Biol Cybern (2007) 96:547–560
DOI 10.1007/s00422-007-0149-1

ORIGINAL PAPER

Biological Cybernetics

Neuronal selectivity, population sparseness, and ergodicity in the inferior temporal visual cortex

Leonardo Franco · Edmund T. Rolls · Nikolaos C. Aggelopoulos · Jose M. Jerez

Received: 12 December 2005 / Accepted: 9 March 2007 / Published online: 5 April 2007

Abstract The sparseness of the encoding of stimuli by single neurons and by populations of neurons is fundamental to understanding the efficiency and capacity of representations in the brain, and was addressed as follows. The selectivity and sparseness of firing to visual stimuli of single neurons in the primate inferior temporal visual cortex were measured to a set of 20 visual stimuli including objects and faces in macaques performing a visual fixation task. Neurons were analysed with significantly different responses to the stimuli. The firing rate distribution of 36% of the neurons was exponential. Twenty-nine percent of the neurons had too few low rates to be fitted by an exponential distribution, and were fitted by a gamma distribution. Interestingly, the raw firing rate distribution taken across all neurons fitted an exponential distribution very closely. The sparseness a^s or selectivity of the representation of the set of 20 stimuli provided by each of these neurons (which takes a maximal value of 1.0) had an average across all neurons of 0.77, indicating a rather distributed representation. The sparseness of the representation of a given stimulus by the whole population of neurons, the population sparseness a^p, also had an average value of 0.77. The similarity of the average single neuron selectivity a^s and population sparseness for any one stimulus taken at any one time a^p shows that the representation is weakly ergodic. For this to occur, the different neurons must have uncorrelated tuning profiles to the set of stimuli.

L. Franco · J. M. Jerez
Depto. de Lenguajes y Cs. de la Computacion,
Universidad de Malaga, Campus de Teatinos S/N,
29071 Malaga, Spain

E. T. Rolls (✉) · N. C. Aggelopoulos
Department of Experimental Psychology,
University of Oxford, South Parks Road, Oxford OX1 3UD, UK
e-mail: Edmund.Rolls@psy.ox.ac.uk
URL: http://www.cns.ox.ac.uk

1 Introduction

The question of how information is encoded by populations of neurons in the brain is fundamental for understanding how the brain operates. Towards the end of the primate ventral visual system, in the inferior temporal visual cortex, neurons respond with some selectivity to different faces or objects (Perrett et al. 1982; Desimone 1991; Tanaka 1996; Rolls 2000, 2005, 2007, 2008; Rolls and Deco 2002), with smaller and smaller firing rates for more and more objects, as illustrated in Fig. 1 (Rolls and Tovee 1995; Baddeley et al. 1997; Rolls et al. 1997b; Treves et al. 1999). The selectivity of neurons and the sparseness of representations in the visual system are topics of great interest in relation to whether the sparseness of the representation matches that of image statistics, and whether the encoding is efficient by utilizing only a small proportion of neurons firing for any one stimulus (Barlow 1961; Barlow et al. 1989; Atick 1992; Field 1994, 1999; Olshausen and Field 1997, 2004; Rolls et al. 1997b; Treves et al. 1999; Vogels 1999; Vinje and Gallant 2000, 2002; Rolls and Deco 2002; Lehky et al. 2005). The single neuron selectivity reflects response distributions of individual neurons across time to different stimuli. Part of the interest of measuring the firing rate probability distributions of individual neurons is that one form of the probability distribution—the exponential—maximizes the entropy of the neuronal responses for a given mean firing rate, which could be used to maximize information transmission consistent with keeping the firing rate on average low, in order to minimize metabolic expenditure (Levy and Baxter 1996; Baddeley et al. 1997). In this paper we show that, while some single neurons do fit an exponential distribution and others do not, there is a very close fit to an exponential distribution of firing rates if all spikes from all the neurons are considered together. The implication is that a neuron with inputs from the inferior

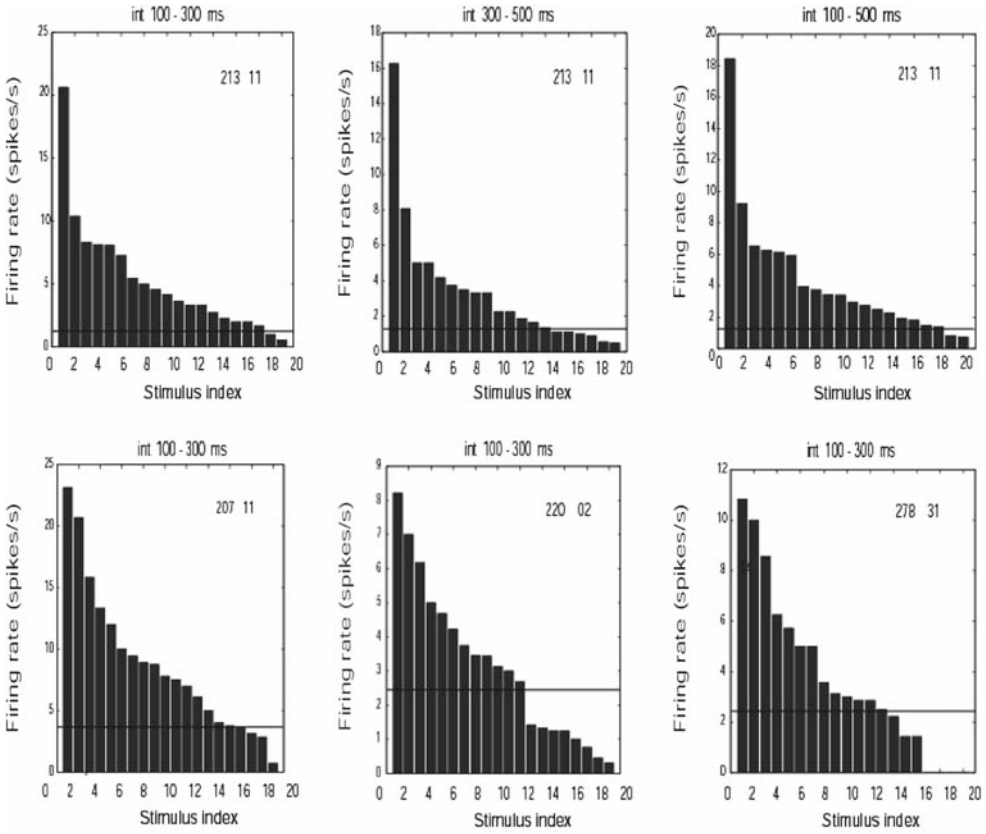

Fig. 1 Firing rates of four neurons to the set of stimuli, with the stimuli ranked according to the firing rates they produced. *Top row* The firing rates for the three intervals 100–300 ms, 300–500 ms, and 100–500 ms for one of the neurons. *Bottom row* The firing rates for three more neurons for the interval 100–300 ms. The *horizontal line* indicates the spontaneous firing rate of each neuron. The number 213 11 etc. in each panel indicates the neuron number

temporal visual cortex will receive an exponential distribution of firing rates on its afferents.

Instead, if we consider the responses of a population of neurons taken at any one time (to one stimulus), we might also expect a sparse graded distribution, with few neurons firing fast to a particular stimulus. It is important to measure the population sparseness, for this is a key parameter that influences the number of different stimuli that can be stored and retrieved in networks such as those found in the cortex with recurrent collateral connections between the excitatory neurons, which can form autoassociation or attractor networks if the synapses are associatively modifiable (Hopfield 1982; Treves and Rolls 1991; Rolls and Treves 1998; Rolls and Deco 2002). Further, in physics, if one can predict the distribution of the responses of the system at any one time (the population level) from the distribution of the responses of a component of the system across time, the system is described as ergodic, and a necessary condition for this is that the components are uncorrelated (Lehky et al. 2005). Considering this in neuronal terms, the average sparseness of a population of neurons over multiple stimulus inputs must equal the average selectivity to the stimuli of the single neurons within the population, provided that the responses of the neurons are uncorrelated (Foldiak 2003).

To address these issues of how populations of neurons encode stimuli, in this paper we describe measurements of the population sparseness. We then examine to what extent ergodicity is shown by neurons in the inferior temporal cortex, and the implications this has for whether the components of the system—the individual neurons—have uncorrelated responsiveness to a set of stimuli. The results described here are the first we know to directly address the issue of the sparseness of the population code of inferior temporal cortex neurons, and the first to directly compare the single cell and population sparsenesses, and show that weak ergodicity is apparent in the neuronal representations in the inferior temporal visual cortex. Previous findings imply that the population code is not very sparse, in that for example qualitative analysis using multidimensional scaling suggests a population code (Young and Yamane 1992); individual neurons do not have very sparse tuning to a set of stimuli (Rolls and Tovee 1995); and each neuron conveys information about many stimuli in the set of stimuli (Rolls et al. 1997b). We note that the analyses described in this paper are concerned

with the mean firing rates to each of a set of stimuli, i.e. with the response profile or tuning of a neuron to a set of stimuli. Evidence on how trial by trial variation in the firing of simultaneously recorded neurons influences the information that can be obtained from neuron populations is a very different issue that has been considered elsewhere using a partially overlapping set of the same neurons (Rolls et al. 2003, 2004, 2006; Franco et al. 2004).

2 Methods

2.1 Recording techniques

The responses of single neurons in the temporal cortical visual areas were measured to a set of 20 visual stimuli in two rhesus macaques performing a visual fixation task using experimental procedures similar except as described below to those described in detail previously (Rolls et al. 1997a). The stimuli in the standard stimulus set included $S = 20$ greyscale images of objects (7), faces (8), natural scenes (3), and geometrical stimuli (2) of the type which produce differential responses from inferior temporal cortex neurons, and examples of which have been illustrated previously (Rolls and Tovee 1995). The resolution of these images was 256 wide $\times$ 256 high with 256 grey levels. Grey level rather than color images were used as we were especially interested in form encoding rather than just color features (Rolls and Deco 2002). From 64 cells recorded simultaneously in groups of 2–4 neurons in 33 experiments, 41 neurons had significant differences in their firing rate to the set of 20 stimuli, as shown by ANOVA ($p < 0.05$). These 41 neurons were used to characterize the firing rate probability distributions described in this paper. When the sparseness of the population was measured, the data were from 29 neurons recorded in one monkey with the standard stimulus set, so that the population firing rate distribution to a given set of stimuli was being measured. When comparisons are made between the single cell and population sparsenesses described below, these were always using the same set of 29 neurons.

The neurons were selected to show responses that differed between the different stimuli (as shown by a one-way ANOVA). Usually, 20 trials for each stimulus were available. The set of stimuli were shown once in random order, a second time in a new random sequence, etc. Populations of two to nine neurons were recorded simultaneously using two to four independently movable single neuron epoxy-insulated tungsten electrodes with uninsulated tip diameters of less than 10 μm (FHC Inc., USA) using an Alpha-Omega (Israel) recording system. Typically, we were able to move the microelectrodes until two to four of the simultaneously recorded neurons responded differentially to the set of stimuli used. The microelectrodes were stereotaxically guided, and the location of the microelectrodes was reconstructed on each track using X-rays and subsequent histological reconstruction using microlesions made on selected tracks as described by Feigenbaum and Rolls (1991). The recording system (Neuralynx Inc., USA) filtered and amplified the signal and stored spike waveforms which were later sorted to ensure that the spike waveforms from each neuron in the small number of cases when there were spikes from more than one neuron on one microelectrode were clearly separated into different waveform clusters using the Datawave (USA) Discovery software. The neurophysiological methods used here have been described in detail by Booth and Rolls (1998). All procedures, including preparative and subsequent ones, were carried out in accordance with the NIH Guide for the Care and Use of Laboratory Animals and the guidelines of The Society for Neuroscience, and were licensed under the UK Animals (Scientific Procedures) Act, 1986.

Eye position was measured to an accuracy of 0.5° with the search coil technique (Judge et al. 1980), and steady fixation of a position on the monitor screen was ensured by use of a (blink version of a) visual fixation task. The timing of the task is described below. The stimuli were static visual stimuli presented at the centre of the video monitor placed at a distance of 53 cm from the eyes. A full-size face image typically subtended 21° in the visual field. The fixation spot position was at the centre of the screen. The monitor was viewed binocularly, with the whole screen visible to both eyes.

2.2 Visual fixation task

Each trial started at −500 ms (with respect to the onset of the test image) with a 500 ms warning tone to allow fixation of the fixation point, which appeared at the same time. At −100 ms the, fixation spot was blinked off so that there was no stimulus on the screen in the 100 ms period immediately preceding the test image. The screen in this period, and at all other times including the inter-stimulus interval, was set at the mean luminance of the test images. At 0 ms, the tone was switched off and the test image was switched on for 500 ms. At the termination of the test stimulus, the fixation spot reappeared, and then after a random interval in the range 150–3,350 ms it dimmed, to indicate that licking responses to a tube in front of the mouth would result in the delivery of fruit juice. The dimming period was 500 ms, and after this, the fixation spot was switched off, and reward availability terminated 500 ms later. [A diagram of the timing of this task is provided by Tovee et al. (1994) and Tovee and Rolls (1995).] The monkey was required to fixate the fixation spot in that, if he licked at any time other than when the spot was dimmed, saline instead of fruit juice was delivered from the tube; in that the dimming was by so little that it could only be detected if the monkey fixated the spot; and in that, if the

eyes moved by more than 0.5° from time 0 until the start of the dimming period, then the trial was aborted. (When a trial aborted, a high frequency tone sounded for 0.5 s, no reinforcement was available for that trial, and the inter-trial interval was lengthened from 8 to 11 s.)

2.3 Profile of responses to the set of stimuli

For each neuron, the firing rate to each stimulus was averaged across trials in selected epochs, as described in the results to produce the response profile to the set of stimuli. It was these mean rates to each stimulus that were used for the majority of the analyses described in this paper, except where as stated the individual rates from every trial were used to generate a distribution.

2.4 Selectivity or sparseness of the tuning of an individual neuron to a set of stimuli

The selectivity or sparseness a^s of the representation of a set of stimuli of an individual neuron can be measured as

$$a^s = \frac{\left(\sum_{s=1}^{S} \left(y_s/S\right)\right)^2}{\sum_{s=1}^{S} (y_s)^2/S},$$

where y_s is the mean firing rate of the neuron to stimulus s in the set of S stimuli. (S is the number of stimuli s in the set.)

The selectivity or sparseness a^s 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 differently from the spontaneous rate), then a value of 0.2 would indicate that 20% of the stimuli produced high firing rates, and 80% had 0 firing rates. In the more general case of a continuous distribution of firing rates, the selectivity measure—a^s,—still provides a quantitative measure of the length of the tail of the firing rate distribution (Treves and Rolls 1991). Although a^s has been named the single cell sparseness or selectivity, we note that because high values of a^s indicate broad tuning of a neuron, the measure might also be called the single cell breadth of tuning (Foldiak 2003). One advantage is that it can be applied to neurons which 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 second 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). Third, it makes no assumption about the mean and the variance of the firing rate. Lehky et al. (2005) are incorrect in their surmize on this. Because it is the ratio of the terms $(\Sigma_{s=1,S}(y_s/S))^2$ and $\Sigma_{s=1,S}(y_s^2/S)$ that is used to calculate a^s, the scaling terms such as that for different mean firing rates cancel, and Lehky et al. (2005) formulated their argument in terms of a difference of these terms, which is not how a^s is defined. Fourth, 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 to all the stimuli in a set of stimuli with the same mean rate. We preferred this measure of selectivity to kurtosis, as kurtosis may be inappropriate for asymmetrical distributions (Olshausen and Field 2004), which the probability distributions are, as shown in Figs. 1, 2, and 3. We also preferred this measure to the entropy measure of Lehky et al. (2005), which requires large numbers of stimuli and measures of the neuronal response to each stimulus, or very careful correction of the type developed for measures of information representation by Treves and Panzeri (1995) and Panzeri and Treves (1996) (see also Vinje and Gallant 2000). We checked whether the calculated value of a^s can be estimated accurately with the number of stimuli used, and showed by Monte Carlo simulations that the estimates given 20 stimuli for every neuron for the neuron sparseness with exponentially distributed firing rates for which the true value is 0.50 is quite close with 20 stimuli (which gave an estimate of a^s of 0.54). A similar argument applies to the population sparseness desribed below, for which 29 and 41 neurons give estimated values of 0.53 and 0.52, respectively.

2.5 Sparseness of the population code

The sparseness a^p of the population code may be quantified (for any one stimulus) as

$$a^p = \frac{\left(\sum_{n=1}^{N} \left(y_n/N\right)\right)^2}{\sum_{n=1}^{N} (y_n)^2/N},$$

where y_n is the mean firing rate of neuron n in the set of N neurons.

This measure, a^p, of the sparseness of the representation of stimulus by a population of neurons has a number of advantages. One is that it is the same measure of sparseness which has proved to be useful and tractable in formal analyses of the capacity of associative neural networks and the interference between stimuli that use an approach derived from theoretical physics (Rolls and Treves 1990, 1998; Treves 1990; Treves and Rolls 1991). We note that high values of a^p indicate broad tuning of the population, and that low values of a^p indicate sparse population encoding. This is a measure of the extent of the tail of the distribution, in this case of the firing rates of a population of neurons to a stimulus. A low value indicates

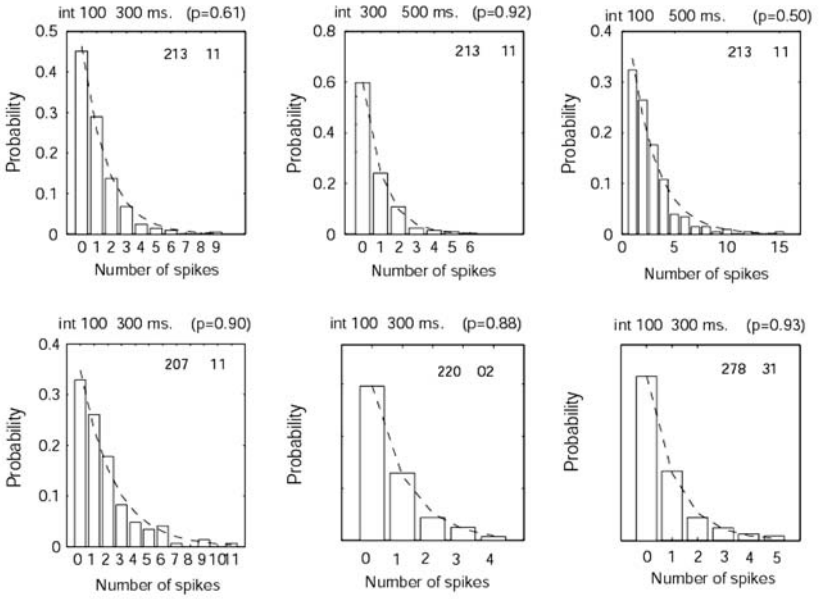

Fig. 2 Firing rate probability distributions for four neurons with a good fit to an exponential firing rate distribution. The neurons are the same as those shown in Fig. 1. *Top row* The distributions for the three intervals 100–300 ms, 300–500 ms, and 100–500 ms for one of the neurons. *Bottom row* The distributions for three more neurons for the interval 100–300 ms. The P value shows the Chi-square measure of the goodness of fit to an exponential firing rate distribution

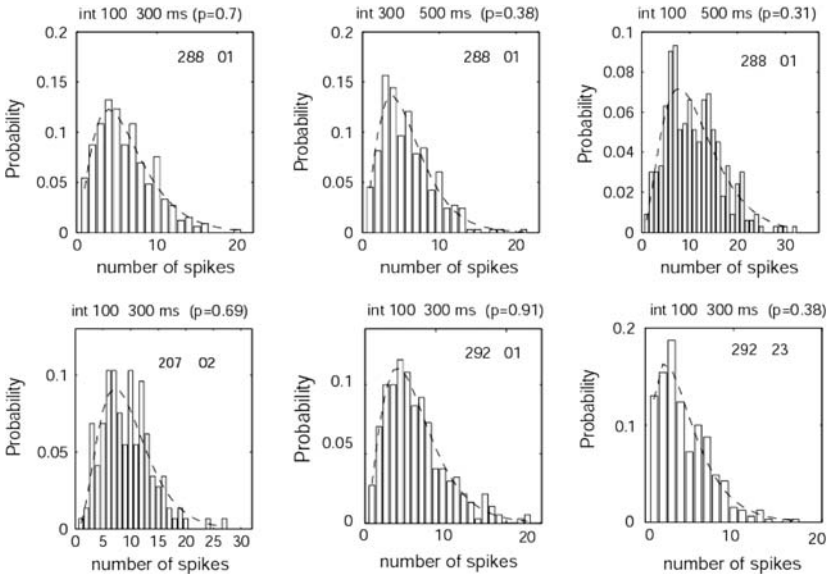

Fig. 3 Firing rate probability distributions for four neurons with a poor fit to an exponential firing rate distribution, but fitted by a gamma distribution. *Top row* The distributions for the three intervals 100–300 ms, 300–500 ms, and 100–500 ms for one of the neurons. *Bottom row* The distributions for three more neurons for the interval 100–300 ms. The P value shows the Chi-square measure of the goodness of fit to a gamma firing rate distribution

that there is a long tail to the distribution, equivalent in this case to only a few neurons with high firing rates. If these neurons were binary (either responding with a high firing rate, or not responding differently from the spontaneous rate), then a value of 0.1 would indicate that 10% of the neurons produced high firing rates, and 90% had 0 firing rates, to a given stimulus. In the more general case of a continuous distribution of firing rates, the selectivity measure—a^{p}—still provides a quantitative measure of the length of the tail of the firing rate distribution. Although a^{p} can be thought of as the population sparseness, we note that because high values of a^{p} indicate broad tuning of the population of neurons, the measure might also be called the population breadth of tuning to a stimulus (or averaged across stimuli).

The measures a^{s} and a^{p} have an analogous formulation, but measure different quantities. We explore the relation between these measures in this paper.

3 Results

From 64 cells recorded simultaneously in groups of 2–4 neurons in 33 experiments, we performed the analyses on 41 neurons that had significant differences in their firing rate to the set of 20 stimuli, as shown by ANOVA ($P < 0.05$). There were typically 20 trials of data available for each stimulus for each neuron. Examples of peristimulus time rastergrams and

histograms for this type of neuron have been shown previously in Fig. 3 of Tovee et al. (1993) and are therefore not reproduced here. Most of the neurons had their best responses to objects, and such neurons sometimes responded to one or two faces in the set.

3.1 Firing rate probability distributions

Firing rate probability distributions based on the mean firing rate to each of the 20 stimuli were calculated for three different post-stimulus time intervals: 100–300 ms, 300–500 ms, and 100–500 ms. The first time interval—100–300 ms,—is the most relevant interval in terms of perceptual processing and the use by other brain areas of the information, in that the onset latency of the neuronal responses to visual stimuli in the inferior temporal visual cortex is in the order of 80–100 ms (Rolls 1984), and correct behavioural responses to visual stimuli can be made in a visual discrimination task with latencies of 350–450 ms (Rolls et al. 1979).

Figure 1 shows the firing rates of four neurons to the set of stimuli. Figure 2 shows the firing rate probability distributions for examples of neurons in which the distribution was closely fitted by an exponential. (The neurons correspond to those shown in Fig. 1.) The distributions shown are for the interval 100–300 ms, but similarly good fits to the exponential were found for these neurons for the other time intervals, as illustrated for one of the neurons. The accuracy of the fits to an exponential was determined using ($P(r) = A \times \exp(B \times r)$) where r is the firing rate and then performing a Chi-square test, and the probability values of a good fit are shown.

Figure 3 shows the firing rate probability distributions for examples of neurons with a poor fit to an exponential firing rate distribution, but fitted by a gamma distribution. The distributions shown are for the interval 100–300 ms, but similarly good fits to the gamma distribution were found for these neurons for the other time intervals, as illustrated for one of the neurons. (The parameters of the gamma distribution were similar for the different time intervals.)

Of the 41 neurons in the dataset, 15 were not rejected (taking $P < 0.05$ as a poor fit) as a good fit to the exponential (i.e. they had exponential distributions), and 12 did not fit an exponential but did fit a gamma distribution (cf. Treves et al. 1999; see Table 1). The remaining 14 neurons could be fitted by a gamma distribution if a third parameter to allow the location of the distribution to shift from zero was included. For the neurons with an exponential distribution, the mean firing rate across the stimulus set was 5.7 spikes/s, and for the neurons with a gamma distribution was 21.1 spikes/s ($t = 4.5$, df $= 25$, $P < 0.001$). It may be that neurons with high mean rates to a stimulus set tend to have few low rates ever, and this accounts for their poor fit to an exponential firing rate probability distribution, which fits when there are many low firing rate values in the distribution.

Table 1 Firing rates and sparseness values for each neuron

Cell	Spontaneous	Mean	Minimum	Maximum	Fit	Sparseness
bj185.01	7.5	33.0	15.0	50.8	2	0.94
bj185.03	21.5	28.2	8.3	89.4	2	0.72
bj207.02	0.0	41.1	24.4	64.5	2	0.94
bj207.11	3.8	8.7	0.7	23.1	1	0.69
bj213.01	10.6	16.6	8.8	48.1	0	0.82
bj213.11	1.3	5.2	0.6	20.6	1	0.58
bj215.11	4.9	7.8	2.8	15.0	2	0.83
bj215.12	11.2	14.9	6.0	19.0	2	0.92
bj220.01	6.8	12.8	5.4	23.5	0	0.88
bj220.02	2.4	3.1	0.3	8.2	1	0.67
bj220.03	17.8	22.2	11.5	33.4	0	0.93
bj220.11	3.1	5.9	0.0	15.0	0	0.63
bj226.02	0.0	10.6	6.7	21.9	2	0.89
bj229.03	7.5	7.4	2.5	15.4	1	0.82
bj229.21	1.6	2.2	0.0	6.3	1	0.68
bj239.11	8.1	5.1	2.8	7.5	0	0.91
bj239.12	3.9	1.7	0.0	5.3	1	0.66
bj643.01	3.9	6.5	1.9	14.0	0	0.81
bj643.21	0.8	1.7	0.0	6.1	1	0.52
bj244.02	10.5	9.7	3.8	20.3	0	0.81
bj244.11	20.7	19.3	12.8	27.6	1	0.97
bj278.31	2.3	3.7	0.0	10.8	1	0.58
bj280.32	1.4	18.8	6.5	63.8	0	0.68
bj280.01	11.6	4.5	2.0	7.8	0	0.92
bj287.01	0.0	10.9	0.0	22.2	0	0.84
bj288.01	8.3	25.2	6.0	47.6	2	0.84
bj690.02	3.4	4.9	0.0	9.2	0	0.80
bj690.21	6.7	6.7	0.0	16.5	1	0.72
bj690.22	10.0	13.4	0.0	18.3	1	0.82
bj291.11	0.0	13.8	8.4	20.0	0	0.95
bj291.24	13.5	13.1	2.3	30.0	2	0.68
bj292.01	25.0	30.9	19.7	57.3	2	0.92
bj292.02	0.8	0.8	0.0	3.1	1	0.53
bj292.21	1.8	2.8	0.3	7.0	1	0.75
bj292.23	17.4	17.2	6.7	38.7	2	0.81
bj293.02	3.7	14.3	3.0	38.0	2	0.67
bj293.11	1.2	2.3	0.0	9.0	1	0.50
bj293.12	4.1	5.7	1.7	13.5	1	0.80
bj293.21	4.6	5.9	1.5	15.0	2	0.78
bs46.01	2.0	23.2	12.5	40.0	0	0.91
bs56.01	0.9	1.9	0.7	3.3	0	0.89
Average	6.5	11.8	4.5	24.5		0.77

Fit: 0 no fit, 1 exponential fit, 2 gamma fit

The spontaneous firing rate was calculated from the firing in the period 0–99 ms after stimulus onset, before the neuron had started to respond to the visual stimuli. The other columns show the mean firing rate across stimuli, and the minimum and maximum firing rate to any stimulus, in spikes/s calculated in the period 100–300 ms after visual stimulus onset. In this 100–300 ms period, the firing rate could be lower than the preceding spontaneous firing rate. The sparseness value is the single cell selectivity a^s

The probability distribution was generated from the firing rate of a neuron to a set of stimuli typical of the stimuli that activate neurons in the inferior temporal visual cortex. The set in a sense was intended to represent a natural set of stimuli for the neurons. The approach is effective as used, in that similar distributions were found when we showed videos of natural scenes to macaques, and measured the firing rate in short temporal windows (e.g. 100 ms) while the monkey watched the video (Baddeley et al. 1997; Treves et al. 1999). Moreover, we were able to confirm that the population of neurons analysed was well tuned to this set of stimuli, in that the neurons with an exponential firing rate distribution conveyed 1.12 bits (± 0.16 SEM calculated over 15 neurons) of information about the most effective stimulus in the set (i.e. the surprise, calculated as described by Rolls et al. 1997b), and the neurons with a gamma firing rate distribution conveyed 1.75 bits (± 0.18 SEM calculated over 12 neurons) of information about the most effective stimulus in the set. Thus both subsets of neurons were well tuned to the stimulus set.

Interestingly, if the probability distribution was calculated for the firing rate counts from all 41 neurons across all stimuli, then the fit to an exponential firing rate distribution was very accurate, as shown in Fig. 4a (Chi-square test, P value = 0.85, time interval 100–300 ms). The raw spike counts, with no normalization or scaling, were used in this composite distribution (as also was the case for the distributions for individual neurons shown in Figs. 2, 3). (Figure 4a thus includes all mean spike counts for each stimulus/neuron combination in the dataset.) The data in Fig. 4a are for the time windows: 100–300 ms, but similar close fits were found for the two other time windows: 100–500 ms ($P = 0.70$) and 300–500 ms ($P = 0.77$). The result shown in Fig. 4a is not trivial, for if the data from each cell was first scaled to the same mean rate, and then a composite distribution was calculated, this was far from exponential as shown in Fig. 4b. Given the result shown in Fig. 4a, we performed a check with an independent data set of inferior temporal cortex neurons recorded while a macaque performed a visual discrimination task in which one of two objects shown in a complex or blank scene had to be touched to obtain a juice reward (Aggelopoulos et al. 2005). We found a similarly exact fit of the raw (unscaled) firing rate probability distribution to an exponential distribution.

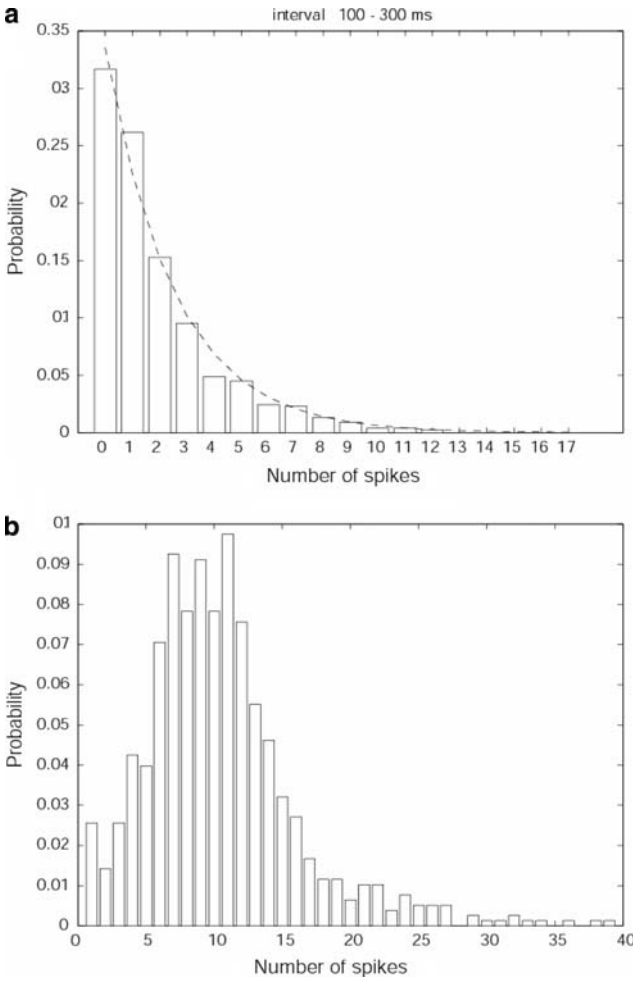

Fig. 4 **a** The population firing rate probability distributions for the 100–300 ms interval formed by adding the spike counts from all 41 neurons, and across all stimuli. The fit to the exponential distribution was high. **b** The population normalized firing rate probability distributions for the 100–300 ms interval formed by adding the spike counts from all 41 neurons, and across all stimuli. The firing rate for each cell was normalized to the same mean before the probability distribution was calculated

3.2 Selectivity of individual cells

The selectivity or sparseness a^s of the representation was calculated across stimuli of individual neurons as described in the methods. The mean single cell selectivity or sparseness is 0.77 ± 0.13 (mean ± SD) for the interval 100–300 ms. (For the interval 100–500 ms, the sparseness a^s was 0.79 ± 0.14, and for 300–500 ms was 0.74 ± 0.18.) [The sparseness a^s of the cells with a fit to an exponential distribution for the 100–300 ms time window was 0.70, and for the neurons that fitted a gamma distribution was 0.83 ($t = 3.09$, df = 25, $P = 0.05$).] These sparsenesses were calculated for the 29 neurons tested with the standard stimulus set so that the values could be compared with the population sparsenesses described next. However, the values were essentially identical for the set of 41 neurons.

3.3 Population sparseness

The distribution of raw (i.e. unscaled, and without any spontaneous value subtracted) firing rates of the population of 29 neurons to one of the stimuli (stimulus 14) is shown is

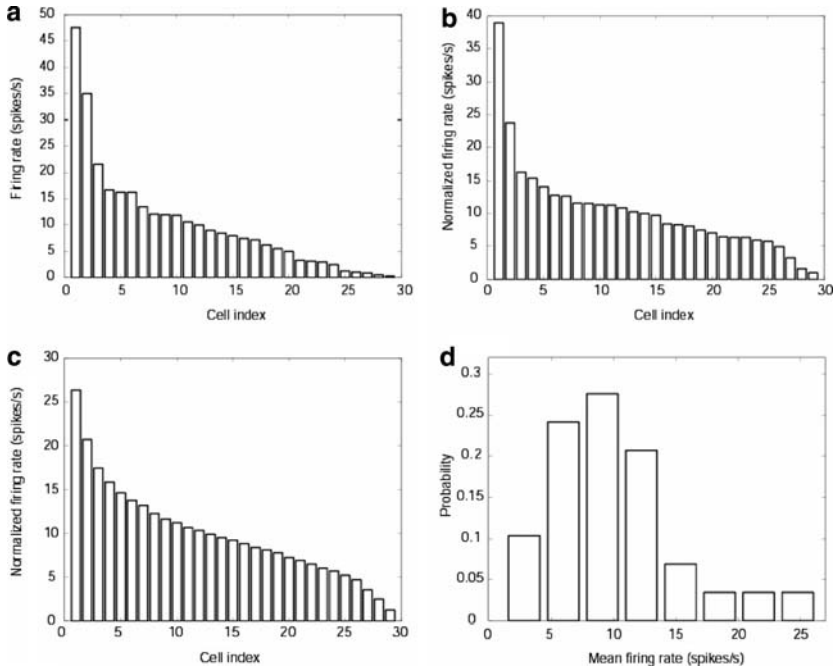

Fig. 5 **a** The (unscaled, raw) firing rates of the population of 29 cells to one of the stimuli, stimulus number 14. **b** The firing rates of the population of 29 cells to the same stimulus as in (a), but with the response for each neuron scaled relative to its average firing rate across the set of stimuli, which was taken as 10 Hz. **c** The firing rates of the population of neurons to any one stimulus. The rates of each neuron were normalized to the same average value of 10, then for each stimulus, the cell firing rates were placed in rank order, and then the mean firing rates of the first ranked cell, second ranked cell, etc. were taken. The graph thus shows how, for any one stimulus picked at random, the expected normalized firing rates of the population of neurons. **d** The population normalized firing rate probability distributions for any one stimulus. This was computed effectively by taking the probability density function of the data shown in Fig. 5c

Fig. 5a. (These were the 29 neurons in the set of 41 neurons that were tested with the identical stimulus set.)

Because each neuron has its own range of firing rates (some may have peak rates of 20 spikes/s, others of 30 spikes/s, and their spontaneous firing rates may be different, as shown in Table 1), what is shown in Fig. 5a needs further analysis. It is useful to normalize the firing rates of each neuron when considering the population code for any stimulus, because otherwise a population of neurons all with the same profile of responses to the set of stimuli would have a population sparseness for any one stimulus that could take any value, depending on the exact raw firing rates of each neuron. That is, for any stimulus, each neuron in the population could have its own level of firing, and thus the population vector would appear to have firing rates that would have nothing to do with encoding the stimulus. To clarify this point further, assume that we can characterize the tuning profile of for example a V1 neuron responding to oriented bars by an angle (preferred orientation) and mean (or peak) firing rate. In order to analyse how the population of neurons encodes a set of stimuli it is more relevant to analyse the distribution of the preferred orientations of the neurons rather than using the firing rate of the neuron. If we then consider a cortical column where all neurons were tuned to the same preferred orientation but where neurons have different firing rates to the preferred stimulus, and then compute the population sparseness, a value will be obtained just depending on the distribution of the firing rate values, even if all the neurons have the same tuning profile. The analogy for inferior temporal cortex neurons is with respect to the tuning profile of each neuron measured across the set of stimuli.

For this reason, some normalization, to allow the neurons to be treated equally even though each may have its own characteristic range of firing, is needed. We therefore normalized the firing rates of each neuron with the range of rates scaled for each neuron so that the mean rate for each neuron was the same value. Which mean rate value is chosen does not matter for the calculation of the sparseness, as long as some normalization is performed. We chose to normalize so that the mean firing rate of each neuron across the range of stimuli was 10 spikes/s. This was chosen instead of scaling so that each neuron had the same scaled spontaneous rate, because some of the neurons had no spontaneous or very low spontaneous firing, which would greatly distort the scaling. We also did not scale so that each neuron had the same maximum and minimum firing rate, because this could be affected by using a restricted set of stimuli, whereas setting to the same average rate takes into account the firing of the neurons to all

stimuli in the set, and is thus more robust. It was also inappropriate to attempt to scale the responses of the neurons (i.e. the firing rate minus the spontaneous firing rate), as this could result in negative values which could lead to ambiguous results when the sparseness is calculated. Moreover, we note that a neuron receives the actual firing rates, not some calculated response that artificially subtracts a spontaneous firing rate, so that use of firing rates to a stimulus, and not the change of firing rate, i.e. the response to a stimulus, is appropriate. For these reasons scaling of the firing rates of each neuron to the same mean rate of 10 spikes/s was chosen for the population analyses.

We therefore show in Fig. 5b the firing rate of the neurons to the same stimulus, with the range of rates scaled for each neuron so that the mean rate for each neuron was 10 spikes/s. It is now possible to see that for this stimulus, each neuron fires to a different extent above or below its mean rate, and this is a population code for that stimulus.

The firing rates of the population of (29) neurons to any one stimulus are shown in Fig. 5c. The rates of each neuron were scaled so that each neuron has the same mean rate of 10, then for each stimulus, the cell firing rates were placed in rank order, and then the mean firing rates of the first ranked cell, second ranked cell, etc. were taken. This is what is shown in Fig. 5c. The graph thus shows, for any one stimulus picked at random, the expected normalized firing rates of the population of neurons. The population sparseness a^{p} of this normalized (i.e. scaled) set of firing rates is 0.77. This is effectively the sparseness of the idealized set of firing rates of the population of neurons to any one stimulus. It is very close to the average value of the population sparseness calculated by taking the mean of the population sparsenesses calculated for each stimulus (which has the value 0.77, as described below.)

Figure 5d shows the probability distribution of the normalized firing rates of the population of (29) neurons to any stimulus from the set. This was calculated by taking the probability distribution of the data shown in Fig. 5c.

The probability distribution of firing rates of the population of neurons to any one stimulus that is shown in Fig. 5d is effectively also a normalized version of what is shown in Fig. 4b, which is the probability distribution of firing rates of all the neurons to all the stimuli. The reason is that the data used to generate the population firing rate distribution to any one stimulus shown in Fig. 5d is generated by averaging across all stimuli, and this contains the normalized firing rates of every neuron to every stimulus. The reason for making this comparison is because it is necessary to normalize (i.e. scale) the firing rates when considering how a population encodes a stimulus (so that differences in the firing rates of each neuron can be factored out); yet, we see from Fig. 4a that the population firing rate distributions averaged across all stimuli do in fact follow in their unscaled form an exponential distribution, allowing the population encoding

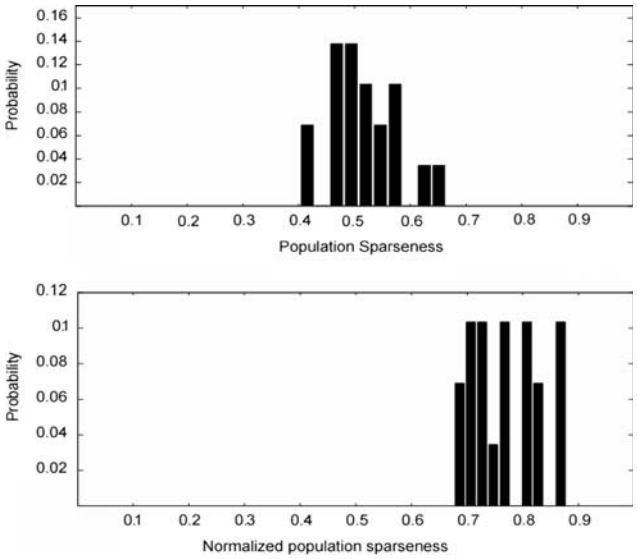

Fig. 6 The sparseness a^{p} of the population code for the 20 different stimuli. This is the probability distribution of the sparseness for each of the 20 stimuli using the 29 cells tested with the same 20 stimuli. **a** Calculated from the raw, normalized, firing rates. **b** The sparseness calculated after the mean firing rate of each neuron to the set of 20 stimuli was normalized to the same mean

to be related to energy efficient arguments about exponential firing rate distributions. We emphasize this point because of its potentially important implications: in the unscaled form, the probability distribution of the firing rates of a population of neurons is on average for a given stimulus very close to an exponential distribution.

The sparseness a^{p} of the population code was calculated as described in the Sect. 2. The mean (across the set of stimuli) population cell sparseness when computed with normalized firing for each neuron was 0.77 ± 0.06 (mean $\pm$ SD) (and when calculated without scaling was 0.52 ± 0.06). The distribution of population sparseness values for the different stimuli in the set (calculated for the 100–300 ms window) is shown in Fig. 6 (top: raw firing rates without normalization; bottom with the sparseness calculated with normalized firing rates, i.e. with the mean rates of each neuron in response to the 20 stimuli normalized to the same mean before the sparseness to any one stimulus of the population was calculated). For the interval 100–500 ms, the population sparseness a^{p} was 0.79 ± 0.14 (mean $\pm$ SD) (and when calculated without scaling was 0.51 ± 0.05).

The single cell and population sparseness values given above were for a set of 29 neurons from among the 41 neurons that were all tested with the identical set of 20 stimuli. To take advantage of the full size of this dataset of 41 neurons, we also calculated the population sparseness for the 41 neurons, now ranking for each cell its responses from the highest to the lowest for whichever stimulus set it was tested on. (All

the sets of stimuli were statistically similar, in that they all contained a set of faces and non-face images that can be effective for different neurons in the inferior temporal visual cortex. This procedure can thus be seen as a check that the results are not due to the particular set of 20 stimuli in the standard set used for 29 of the 41 neurons, but may be generalized to other exemplars of typically effective stimuli for inferior temporal cortex neurons.) For the 41 neurons, the population sparseness a^{p} was 0.79, very similar to the value of 0.77 for the set of 29 neurons tested with the standard stimulus set. For direct comparison, the mean of the single cell sparsenesses a^{s} was 0.78 ± 0.13, also very close to the mean of the single cell sparseness for the 29 neurons of 0.77. These values are for the 100–300 ms time interval. For the 100–500 ms time interval, for the 41 neurons the population sparseness was 0.81, and the mean of the single cell sparsenesses a^{s} was 0.80 ± 0.13. For the 300–500 ms time interval, for the 41 neurons the population sparseness was 0.76, and the mean of the single cell sparsenesses a^{s} was 0.75 ± 0.16. Thus for the large dataset of 41 neurons, the single cell and population sparseness measures were close to each other, as they also were for the 29 neurons tested with the standard set of stimuli. This extension of the analysis to the larger number of neurons (41) helps to establish that the sample was sufficiently large to obtain very similar values when the sample size was increased, and that the results hold for different stimulus sets containing stimuli of the type to which inferior temporal cortex neurons are tuned.

We note that weak ergodicity necessarily occurs if a^{s} and a^{p} are the same and the neurons are uncorrelated, i.e. each neuron is independently tuned to the set of stimuli (Lehky et al. 2005; see Sect. 4). Given that the values of a^{s} and a^{p} are very similar for this population of neurons, an implication is that this could arise if the response profiles of the neurons are uncorrelated. We tested this in two ways. In a first test, we measured whether the response profiles of pairs of neurons to the set of 20 stimuli were uncorrelated. We found that the mean (Pearson) correlation computed over the 406 neuron pairs was low, 0.049 ± 0.013 (SEM). In a second test, we computed how the multiple cell information available from these neurons about which stimulus was shown increased as the number of neurons in the sample was increased, using the method described previously (Rolls et al. 1997a; Franco et al. 2004). It was found, as shown in Fig. 7, that the information increased approximately linearly with the number of neurons in the ensemble. The implication is that the neurons convey independent (non-redundant) information, and this would be expected to occur if the response profiles of the neurons to the stimuli are uncorrelated. As shown in Sect. 4, this evidence that the response profiles of the neurons are uncorrelated can be taken as a contributing factor to the finding that the responses of this population of neurons are weakly ergodic.

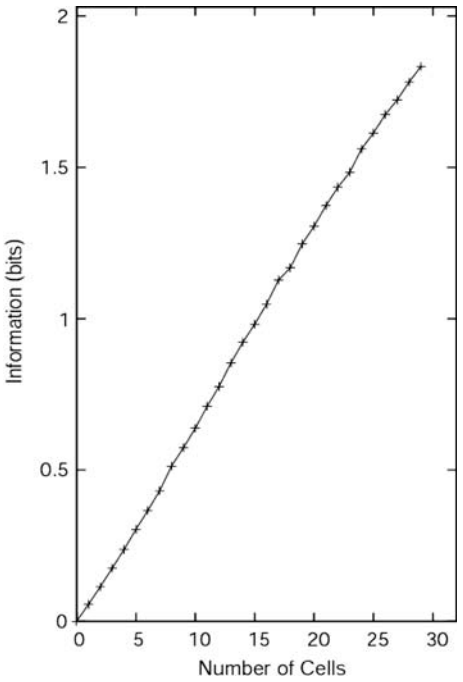

Fig. 7 Multiple cell information available in a 200 ms window starting 100 ms after stimulus onset about which of 20 stimuli was shown as a function of the number of neurons in the ensemble

The recording sites of the neurons analysed in this paper in the cortex in the inferior temporal visual cortex are shown in Fig. 8; 0 mm with respect to the sphenoid corresponds approximately to the antero-posterior level of the anterior commissure, and is approximately 18 mm anterior to the auditory meatus.

4 Discussion

We found that the firing rate probability distribution of some of these inferior temporal cortex neurons was closely exponential (see Fig. 2). For other neurons, the fit to an exponential was poor, this was associated with a higher mean firing rate (which may tend to take the distribution away from exponential because there are few very low firing rate counts), and the distributions could be fitted by a gamma function (see Fig. 3).

What was remarkable about the firing rate probability distribution was that when all the firing rate counts were collected together across all stimuli and across all neurons, the resulting firing rate probability distributions were very close to exponential, as shown in Fig. 4a. The implication is that across a large number of neurons in a given brain area (at least the inferior temporal visual cortex), the firing rate probability density function (pdf) in any short time window from all the neurons has a rather strict form, that of an exponential distribution. This is an interesting new finding that is

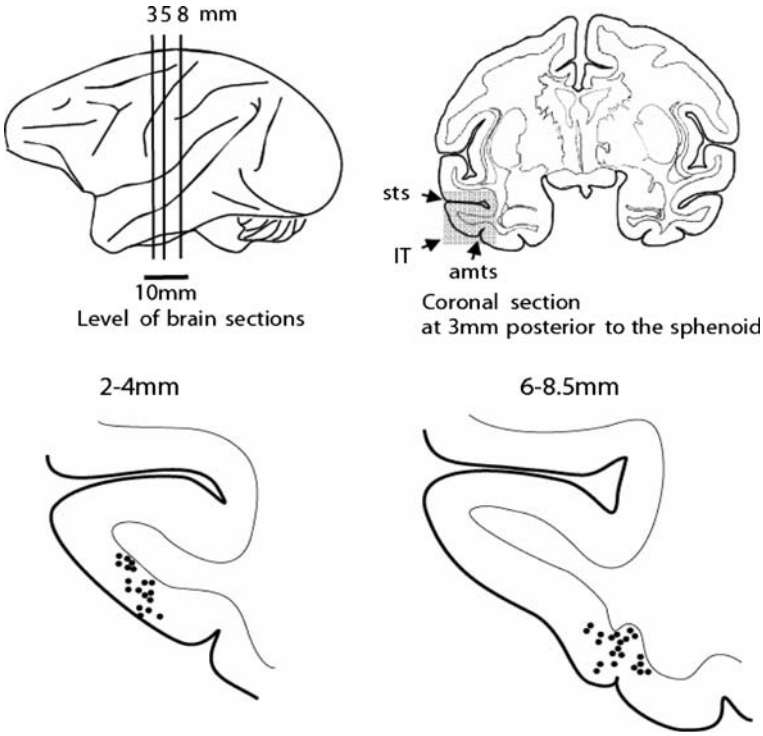

Fig. 8 The recording sites shown on coronal sections of the neurons included in this study. The positions of the coronal sections are shown on a lateral view of the macaque brain. The distances refer to mm posterior (P) to the sphenoid reference plane (see text). STS superior temporal sulcus, IT inferior temporal cortex

reported here, and its implications include the following. At the level of single neurons, an exponential pdf is consistent with minimizing energy utilization and maximizing information transmission, for a given mean firing rate (Levy and Baxter 1996; Baddeley et al. 1997). An additional implication is that a neuron that receives inputs from this population of neurons will see a firing rate distribution on its afferents that is closely exponential, and this is therefore the type of input that needs to be considered in theoretical models of neuronal network function in the brain (Rolls and Deco 2002).

The selectivity of individual cells for the set of stimuli, or single cell sparseness a^s, had a mean value of 0.77. This is close to a previously measured estimate, 0.65, which was obtained with a larger stimulus set of 68 stimuli (Rolls and Tovee 1995). Thus, the single neuron pdfs (see Figs. 2, 3) do not produce very sparse representations. Therefore, the goal of the computations in the inferior temporal visual cortex may not be to produce sparse representations [as has been proposed for V1 (Field 1994; Vinje and Gallant 2000)]. Instead one of the goals of the computations in the inferior temporal visual cortex may be to compute invariant representations of objects and faces (Rolls 2000; Rolls and Deco 2002).

The fact that the single cell sparseness values a^s found here, and the firing rate distributions, are similar to those measured previously in the same brain region with a larger set of 68 stimuli of the type found to activate inferior temporal visual cortex neurons (Rolls and Tovee 1995) indicates that the set of 20 stimuli used here is sufficiently large and representative to provide good estimates of the firing rate pdfs and sparseness a^s of the neurons analysed in this paper. Further evidence that the results described here are representative of what happens with a large set of natural stimuli is that the firing rate pdf takes approximately the same form when the firing rates of inferior temporal visual cortex neurons are measured when macaques watch videos lasting several minutes of the natural scenes that they normally encounter (Baddeley et al. 1997). However, we note that the exact sparseness value obtained is likely to be influenced by the exact stimulus set used and by the criteria used to include neurons in the study; that the distributions described in this paper are similar to those with larger stimulus sets representative of what is normally seen by monkeys (Rolls and Tovee 1995), and to those with the natural video images (Baddeley et al. 1997); and that similar distributions were obtained with the standard and other stimulus sets used in this paper. We note that part of the value of the present results is that we directly compare the single cell and population sparseness measures on an identical stimulus set, with identical neuronal populations.

The major issue of the sparseness of the representation provided by populations of neurons is also directly addressed by this investigation. Figure 5c shows the firing rates of the population of 29 cells to a typical stimulus (averaged across the set of stimuli). The population firing rate probability distribution for this average stimulus is shown in Fig. 5d. The implication of this probability distribution is that to represent any one stimulus, a few neurons will have relatively high activity, more will have moderate activity, and most

will have little activity. The population sparseness values a^{p} for each stimulus are shown in Fig. 6, and have a mean value (across stimuli) of 0.77 with normalization of each neuron to the same mean firing rate (and of 0.52 without normalization). It is this (normalized) population sparseness which is important for considering the storage capacity of associative neural networks, and interference between patterns stored in these networks (Rolls and Treves 1990; Treves 1990; Treves and Rolls 1991; Rolls and Treves 1998). (The normalized form is appropriate in that otherwise each neuron would contribute unequally to the information stored in the associative memory. The normalization could be performed by the synaptic weights.) This value of 0.77 for the population sparseness is not low, and this relatively non-sparse population representation indicates that a very large number of separate memories may not be maintainable by the recurrent collateral connections between pyramidal cells if they are associatively modifiable. On the other hand, the information that could be made available to perform for example constraint satisfaction between neurons in a small cortical area would be large due to this high value of the population sparseness a^{p} (i.e. the non-sparse representation) (see Rolls and Treves 1998; Rolls and Deco 2002), and high information transmission, rather than storing a large number of patterns, could be a useful property of this type of encoding found in the inferior temporal visual cortex.

We now consider the relation between the selectivity or sparseness of the representation of different stimuli by a single neuron—a^{s}—and the population sparseness a^{p}. They are not necessarily very similar. [For example, if each of the single neuron selectivities is high (i.e. a low value of a^{s}), but each neuron is tuned to respond to have the same profile of responsiveness across the set of stimuli, then the population firing rate distribution for any one stimulus will be flat, and the value for a^{p} will be 1.0. If the firing rates of the neurons were not normalized to the same average firing rate, then a^{p} (and the population firing rate distribution) could take any value depending on how different the average firing rates of the different neurons were.] However, we found that the sparseness a^{p} of the (normalized) firing rate distribution for the population of neurons firing to any one stimulus (see Fig. 6b) has an average value (across stimuli) of 0.77 ± 0.06 ($\pm$SD) that is essentially the same value as the single cell selectivity or sparseness $a^{\mathrm{s}} = 0.77 \pm 0.13$.

The single neuron selectivity, a^{s}, reflects response distributions of individual neurons across time. [Willmore and Tolhurst (2001) described this as "lifetime sparseness".] The population sparseness a^{p} reflects response distributions across all neurons in a population measured simultaneously. The similarity of the average values a^{p} and a^{s} (both 0.77) indicates, we believe for the first time experimentally, that the representation (at least in the inferior temporal cortex) is ergodic. The representation is ergodic in the sense of statistical physics, where the average of a single component (in this context a single neuron) across time is compared with the average of an ensemble of components at one time (cf. Masuda and Aihara 2003; Lehky et al. 2005). In comparing the neuronal selectivities a^{s} and population sparsenesses a^{p}, we formed a table in which the columns represent different neurons and the stimuli different rows (Foldiak 2003). We are interested in the probability distribution functions (and not just their summary values a^{s}, and a^{p}), of the columns (which represent the individual neuron selectivities) and the rows (which represent the population tuning to any one stimulus). We could call the system *strongly ergodic* (cf. Lehky et al. 2005) if the selectivity (probability density or distribution function) of each individual neuron is the same as the average population sparseness (pdf). (Each neuron would be tuned to different stimuli, but have the same shape of the pdf.) We have seen that this is not the case, in that the probability distribution functions of different neurons are different (Figs. 1, 2). We can call the system *weakly ergodic* if individual neurons have different selectivities (i.e. different response pdfs), but the average selectivity (measured in our case by $\langle a^{\mathrm{s}} \rangle$) is the same as the average population sparseness (measured by $\langle a^{\mathrm{p}} \rangle$), where $\langle \cdots \rangle$ indicates the ensemble average. We have seen that the neuron selectivity pdfs are different (Figs. 1, 2), but that their average $\langle a^{\mathrm{s}} \rangle$ is the same as the average $\langle a^{\mathrm{p}} \rangle$ of the population sparseness, 0.77, and thus conclude that the representation in the inferior temporal visual cortex of objects and faces is weakly ergodic. We note that, although Lehky et al. (2005) have modelled such a situation of weak ergodicity for V1, the present investigation may be the first in which this has been shown to apply in any brain region.

We note that weak ergodicity necessarily occurs if a^{s} and a^{p} are the same and the neurons are uncorrelated, i.e. each neuron is independently tuned to the set of stimuli (Lehky et al. 2005). (The independence in this case refers to the fact that the mean response profiles of the neurons to a set of stimuli are uncorrelated, and this was shown to be the case in the analysis described in the penultimate paragraph of Sect. 3 (Rolls et al. 2003, 2004).) The fact that both hold for this population of neurons thus indicates that their responses are uncorrelated, and this is potentially an important conclusion about the encoding of stimuli by these neurons. This conclusion is confirmed by the linear increase in the information with the number of neurons which is the case not only for this set of neurons (Fig. 7), but also in other datasets for the inferior temporal visual cortex (Rolls et al. 1997a). Both types of evidence thus indicate that the encoding provided by at least small subsets (up to e.g. 20 neurons) of inferior temporal cortex neurons is approximately independent (non-redundant), which is an important principle of cortical encoding.

In conclusion, the results described here are the first we know to directly address the issue of the sparseness of the

population code of inferior temporal cortex neurons, and the first to directly compare the single cell and population sparsenesses, and to show that weak ergodicity is apparent in the neuronal representations in the inferior temporal visual cortex. Moreover, these findings provide a new type of evidence that the encoding provided by different inferior temporal cortex neurons is, up to tens of neurons, approximately independent. The findings further show that if all the neuronal spiking is considered in the inferior temporal visual cortex (i.e. including from all neurons and stimuli analysed), then the pdf is very close to exponential. This has the potentially important implication that an exponential distribution of firing rates will be seen by any neuron receiving information from the inferior temporal visual cortex. The fit to an exponential firing rate distribution also has the interesting interpretation that in a large region of cortex such as the inferior temporal visual cortex, there will be, in response to effective stimuli, a high probability of low firing rates (which might be metabolically efficient) and a long thin tail of high firing rates.

Acknowledgments This research was supported by Wellcome Trust and Medical Research Council Grants to Professor E. T. Rolls. LF acknowledges support through a Ramon y Cajal Fellowship from the Spanish Ministry of Education and Science. LF and JMJ acknowledge support from grant CICYT-TIN2005-02984 (FEDER). We acknowledge helpful discussions with Drs. Alessandro Treves (Trieste) and Gustavo Deco (Barcelona).

References

Aggelopoulos NC, Franco L, Rolls ET (2005) Object perception in natural scenes: encoding by inferior temporal cortex simultaneously recorded neurons. J Neurophysiol 93:1342–1357

Atick JJ (1992) Could information theory provide an ecological theory of sensory processing? Nature 3:213–251

Baddeley RJ, Abbott LF, Booth MJA, Sengpiel F, Freeman T, Wakeman EA, Rolls ET (1997) Responses of neurons in primary and inferior temporal visual cortices to natural scenes. Proc R Soc Lond B 264:1775–1783

Barlow HB (1961) Possible principles underlying the transformation of sensory messages. In: Rosenblith W (ed) MIT Press, Sensory Communication. Cambridge

Barlow HB, Kaushal TP, Mitchison GJ (1989) Finding minimum entropy codes. Neural Comput 1:412–423

Booth MCA, Rolls ET (1998) View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex. Cereb Cortex 8:510–523

Desimone R (1991) Face-selective cells in the temporal cortex of monkeys. J Cogn Neurosci 3:1–8

Feigenbaum JD, Rolls ET (1991) Allocentric and egocentric spatial information processing in the hippocampal formation of the behaving primate. Psychobiology 19:21–40

Field DJ (1994) What is the goal of sensory coding? Neural Comput 6:559–601

Field DJ (1999) Wavelets, vision, and the statistics of natural scenes. Philos Trans R Soc Lond A 357:2527–2542

Foldiak P (2003) Sparse coding in the primate cortex. In: Arbib MA (ed) The handbook of brain theory and neural networks. MIT Press, Cambridge, pp 1064–1068

Franco L, Rolls ET, Aggelopoulos NC, Treves A (2004) The use of decoding to analyze the contribution to the information of the correlations between the firing of simultaneously recorded neurons. Exp Brain Res 155:370–384

Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. Proc Nat Acad Sci USA 79:2554–2558

Judge SJ, Richmond BJ, Chu FC (1980) Implantation of magnetic search coils for measurement of eye position: An improved method. Vis Res 20:535–538

Lehky SR, Sejnowski TJ, Desimone R (2005) Selectivity and sparseness in the responses of striate complex cells. Vis Res 45:57–73

Levy WB, Baxter RA (1996) Energy efficient neural codes. Neural Comput 8:531–543

Masuda N, Aihara K (2003) Ergodicity of spike trains: when does trial averaging make sense? Neural Comput 15:1341–1372

Olshausen BA, Field DJ (1997) Sparse coding with an overcomplete basis set: a strategy employed by V1. Vis Res 37:3311–3325

Olshausen BA, Field DJ (2004) Sparse coding of sensory inputs. Curr Opin Neurobiol 14:481–487

Panzeri S, Treves A (1996) Analytical estimates of limited sampling biases in different information measures. Network 7:87–107

Perrett DI, Rolls ET, Caan W (1982) Visual neurons responsive to faces in the monkey temporal cortex. Exp Brain Res 47:329–342

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 (2000) Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition. Neuron 27:205–218

Rolls ET (2005) Emotion Explained. Oxford University Press, Oxford

Rolls ET (2007) The representation of information about faces in the temporal and frontal lobes. Neuropsychologia 45:125–143

Rolls ET (2008) Memory, attention, and decision-making: a unifying computational neuroscience approach. Oxford University Press, Oxford

Rolls ET, Tovee MJ (1995) Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. J Neurophysiol 73:713–726

Rolls ET, Deco G (2002) Computational neuroscience of vision. Oxford University Press, Oxford

Rolls ET, Treves A (1998) Neural networks and brain function. Oxford University Press, Oxford

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, Sanghera MK, Roper-Hall A (1979) The latency of activation of neurons in the lateral hypthalamus and substantia innominata during feeding in the monkey. Brain Res 164:121–135

Rolls ET, Treves A, Tovee MJ (1997a) The representational capacity of the distributed encoding of information provided by populations of neurons in the primate temporal visual cortex. Exp Brain Res 114:177–185

Rolls ET, Treves A, Tovee MJ, Panzeri S (1997b) Information in the neuronal representation of individual stimuli in the primate temporal visual cortex. J Comput Neurosci 4:309–333

Rolls ET, Franco L, Aggelopoulos NC, Reece S (2003) An information theoretic approach to the contributions of the firing rates and correlations between the firing of neurons. J Neurophysiol 89:2810–2822

Rolls ET, Aggelopoulos NC, Franco L, Treves A (2004) Information encoding in the inferior temporal cortex: contributions of the firing rates and correlations between the firing of neurons. Biol Cybern 90:19–32

Rolls ET, Franco L, Aggelopoulos NC, Perez JM (2006) Information in the first spike, the order of spikes, and the number of spikes provided by neurons in the inferior temporal visual cortex. Vis Res 46:4193–4205

Tanaka K (1996) Inferotemporal cortex and object vision. Annu Rev Neurosci 19:109–139

Tovee MJ, Rolls ET (1995) Information encoding in short firing rate epochs by single neurons in the primate temporal visual cortex. Vis Cogn 2:35–58

Tovee MJ, Rolls ET, Azzopardi P (1994) Translation invariance in the responses to faces of single neurons in the temporal visual cortical areas of the alert macaque. J Neurophysiol 72:1049–1060

Tovee MJ, Rolls ET, Treves A, Bellis RP (1993) Information encoding and the responses of single neurons in the primate temporal visual cortex. J Neurophysiol 70:640–654

Treves A (1990) Graded-response neurons and information encodings in autoassociative memories. Physical Review A 42:2418–2430

Treves A, Rolls ET (1991) What determines the capacity of autoassociative memories in the brain? Network 2:371–397

Treves A, Panzeri S (1995) The upward bias in measures of information derived from limited data samples. Neural Comput 7:399–407

Treves A, Panzeri S, Rolls ET, Booth M, Wakeman EA (1999) Firing rate distributions and efficiency of information transmission of inferior temporal cortex neurons to natural visual stimuli. Neural Comput 11:611–641

Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. Science 287:1273–1276

Vinje WE, Gallant JL (2002) Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1. J Neurosci 22:2904–2915

Vogels R (1999) Categorization of complex visual images by rhesus monkeys. Part 2: single-cell study. Eur J Neurosci 11:1239–1255

Willmore B, Tolhurst DJ (2001) Characterizing the sparseness of neural codes. Network 12:255–270

Young MP, Yamane S (1992) Sparse population coding of faces in the inferotemporal cortex. Science 256:1327–1331