firing rate to all stimuli. The relationship was stronger when the number of standard deviations

of the response to a stimulus from the mean response to all stimuli was considered (Figs. 3–6*C*). This relationship arises because the information measure reflects how probable the response is, and most information is gained when improbable responses occur. We note that because the question being asked is which stimulus was delivered, the relevant comparison is the mean response to all stimuli. If the question had been about whether an olfactory stimulus was delivered, then the data would have been collected differently, with many trials with no olfactory stimulus, and the relevant comparison would have been to the spontaneous firing rate. At the same time, we can note that the information provided by a neuron, both on average across the stimulus set, and about each odor in the set, does depend on the stimulus set used. For example, if a stimulus set with very similar odors is used, then the average information provided by a neuron will be low. In another example, if there were only two odors in the set, the maximum information that could be provided would be 1 bit.

### *Encoding of stimulus quality and reinforcement value*

The indication that many of these orbitofrontal olfactory neurons are concerned more with the representation of the reinforcement association of the odorants rather than the discrimination of odorants of equal reward values comes from the analysis of the average amount of information held by the neurons about the reward association, and about the identity of odorants independent of reward association (see Table 2 and Fig. 7). Some neurons reflected information primarily about odor identity, and others about odor association. In a number of cases, when very selective responses of these orbitofrontal neurons do occur, this is usually due to differences in the reinforcement association of the stimuli to which the neuron responds differently (Critchley and Rolls 1996a). This is important in understanding the purpose of the olfactory representation in the orbitofrontal cortex. One means by which this might occur is by associative learning between odor and taste, in such a way that an initially broadly tuned olfactory neuron might by virtue of specific taste inputs come to respond to the odors associated with one but not with another taste. This would give rise to the cross-modal correspondence described by Rolls and Baylis (1994). Further evidence for such a mechanism in primates is provided elsewhere, although associative learning between odor and taste is relatively slow and inflexible compared with visual-to-taste association learning (see Rolls et al. 1996).

We are very grateful for helpful comments and assistance provided by Prof. L. A. Abbott and Dr. R. Baddeley.

This research was supported by Medical Research Council Grant PG8513790 to Dr. E. T. Rolls, by the Human Frontier Science Program, and by EC Human Capital and Mobility Grant ERB-CHRX-CT93-0245.

Present address of A. Treves: S.I.S.S.A.-Biofisica, via Beirut 2-4, 34103 Trieste, Italy.

Address reprint requests to E. T. Rolls.

Received 6 February 1995; accepted in final form 25 October 1995.

## REFERENCES

AGGLETON, J. P. AND PASSINGHAM, R. E. Syndrome produced by lesions of the amygdala in monkeys (*Macaca mulatta*). *J. Comp. Physiol. Psychol.* 95: 961–977, 1981.

AMBROS-INGERSON, J., GRANGER, R., AND LYNCH, G. Simulation of paleocortex performs hierarchical clustering. *Science Wash. DC* 247: 1344–1348, 1990.

AMOORE, J. L. Specific anosmia and the concept of primary odors. *Chem. Senses Flavor* 2: 267–281, 1977.

BARBAS, H. Organization of cortical afferent input to orbitofrontal areas in the rhesus monkey. *Neuroscience* 56: 841–864, 1993.

BARLOW, H. B. Single units and sensation: a neuron doctrine for perceptual psychology? *Perception* 1: 371–394, 1972.

BAYLIS, L. L., ROLLS, E. T., AND BAYLIS, G. C. Afferent connections of the caudolateral orbitofrontal cortex taste area of the primate. *Neuroscience* 64: 801–812, 1994.

CHEE-ORTS, M. N. AND OPTICAN, L. M. Cluster method for analysis of transmitted information in multivariate neuronal data. *Biol. Cybern.* 69: 29–35, 1993.

CRITCHLEY, H. D. AND ROLLS, E. T. Olfactory neuronal responses in the primate orbitofrontal cortex: analysis in an olfactory discrimination task. *J. Neurophysiol.* 75: 1659–1672, 1996a.

CRITCHLEY, H. D. AND ROLLS, E. T. Hunger and satiety modify the responses of olfactory and visual neurons in the primate orbitofrontal cortex. *J. Neurophysiol.* 75: 1673–1686, 1996b.

DAVIS, J. L. AND EICHENBAUM, H. (Editors). *Olfaction: A Model System for Computational Neuroscience.* Cambridge, MA: Bradford, 1991.

FEIGENBAUM, J. D. AND ROLLS, E. T. Allocentric and egocentric spatial information processing in the hippocampal formation of the behaving primate. *Psychobiology* 19: 21–40, 1991.

GRANGER, R., AMBROS-INGERSON, J., AND LYNCH, G. Derivation of encoding characteristics of layer II cerebral cortex. *J. Cognit. Neurosci.* 1: 61–87, 1989.

HABERLY, L. B. AND BOWER, J. M. Olfactory cortex: model circuit for study of associative memory? *Trends Neurosci.* 12: 258–264, 1989.

HERTZ, J. A., KJAER, T. W., ESKLANDER, E. N., AND RICHMOND, B. J. Measuring natural neural processing with artificial neural networks. *Int. J. Neural Syst.* 3, *Suppl.*: 91–103, 1992.

HINTON, G. E., MCCLELLAND, J. L., AND RUMELHART, D. E. Distributed representations. In: *Parallel Distributed Processing,* edited by D. E. Rumelhart and J. L. McClelland. Cambridge, MA: MIT Press, 1986, chapt. 3, p. 77–109.

HORNAK, J., ROLLS. E. T., AND WADE, D. Face and voice expression identification in patients with emotional and behavioural changes following ventral frontal lobe damage. *Neuropsychologia* 34: 247–261, 1996.

JONES-GOTMAN, M. AND ZATORRE, R. J. Olfactory identification in patients with focal cerebral excision. *Neuropsychologia* 26: 387–400, 1988.

KAUER, J. S., NEFF, S. R., HAMILTON, K. A., AND CINELLI, A. R. The salamander olfactory pathway: visualizing and modeling circuit activity. In: *Olfaction: A Model System for Computational Neuroscience,* edited by J. L. Davis and H. Eichenbaum. Cambridge, MA: Bradford, 1991, chapt. 2, p. 43–68.

LINSKER, R. Local synaptic learning rules suffice to maximize mutual information in a linear-network. *Neural Comput.* 4: 691–702, 1992.

MERRILL, E. G AND AINSWORTH, A. Glass-coated platinum-plated tungsten microelectrodes. *Med. Biol. Eng.* 10: 662–672, 1972.

MORECRAFT, R. J., GEULA, C., AND MESULAM, M.-M. Cytoarchitecture and neural afferents of orbitofrontal cortex in the brain of the monkey. *J. Comp. Neurol.* 323: 341–358, 1992.

OPTICAN, L. M., GAWNE, T. J., RICHMOND, B. J., AND JOSEPH, P. J. Unbiased measures of transmitted information and channel capacity from multivariate neuronal data. *Biol. Cybern.* 65: 305–310, 1991.

OPTICAN, L. AND RICHMOND, B. J. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J. Neurophysiol.* 57: 132–146, 1987.

PANZERI, S. AND TREVES, A. Analytical estimates of limited sampling biases in different information measures. *Network.* In press.

PRICE, J. L., CARMICHAEL, S. T., CARNES, K. M., CLUGNET M.-C., AND KURODA, M. Olfactory input to the prefrontal cortex. In: *Olfaction: A Model System for Computational Neuroscience,* edited by J. L. Davis and H. Eichenbaum. Cambridge, MA: MIT Press, 1991, p. 101–120.

ROLLS, E. T. Central taste anatomy and physiology. In: *Handbook of Olfaction and Gustation,* edited by R. L. Doty. New York: Dekker, 1995, chapt. 24, p. 549–573.

ROLLS, E. T. AND BAYLIS, L. L. Gustatory, olfactory and visual convergence within the primate orbitofrontal cortex. *J. Neurosci.* 14: 5437–5452, 1994.

ROLLS, E. T., BURTON, M. J., AND MORA, F. Hypothalamic neuronal responses associated with the sight of food. *Brain Res.* 111: 53–60, 1976.

ROLLS, E. T., CRITCHLEY, H., MASON, R., AND WAKEMAN, E. A. Orbitofrontal cortex neurons: role in olfactory and visual association learning. *J. Neurophysiol.* 75: 1870–1881, 1996.

ROLLS, E. T. AND TOVEE, M. J. The sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J. Neurophysiol.* 73: 713–726, 1995.

ROLLS, E. T. AND TREVES, A. The relative advantages of sparse versus distributed encoding for associative neuronal networks in the brain. *Network* 1: 407–421, 1990.

ROLLS, E. T., YAXLEY, S., AND SIENKIEWICZ, Z. J. Gustatory responses of single neurons in the orbitofrontal cortex of the macaque monkey. *J. Neurophysiol.* 64: 1055–1066, 1990.

SCOTT, T. R., YAXLEY, S., SIENKIEWICZ, Z. J., AND ROLLS, E. T. Gustatory responses from the frontal opercular cortex of the alert cynomolgus monkey. *J. Neurophysiol.* 56: 876–890, 1986.

TANABE, T., IIONO, M., OOSHIMA, Y., AND TAKAGI, S. F. An olfactory area in the prefrontal lobe. *Brain Res.* 80: 127–130, 1974.

TANABE, T., IIONO, M., OOSHIMA, Y., AND TAKAGI, S. F. An olfactory projection area in the orbitofrontal cortex of the monkey. *J. Neurophysiol.* 38: 1269–1283, 1975a.

TANABE, T., IIONO, M., AND TAKAGI, S. F. Discrimination of odors in olfactory bulb, pyriform-amygdaloid areas and orbitofrontal areas. *J. Neurophysiol.* 38: 1284–1296, 1975b.

TOVEE, M. J. AND ROLLS, E. T. Information encoding in short firing rate epochs by single neurons in the primate temporal visual cortex. *Visual Cognit.* 2: 35–58, 1995.

TOVEE, M. J., ROLLS, E. T., TREVES, A., AND BELLIS, R. P. Information encoding and the responses of single neurons in the primate temporal visual cortex. *J. Neurophysiol.* 70: 640–654, 1993.

TREVES, A. Dilution and sparse encoding in threshold-linear nets. *J. Phys. A Math. Gen.* 23: 1–9, 1990.

TREVES, A. AND PANZERI, S. The upward bias in measures of information derived from limited data samples. *Neural Comput.* 7: 399–407, 1995.

TREVES, A. AND ROLLS, E. T. What determines the capacity of autoassociative memories in the brain? *Network* 2: 371–397, 1991.

TREVES, A. AND ROLLS, E. T. A computational analysis of the role of the hippocampus in memory. *Hippocampus* 4: 374–391, 1994.

YARITA, H., IIONO, M., TANABE, T., KOGURE S., AND TAKAGI, S. F. A transthalamic olfactory pathway to orbitofrontal cortex in the monkey. *J. Neurophysiol.* 43: 69–85, 1980.

YAXLEY, S., ROLLS, E. T., AND SIENKIEWICZ, Z. J. Gustatory responses of single neurons in the insula of the macaque monkey. *J. Neurophysiol.* 63: 689–700, 1990.

ZATORRE, R. J. AND JONES-GOTMAN, M. Human olfactory discrimination after unilateral frontal or temporal lobectomy. *Brain* 114: 71–84, 1991.

ZATORRE, R. J., JONES-GOTMAN, M., EVANS, A. C., AND MEYER, E. Functional localization of human olfactory cortex. *Nature Lond.* 360: 339–340, 1992.