Available online at www.sciencedirect.com
ScienceDirect
Vision Research 46 (2006) 4193–4205
Vision Research
www.elsevier.com/locate/visres

Information in the first spike, the order of spikes, and the number of spikes provided by neurons in the inferior temporal visual cortex

Edmund T. Rolls *, Leonardo Franco [1], Nikolaos C. Aggelopoulos, Jose M. Jerez [1]

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

Received 26 October 2005; received in revised form 18 July 2006

Abstract

Information theoretic analyses showed that for single inferior temporal neurons and neuronal populations, more information was encoded in 20 or more ms by all the spikes available than just by the first spike in the same time window about which of 20 objects or faces was shown. Further, the temporal order in which the first spike arrived from different simultaneously recorded neurons did not encode more information than was present in the first spike or the spike counts. Thus information transmission in the inferior temporal cortex by the number of spikes in even short time windows is fast, and provides more information than only the first spike, or the spike order from different neurons.

Keywords: Object recognition; Inferior temporal cortex; First spike; Spike order

1. Introduction

The question of how information is encoded by neuronal activity 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 (Baddeley et al., 1997; Desimone, 1991; Perrett, Rolls, & Caan, 1982; Rolls, 2000, 2005b; Rolls & Deco, 2002; Rolls & Tovee, 1995; Rolls, Treves, Tovee, & Panzeri, 1997b; Tanaka, 1996; Treves, Panzeri, Rolls, Booth, & Wakeman, 1999). An important issue is how rapidly information can be read from these neurons, or from any stage of visual cortical processing. A rapid readout of information from any one stage is important, for the ventral visual system is organised as a hierarchy of cortical areas, and the neuronal response latencies are approximately 100 ms in the inferior temporal visual cortex, and 40–50 ms in the primary visual cortex, allowing only approximately 50–60 ms of processing time for V1–V2–V4–inferior temporal cortex (Baylis, Rolls, & Leonard, 1987; Nowak & Bullier, 1997; Rolls & Deco, 2002). There is much evidence that the time required for each stage of processing is relatively short. For example, visual stimuli presented in succession approximately 15 ms apart can be separately identified (Keysers & Perrett, 2002); the reaction time for identifying visual stimuli is relatively short (Bacon-Mace, Mace, Fabre-Thorpe, & Thorpe, 2005; Thorpe, Fize, & Marlot, 1996; VanRullen & Thorpe, 2001a); a considerable amount of information about which stimulus was shown is available in the spike counts of single inferior temporal cortex neurons in 20 ms (Tovee & Rolls, 1995); and we have shown in a backward masking paradigm that neurons in the inferior temporal cortex fire for only approximately 30 ms to a visual stimulus presented for 16 ms which can be identified above chance (Rolls, 2003; Rolls & Tovee, 1994; Rolls, Tovee, & Panzeri, 1999; Rolls, Tovee, Purcell, Stewart, & Azzopardi, 1994). Delorme and Thorpe (2001) have suggested that just one spike from each neuron is sufficient, and indeed it has been suggested

* Corresponding author. Fax: +44 1865 310447.
E-mail address: Edmund.Rolls@psy.ox.ac.uk (E.T. Rolls).
[1] Present address: Universidad de Malaga, Depto. de Lenguajes y Cs. de la Computacion, Campus de Teatinos S/N, 29071 Malaga, Spain.

0042-6989/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2006.07.026

that the order of the first spike in different neurons may be part of the code (Thorpe, Delorme, & Van Rullen, 2001; VanRullen, Guyonneau, & Thorpe, 2005; VanRullen & Thorpe, 2001b). An alternative view is that the number of spikes in a fixed time window over which a postsynaptic neuron could integrate information is more realistic, and this time might be in the order of 20 ms for a single receiving neuron, or much longer if the receiving neurons are connected by recurrent collateral associative synapses and so can integrate information over time (Deco & Rolls, 2006; Rolls & Deco, 2002; Rolls et al., 1999; Tovee & Rolls, 1995; Tovee, Rolls, Treves, & Bellis, 1993). Although the number of spikes in a short time window of e.g. 20 ms is likely to be 0, 1 or 2, the information available may be more than that from the first spike alone, and we examine this by actually measuring neuronal activity, and then applying quantitative information theoretic methods to measure the information transmitted by single spikes, and within short time windows. Moreover, we perform this analysis both for single neurons, and for populations of neurons, and measure whether more information is available if the order of the spike arrival times from each neuron is taken into account in addition to whether any spike is present or not from each neuron.

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 a rhesus macaque performing a visual fixation task using experimental procedures similar except as described below to those described in detail previously (Rolls, Treves, & Tovee, 1997a). The stimuli included $S = 20$ grayscale 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 & Tovee, 1995). The resolution of these images was 256 wide by 256 high with 256 grey levels.

The neurons were selected where possible 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, then a second time in a new random sequence, etc. Populations of 2–9 neurons were recorded simultaneously using 2–4 independently movable single neuron epoxy-insulated tungsten electrodes with uninsulated tip diameters of less than 10 microns (FHC Inc., USA) using an Alpha-Omega (Israel) recording system. Typically, we were able to move the microelectrodes until 2–4 of the simultaneously recorded neurons responded differentially to the set of stimuli used. This differential firing of the 2–4 neurons was evident in the firing rates to the different stimuli. For the other typically 2–4 neurons being simultaneously recorded in the same brain region, there were no differential firing rates to the stimuli, but these were also included in some of the analyses as described below. The recordings were made as part of the experimental design in one rhesus macaque, *Macaca mulatta*, so that an analysis could be performed of non-simultaneously recorded neurons in which the information from all the recordings made from different neurons in different sessions in the same animal could be analysed as described by Rolls et al. (1997a). 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 more than two spikes on one microelectrode were clearly separated into different waveform clusters using the Datawave (CO, 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, 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 degrees with the search coil technique (Judge, Richmond, & Chu, 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 degrees 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 150–3350 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 & Rolls, 1995; Tovee, Rolls, & Azzopardi, 1994).) 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 1 degree 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. Measuring the information from many recorded neurons using a decoding procedure

Estimates of how much information was available from a population of neurons in a fixed time window about which stimulus was shown were calculated using the decoding method described by Rolls et al. (1997a) to analyse the information. This method measures the amount of information available from the number of spikes from each neuron assuming non-simultaneous recording so that it can be applied to the data accumulated from different cells recorded in the same macaque over different days. The method can be used for very large numbers of cells, and when there are many spikes in a time window. The method uses a decoding procedure in which on each trial the probability that each stimulus (called s') was shown is estimated from the vector of neuronal responses. This estimate is made by comparing the vector of neuronal responses on that trial to the average response vectors to each stimulus. Then, knowing the actual stimulus shown on that trial, the mutual information $\langle I_p \rangle$ between the estimated stimulus s' and the real stimulus s over the set of stimuli S can be calculated as:

$$I_{\mathrm{p}} = \sum_{s=1}^{S} \sum_{s'=1}^{S} P(s,s') \log_2 \frac{P(s,s')}{P(s)P(s')}.$$

The decoding procedure used for the results presented here is Bayesian probability estimate (PE) decoding using a Gaussian fit, as described by Rolls et al. (1997a), Rolls and Treves (1998), Rolls and Deco (2002) and Franco, Rolls, Aggelopoulos, and Treves (2004), and includes cross-validation procedure and bias correction procedures as described in detail by Rolls et al. (1997a) and Franco et al. (2004). The same analysis program was used for the calculation of the information from single cells.

Further details of the decoding procedures are as follows (see also Rolls et al. (1997a) and Franco et al. (2004)). The full probability table estimator (PE) algorithm uses a Bayesian approach to extract $P(s'|\mathbf{r})$ for every single trial from an estimate of the probability $P(\mathbf{r}|s')$ of a stimulus-response pair made from all the other trials (as shown in Bayes' rule shown in Eq. (1)) in a cross-validation procedure described by Rolls et al. (1997a)

$$P(s'|\mathbf{r}) = \frac{P(\mathbf{r}|s')P(s')}{P(\mathbf{r})}, \tag{1}$$

where $P(\mathbf{r})$ (the probability of the vector containing the firing rate of each neuron, where each element of the vector is the firing rate of one neuron) is obtained as:

$$P(\mathbf{r}) = \sum_{s'} P(\mathbf{r}|s')P(s'). \tag{2}$$

This requires knowledge of the response probabilities $P(\mathbf{r}\ |s')$ which can be estimated for this purpose from $P(\mathbf{r},s')$, which is equal to $P(s')\prod_c P(r_c|s')$, where r_c is the firing rate of cell c . We note that $P(r_c|s')$ is derived from the responses of cell c from all of the trials except for the current trial for which the probability estimate is being made. The probabilities $P(r_c|s')$ are fitted with a Gaussian (or Poisson) distribution whose amplitude at r_c gives $P(r_c|s')$. By summing over different test trial responses to the same stimulus s, we can extract the probability that by presenting stimulus s the neuronal response is interpreted as having been elicited by stimulus s',

$$P(s'|s) = \sum_{\mathbf{r}\in\mathbf{test}} P(s'|\mathbf{r})P(\mathbf{r}|s) \tag{3}$$

After the decoding procedure, the estimated relative probabilities (normalized to 1) were averaged over all 'test' trials for all stimuli, to generate a (regularized) table $P_N^R(s|s')$ describing the relative probability of each pair of actual stimulus s and posited stimulus s' (computed with N trials). From this probability table the mutual information measure I_{p} was calculated as described above.

Because the probability tables from which the information is calculated may be unregularized with a small number of trials, a bias correction procedure to correct for the undersampling is applied, as described in detail by Rolls et al. (1997a) and Panzeri and Treves (1996). In practice, the bias correction that is needed with information estimates using the decoding procedures described here and by Rolls et al. (1997a, 1997b) is small, typically less than 10% of the uncorrected estimate of the information, provided that the number of trials for each stimulus is in the order of the number of stimuli. We also note that the distortion in the information estimate from the full probability table needs less bias correction than that from the predicted stimulus table (i.e. maximum likelihood) method, as the former is more regularized because every trial makes some contribution through much of the probability table (see Rolls et al. (1997a)).

2.4. Measuring the information in the order of spike arrival times

The decoding method as just described measures how much information is present in the first spike, or in the number of spikes in a given time window. It is possible that there is in addition extra information in the order in which the first spike arrives from different neurons about which stimulus is shown. Although normally the order of spike arrival would be expected to depend on the average firing rate of neurons (with for example a neuron with a high average firing rate to a particular stimulus likely to produce an early spike), in principle the order information could be distinct from the rate information (with for example a neuron with a low firing mean firing rate to a stimulus nevertheless producing an early spike to that stimulus), thus providing additional information independent from the mean response about which stimulus was shown.

We were able to test for this type of order information by ordering the first spike arrival times for each neuron on each trial to each stimulus, and measuring the information when it was this rank ordering that was being decoded by the dot product algorithm described below (see also Franco et al., 2004; Rolls et al., 1997a). To measure the information contained in the relative order of the spikes from different neurons, for each recorded trial we constructed a vector with a number of components equal to the number of simultaneously recorded neurons. Each component contained the rank order of the first spike of that neuron relative to the other neurons. (The component contained one if the first spike from that neuron was the first among all neurons to arrive, 2 if it was the second, etc. If there was no spike from that neuron on that trial, the component was set to one greater than the number of spikes that had arrived.) To measure the information contained in the order of the arrival of the spikes of the different neurons, dot product decoding then compares each trial with the average of the trials to each stimulus, to determine how close the current trial is to the data obtained to each stimulus. (The average vector for each stimulus is calculated excluding the data for the current trial in what is a cross-validation procedure (Franco et al., 2004; Rolls et al., 1997a).) The dot product (DP) algorithm computes the normalized dot product between the current firing vector $\mathbf{r}$ on a 'test' (i.e. the current) trial and each of the mean firing rate response vectors in the 'training' trials for each stimulus s' in the cross-validation procedure. (The normalized dot product is the dot or inner product of two vectors divided by the product of the length of each vector. The length of each vector is the square root of the sum of the squares.) Thus, what is computed are the cosines of the angles of the test vector of cell rates with, in turn for each stimulus, the mean response vector to that stimulus. The highest dot product indicates the most likely stimulus that was presented, and this is taken as the best guess or (the predicted stimulus $\mathbf{S}^{\mathrm{p}}$) for the probability table $P(\mathbf{S},\mathbf{S}^{\mathrm{p}})$ with maximum likelihood information measurement (Rolls et al., 1997a). We verified that using this rank ordering in this information measurement procedure was effective by simulating spike trains with the same mean firing rate of each neuron to any stimulus, and then introducing rank ordering so that the order of spikes from the different neurons was different for each stimulus. The control condition was to take the same number of spikes in the time window in which the information was being measured, but to allocate the time of any spikes in that window to random times (with a uniform probability distribution), thus preserving any information present in the number of spikes in the time window (i.e. the "rate" information), but removing any information present in the relative time at which spikes arrived from different neurons for each stimulus. In both cases, just the first spike in the time window was used in the information analysis. The reason for using just the first spike is that the order information becomes arbitrary after the first spike from each neuron, for it is not assumed that the order encoding in the brain acts like a counting processor that is able to take into account for example that the second spike from neuron x has to follow the first spike from neuron y. Indeed, when suggesting that order information was encoded, Delorme and Thorpe (2001) used just one spike per neuron.

3. Results

From 62 cells recorded simultaneously in groups of 2–4 neurons in 31 experiments, we performed the analyses on 21 neurons that had significant differences in their firing rate to the set of 20 stimuli, as shown by ANOVA ($p < .001$). There were typically 20 trials of data available for each stimulus for each neuron. Examples of peristimulus time rastergrams and histograms for a neuron are illustrated in Fig. 1. Most of the neurons had their best

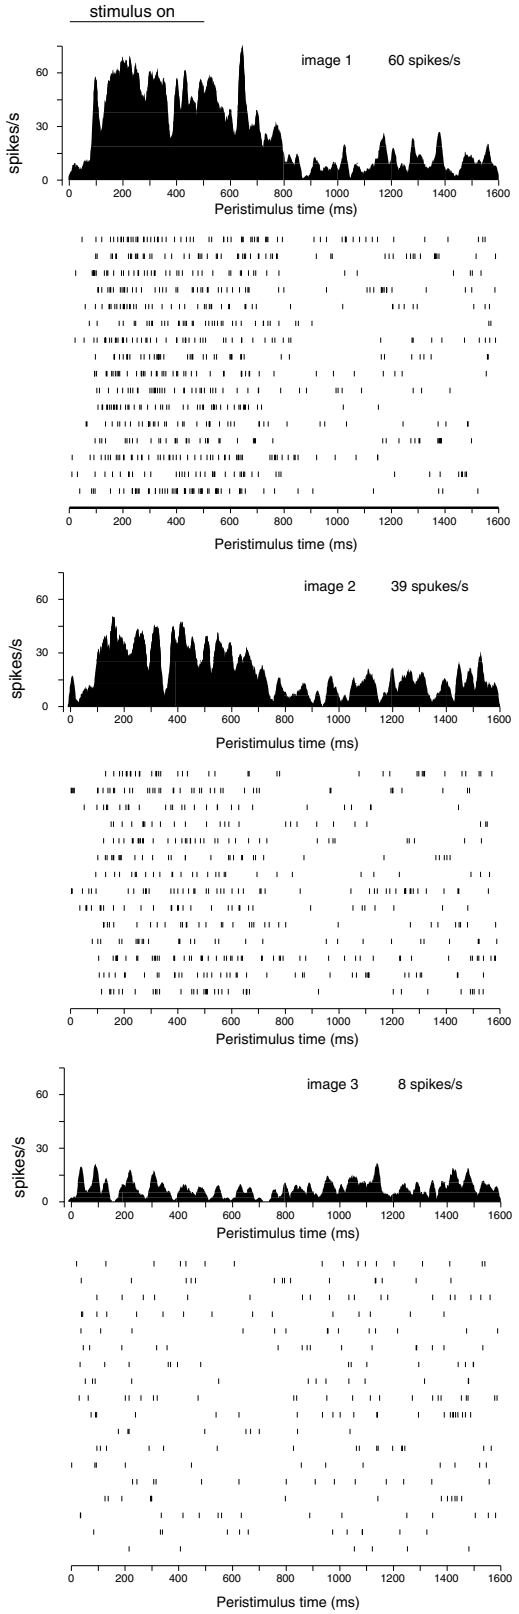

Fig. 1. Peristimulus rastergrams and time histograms for the responses of a single neuron to three different stimuli. The stimuli were presented at time 0, and had a duration of 500 ms. The mean firing rates calculated in the period 100–500 ms after stimulus onset are indicated. The spontaneous firing rate of the neuron was 11.6 spikes/s, close to that for the ineffective stimulus illustrated, image 3 (which produced a small decrease in firing rate in the period 100–500 ms post-stimulus).

responses to objects, and such neurons sometimes responded to one or two faces in the set.

We began by measuring the information in the first spike of neurons about which stimulus was present, and compared this with the information present in time windows starting at the same time but extending for longer durations. We did this for both single cells, and for multiple simultaneously recorded cells. This addresses the issue of how much information is available in the first spike compared to longer spike trains. To further analyze the speed and nature of cortical processing, we measured the information available in short fixed time windows (of 20 and 50 ms), and compared this to the information available from the first spike. To measure whether the order of the arrival times of the spikes of different neurons is part of the code, as has been hypothesized (Thorpe et al., 2001), we measured the information available when the order was retained, and compared it to the information present when the order of the spike times of each neuron for each stimulus was made random as described in Section 2.

The cumulative single cell information about which of the 20 stimuli was shown from all spikes and from the first spike starting at 100 ms after stimulus onset is shown in Fig. 2a. One hundred milliseconds is just longer than the shortest response latency of the neurons from which recordings were made, so starting the measure at this time provides the best chance for the single spike measurement to catch a spike that is related to the stimulus. The means and standard errors across the 21 different neurons are shown. The cumulated information from the total number of spikes is larger than that from the first spike only, and this is evident and significant within 50 ms of the start of the time epoch. In calculating the information from the first spike, just the first spike in the analysis window starting in this case at 100 ms after stimulus onset was used.

The brain does not know when to start looking for the first spike to a stimulus, so we show a re-analysis starting at the time at which the stimulus appeared in Fig. 2b. As expected, the information measure starts to increase at approximately 100 ms. The information measures do not in general reach such high values as when the analysis starts at 100 ms, and the reason for this is that any spikes due to spontaneous neuronal firing in the period 0–100 ms after stimulus presentation are effectively noise in the system. This effect is particularly a problem for the single spike measurement condition, for the system just utilizes the first spike after time 0, and this might well be a spontaneous activity-related spike with no information at all about the stimulus. This underlines the fact that the single spike measure does rather make the assumption that one knows when to start measuring the effects of the stimulus, if much information about the stimulus is to be extracted.

Because any one neuron receiving information from the population being analyzed has multiple inputs, we show in Fig. 3 the cumulative information that would be available from multiple cells (21) about which of the 20 stimuli was

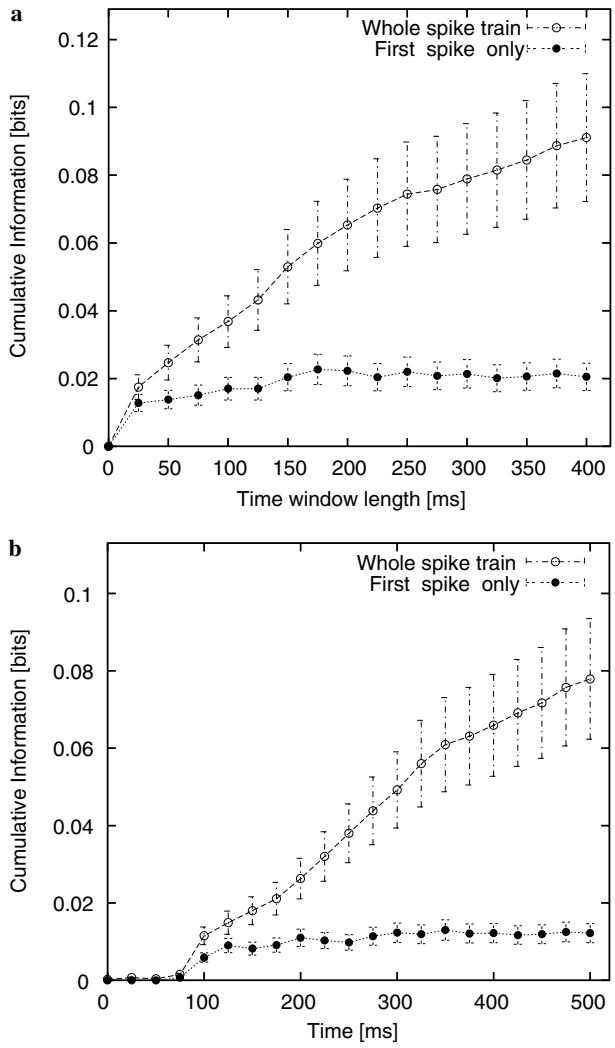

Fig. 2. (a) Cumulative single cell information from all spikes and from the first spike with the analysis starting at 100 ms after stimulus onset. The mean and SEM over the 21 neurons are shown. (b) Cumulative single cell information from all spikes and first spike starting at 0 ms with respect to stimulus onset. The mean and SEM over the 21 neurons are shown.

shown, taking both the first spike after 100 ms, and the total number of spikes after 100 ms from each neuron (Fig. 3a). The cumulative information even from multiple cells is much greater when all the spikes rather than just the first spike are used. The same conclusion is reached if the measurement starts at 0 ms with respect to the stimulus onset, as shown in Fig. 3b. As in Fig. 2, the information particularly with the first spike is less when starting at 0 ms than at 100 ms, because in the 0 ms condition the first spike may not reflect any firing produced by the stimulus, as it may just be spontaneous activity.

An attractor network might be able to integrate the information arriving over a long time period of several hundred ms, and might produce the advantage shown in Fig. 3 for the whole spike train compared to the first spike only. However a single layer pattern association network might only be able to integrate the information over the time constants of its synapses and cell membrane, which might be in the order of 15–30 ms (Panzeri, Rolls, Battaglia, & Lavis, 2001; Rolls & Deco, 2002). In a hierarchical processing system such as the visual cortical areas, there may only be a short time during which each stage may decode the information from the preceding stage, and then pass on information sufficient to support recognition to the next stage (Rolls & Deco, 2002). We therefore analyzed the information that would be available in short epochs from multiple inputs to a neuron, and show the multiple cell information for the population of 21 neurons in Fig. 4a (for a 20 ms epoch) and in Fig. 4b (for a 50 ms epoch). The epochs started at 0 ms with respect to stimulus onset. We see in this case that the first spike information, because it is being made available from many different neurons (in this case 21 selective neurons discriminating between the stimuli each with $p < .001$ in the ANOVA), fares better relative to the information from all the spikes in these short epochs, but is still less than the information from all the spikes, particularly in the 50 ms epoch. In particular, for the epoch starting 100 ms after stimulus onset in Fig. 4a the information in the 20 ms epoch is 0.37 bits, and from the first spike is 0.24 bits. Correspondingly, for the 50 ms epoch, the values in the epoch starting at 100 ms post stimulus are 0.66 bits for the 50 ms epoch, and 0.40 bits for the first spike. Thus, with a population of neurons, having just one spike from each can allow considerable information to be read if only a limited period (of e.g. 20 or 50 ms) is available for the readout, though even in these cases, more information was available if all the spikes in the short window are considered (Fig. 4a and b).

To show how the information increases with the number of neurons in the ensemble in these short epochs, we show in Fig. 4c the information from different numbers of neurons for a 20 ms epoch starting at time = 100 ms with respect to stimulus onset, for both the first spike condition and the condition with all the spikes in the 20 ms window. The linear increase in the information in both cases indicates that the neurons provide independent information, which could be because there is no redundancy or synergy, or because these cancel (Rolls, Aggelopoulos, Franco, & Treves, 2004; Rolls, Franco, Aggelopoulos, & Reece, 2003b). It is also clear from Fig. 4c that even with the population of neurons, and with just a short time epoch of 20 ms, more information is available from the population if all the spikes in 20 ms are considered, and not just the first spike. The 20 ms epoch analyzed for Fig. 4c is for the post-stimulus time period of 100–120 ms. (Given this, the values for 21 neurons in Fig. 4c correspond to the values shown for time 100–120 ms in Fig. 4a.) Thus, even with a population of neurons selected to be tuned to the stimulus set, at 21 neurons the information available from the first spike only has not reached that obtainable from all the spikes in the 20 ms epoch. (The first spike information measures the information obtained by the presence or absence of a spike in the given time window.)

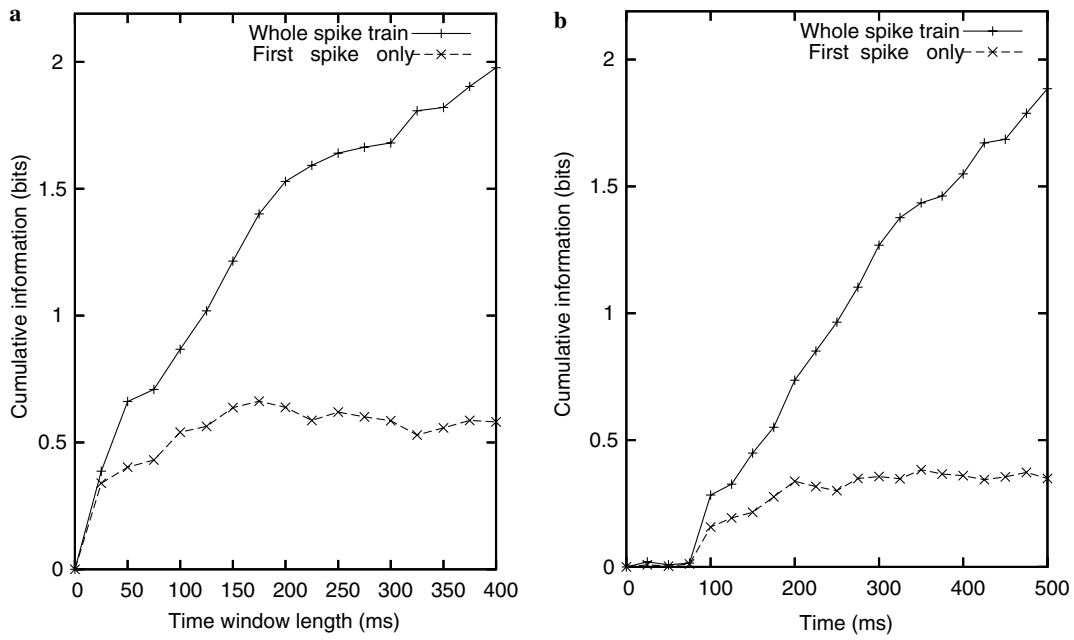

Fig. 3. (a) Cumulative multiple cell information from all spikes and first spike starting at 100 ms for the population of 21 neurons about the set of 20 stimuli. (b) Cumulative multiple cell information from all spikes and first spike starting at 0 ms for the population of 21 neurons about the set of 20 stimuli.

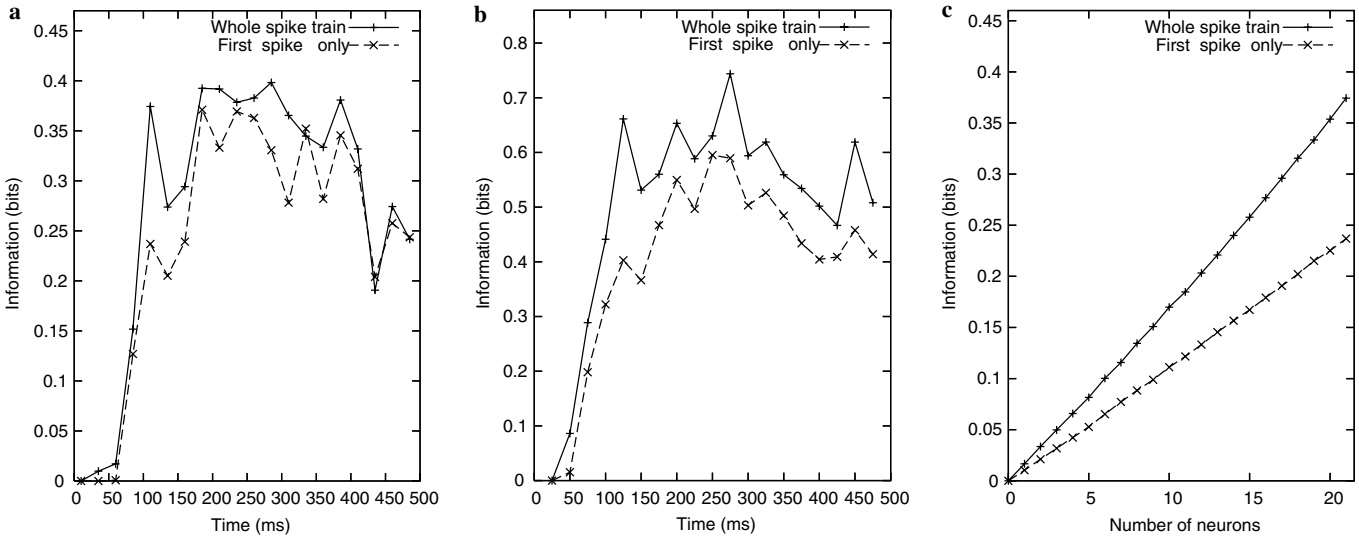

Fig. 4. (a) Multiple cell information from all spikes and 1 spike in 20 ms time windows taken at different post-stimulus times starting at time 0. (b) Multiple cell information from all spikes and 1 spike in 50 ms time windows taken at different post-stimulus times starting at time 0. (c) Multiple cell information from all spikes and 1 spike in a 20 ms time windows starting at 100 ms after the stimulus onset as a function of the number of neurons in the ensemble.

To elucidate the spike train properties that give rise to the information just described, we show in Fig. 5a the probability distribution of the number of spikes on each trial from each neuron for the most effective stimulus for each neuron, for a 20 ms epoch starting at 100 ms after stimulus onset. Because the data in Fig. 5a is computed across all 21 neurons, we note that some individual neurons will have relatively higher probabilities of 2 or more spikes on each trial in this time window. For comparison, we show in Fig. 5b the probability distribution of the number of spikes on each trial from each neuron across all stimuli for each neuron, for the same 20 ms epoch.

To assess whether there is information that is specifically related to the order in which the spikes arrive from the different neurons, we computed for every trial the order across the different simultaneously recorded neurons in which the first spike arrived to each stimulus, and used this in the information theoretic analysis described in Section 2. The control condition was to randomly allocate the order

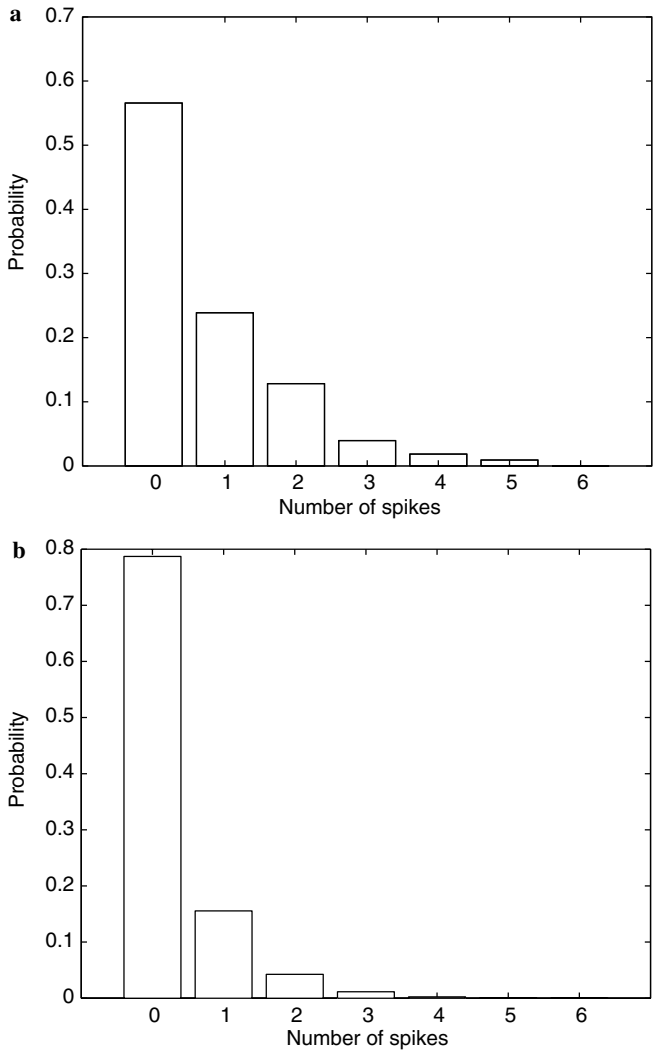

Fig. 5. (a) The probability distribution of the number of spikes on each trial from each neuron for the most effective stimulus for each neuron, for a 20 ms epoch starting at 100 ms after stimulus onset. (b) The probability distribution of the number of spikes on each trial from each neuron across all stimuli for each neuron, for the same 20 ms epoch.

values for each trial between the neurons that had any spikes on that trial, thus shuffling or scrambling the order of the spike arrival times in the time window. In both cases, just the first spike in the time window was used in the information analysis. (In both the order and the shuffled control conditions, on some trials some neurons had no spikes, and this itself in comparison with the fact that some neurons had spiked on that trial, provided some information about which stimulus had been shown. However, by explicitly shuffling in the control condition the order of the spikes for the neurons that had spiked on that trial, comparison of the control with the unshuffled order condition provides a clear measure of whether the order of spike arrival from the different neurons itself carries useful information about which stimulus was shown.) The dataset was 36 cells with significantly different ($p < .05$) responses to the stimulus set where it was possible to record simultaneously from groups of 3 and 4 cells (so that the order on each trial could be measured) in 11 experiments. Taking a 75 ms time window starting 100 ms after stimulus onset, the information with the order of arrival times of the spikes was 0.142 ± 0.02 bits, and in the control (shuffled order) condition was 0.138 ± 0.02 bits (mean across the 11 experiments $\pm$SEM). Thus the information increase by taking into account the order of spike arrival times relative to the control condition was only $(0.142 - 0.138) = 0.004$ bits per experiment (which was not significant, $p = .7$, $t = 0.33$, $\mathrm{df} = 10$, paired t test). For comparison, the information calculated for the first spike using the same dot product decoding, as described above was, 0.136 ± 0.03 bits per experiment. Analogous results were obtained for different time windows, as follows. For a time window of 50 ms, when the spike order was taken into account, the information was 0.104 ± 0.03 bits, and when the spike order was scrambled in the control condition, the information was 0.115 ± 0.03 bits. For a time window of 25 ms, when the spike order was taken into account, the information was 0.108 ± 0.03 bits, and when the spike order was scrambled in the control condition, the information was 0.091 ± 0.03 bits. As before, these differences were not significant. Thus, taking the spike order into account compared to a control condition in which the spike order was scrambled made essentially no difference to the amount of information that was available from the populations of neurons about which stimulus was shown.

We acknowledge that in single neuron neurophysiology some of the neurons studied are selected to have changes in firing rate to the set of stimuli, and that if a specialized population of neurons might carry information in other ways such as the relative temporal order of spike arrival from different neurons and specifically without any change of firing rate, this could bias the analysis against these other possible types of encoding. First, we note that it is a priori unlikely, though possible, that only neurons with no change of firing rate would convey such relative temporal order information, as many neurons in visual cortical areas do respond to visual stimuli by altering their firing rates. Second, we explicitly took steps in the present study to include some neurons in the analysis that did not have changes in firing rate to the set of stimuli, as described next, just in case relative temporal order information then became apparent in the analyses. So as not to bias this analysis of information present in spike order only towards a dataset of neurons that is highly stimulus selective (in case the order information may be expressed in other neurons than these), we set the criterion for neurons to be included in the analysis as having a p value in the ANOVA of only .05 for the significance of the difference of the firing rates to the different stimuli. To further address this issue, we performed further analyses of two types. In one, we measured the order information for datasets from neurons where the ANOVA was not significant at all, although the neurons were recorded in the same brain area. In a typical case, we found very low

amounts of information present in the order (e.g. 0.03 bits), found that this was not reduced by shuffling the order on each trial between neurons, and that the information from the first spike analysis was also (as expected) similarly low. In a second approach, we included additional neurons that had no significant rate information in datasets that were included in the order analyses described above. Again, in a typical case, adding such neurons made no significant difference to the order information analysis. The evidence from these two approaches, and the fact that a low criterion was set for the ANOVA in the main order analyses, thus does not suggest that order information might have been missed in the present study because it is not present in neurons that have high selectivity to the stimulus set. We have thus introduced a systematic approach to measuring the information present in the order of spike arrival times of different neurons, have found that there is little information present in the order in the current dataset, and hope that similar quantitative approaches may be applied in future studies.

The recording sites of the neurons analyzed in this paper in the inferior temporal visual cortex are shown in Fig. 6. Zero millimeter 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.

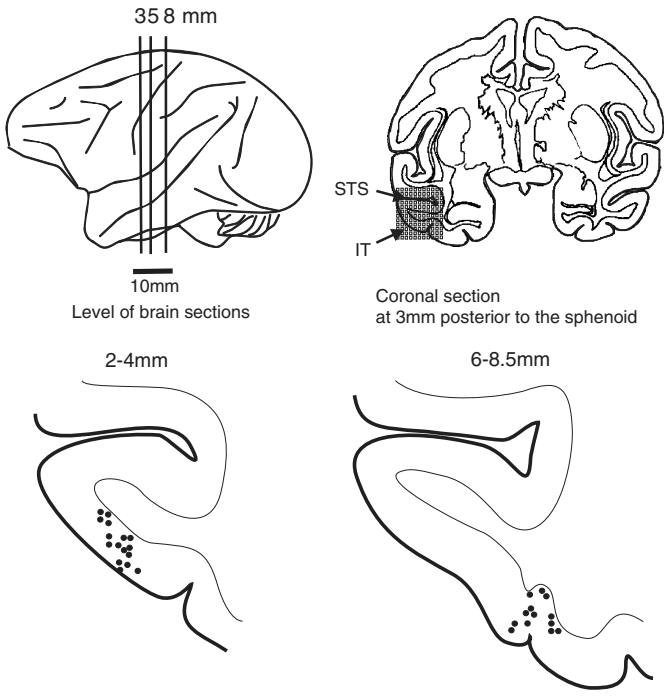

Fig. 6. The recording sites shown on coronal sections of the neurons included in this study. The positions of the coronal sections are shown on a lateral view of the macaque brain. The distances refer to mm posterior (P) to the sphenoid reference plane (see text). STS, superior temporal sulcus; IT, inferior temporal cortex.

4. Discussion

The information analyses shown in Figs. 2–4 show that the information available from the first spike is not as great as that from all the spikes, even in short time windows. In particular, Fig. 2 shows that the cumulated single cell information is much less with the single spike than with all the spikes in the time window. (Information could be cumulated across time by for example an attractor network (Rolls & Deco, 2002).) Fig. 3 shows that even from multiple neurons (21), the cumulated firing rate information is higher if all the spikes rather than only the first spike from each neuron in the same time window is utilized. The point is also made in Figs. 2 and 3b that if one does not know when to start counting the first spike, spontaneous activity adds noise to the information estimate, which is low. In a sense, to utilise the information from the first spike, one needs to be told by a teacher when to start counting, which seems biologically implausible. Fig. 4c shows that even in a short epoch (20 ms), the information from the presence or absence of a first spike from each neuron is somewhat lower than the information computed when all the spikes in the same time window are considered.

These results provide the first quantitative test we know at both the single neuron and population of neuron levels of the hypothesis that the first spike is all that is necessary to decode the stimulus. First, we show that a considerable amount of information is present in the first spike at both the single neuron (Fig. 2) and the population (Fig. 4) levels. Second, the results show that although considerable information is present in the first spike, more information is available under the more biologically realistic assumption that neurons integrate spikes over a short time window (depending on their time constants) of for example 20 ms. The results shown in Fig. 4c are of considerable interest, for they show that even when one increases the number of neurons in the population, the information available from the number of spikes in a 20 ms time window is larger than the information available from just the first spike. Thus although intuitively one might think that one can compensate by taking a population of neurons rather than just a single neuron for using just the first spike instead of the number of spikes available in a fixed time window, this compensation by increasing neuron numbers is insufficient to make the code as efficient as taking the number of spikes. Of course, if the total amount of information required to discriminate a stimulus set of a fixed size (4 bits if 16 stimuli) was reached by both the first spike and spike count codes asymptotically as the number of neurons in the ensemble was reached (Rolls et al., 1997a), then the codes would still be unequal, for more neurons would be required in the first spike only case, and in this sense the first spike only code would be inefficient, and require neurons that with the more efficient code would be redundant.

The encoding of information that uses the number of spikes in a short time window that is supported by the analyses presented here deserves further elaboration. It could

be thought of as a rate code, in that the number of spikes in a short time window is relevant, but is not a rate code in the rather artificial sense considered by Thorpe et al. (Delorme & Thorpe, 2001; Thorpe et al., 2001; VanRullen et al., 2005; VanRullen & Thorpe, 2001b) in which a rate is estimated from the interspike interval. This is not just artificial, but also begs the question of how, once the rate is calculated from the interspike interval, this decoded rate is passed on to the receiving neurons, or how, if the receiving neurons calculate the interspike interval at every synapse, they utilize it. In contrast, the spike count code in a short time window that we consider here is very biologically plausible, in that each spike would inject current into the post-synaptic neuron, and the neuron would integrate all such currents in a dendrite over a time period set by the synaptic and membrane time constants, which will result in an integration time constant in the order of 15–20 ms. Explicit models of exactly this dynamical processing at the integrate-and-fire neuronal level have been described to define precisely the operations to which we refer (Deco & Rolls, 2003, 2005a, 2005c; Deco, Rolls, & Horwitz, 2004; Rolls & Deco, 2002). Even though the number of spikes in a short time window of e.g. 20 ms is likely to be 0, 1 or 2, it can be 3 or more for effective stimuli (see Fig. 5), and this is more efficient than using the first spike. To add some detail here, a neuron receiving information from a population of inferior temporal cortex neurons of the type described here would have a membrane potential that varied continuously in time reflecting with a time constant in the order of 15–20 ms (resulting from a time constant of order 10 ms for AMPA synapses, 100 ms for NMDA synapses, and 20 ms for the cell membrane) a dot (inner) product over all synapses each spike count and the synaptic strength. This continuously time varying membrane potential would lead to spikes whenever the results of this integration process produced a depolarization that exceeded the firing threshold. The result is that the spike train of the neuron would reflect continuously with a time constant in the order of 15–20 ms the likelihood that the input spikes it was receiving matched its set of synaptic weights. The spike train would thus indicate in continuous time how closely (for a dot product is essentially a correlation) the stimulus or input matched its most effective stimulus. In this sense, no particular starting time is needed for the analysis, and in this respect it is a much better component of a dynamical system than is a decoding that utilizes an order in which the order of the spike arrival times is important and a start time for the analysis must be assumed. However, we may note that an autoassociation or attractor network implemented by recurrent collateral connections between the neurons can using its short term memory integrate its inputs over much longer periods, for example over 500 ms in a model of how decisions are made (Deco & Rolls, 2006), and thus if there is time, the extra information available in more than the first spike of few spikes that is evident in Figs. 2 and 3 could be used by the brain.

To assess whether there is information that is specifically related to the order in which the spikes arrive from the different neurons, which has been proposed by Thorpe et al. (Delorme & Thorpe, 2001; Thorpe et al., 2001; VanRullen et al., 2005; VanRullen & Thorpe, 2001b), we computed the information present when the order across the different simultaneously recorded neurons in which the first spike arrived to each stimulus was used. We compared this to a control condition in which the spike arrival times in the time window were allocated at random, but leaving the same number of spikes. The control condition takes into account the fact that if the mean spike rate or count is higher in the time interval, then on average a spike will arrive earlier in that condition. We found that the increase of information by taking the order into account was not significant. Thus, in this first empirical test of which we know of the hypothesis, we found that in the inferior temporal visual cortex there was no significant evidence that the order of the spike arrival times from different simultaneously recorded neurons is important. Indeed, the evidence found in the experiments described here is that the number of spikes in the time window is the important property that is related to the amount of information encoded by the spike trains of simultaneously recorded neurons. Of course it would be interesting to test this at earlier stages of cortical visual processing with the rigorous methods applied to empirical neurophysiological data described here, but as far as we know this has not yet been attempted. We believe that because the thinking of Thorpe et al. (Delorme & Thorpe, 2001; Thorpe et al., 2001; VanRullen et al., 2005; VanRullen & Thorpe, 2001b) did not take into account the spontaneous firing that is typical of cortical neurons, this may have influenced their conclusions. In particular, if an integrate-and-fire conceptualization starts with a hyperpolarized neuron, then inevitably the time at which a neuron fires is causally linked to the average rate at which spikes arrive, and there is not independence of the spike arrival time or the order of the spikes from the average spike count in a short time window hypothesis. Indeed, the rank order of spike arrival times is inevitably linked to the number of counts in any short time window if that type of conceptualization is used. (If the input to a neuron is stronger, it will charge up from reset potential to reached its threshold for firing more quickly giving a low rank for its order of firing, and it will also fire more spikes in a short count period.)

As we have just seen, using the number of spikes in a short time window is very biologically plausible, and more efficient as well as more plausible than just taking the first spike. If one were to take literally the further suggestion that it is the order of arrival of spikes from different neurons that is the code (Delorme & Thorpe, 2001; Thorpe et al., 2001; VanRullen et al., 2005; VanRullen & Thorpe, 2001b), one would need special mechanisms for decoding the ranking, which are not biologically plausible. Further, the spike order code considered did not take into account the spontaneous firing of cortical neurons,

together with its approximately Poisson spike time characteristics. For the neurons described here, the mean spontaneous firing rate was 7.1 spikes/s (SD = 6.7). This spontaneous activity of cortical neurons would make utilization of the order of arrival of spikes in the cortex very difficult to use as a code, just as it made difficult the simpler case shown in Figs. 2 and 3b where one does not know when to start taking the first spike as there is no independent indicator of when the stimulus was shown. The analogous problem for the spike order code hypothesis is from when does one start taking the order of arrival, given that spontaneous spiking is present in the cerebral cortex all the time? We note that spontaneous firing is not only present as described in visual fixation tasks such as this, but is also present with a very similar rate when macaques are searching for a target in a natural visual scene (Aggelopoulos, Franco, & Rolls, 2005; Rolls, Aggelopoulos, & Zheng, 2003a).

In fact, the spontaneous spiking of cortical neurons is part of the solution that we propose to how information is processed fast by the cerebral cortex (Battaglia & Treves, 1998; Rolls & Deco, 2002; Rolls & Treves, 1998; Treves, 1993). If cortical neurons were continually in a hyperpolarized state as they are just after spiking, then the cell membrane would have to charge up before the neuron could respond to a new input, and given that the membrane time constant may be in the order of 20 ms, this could delay information transmission by 20 or more ms. However, by having random (Poisson) spontaneous spiking activity, whenever a new input is received, some neurons will be close to their firing threshold, and will be able to emit a spike almost immediately (within 1–2 ms) in response to the new input. Thus some neuronal responses to the new stimulus can occur within a very few ms of synaptic input to an area. (These new responses will of course be selective, in that if the dot product of the new synaptic input with the synaptic strengths is larger than previously the neuron will be more likely to fire, whereas if the dot product of the new synaptic input with the strengths is smaller than previously, or because of feedforward or even feedback inhibition, the neuron will be less likely to fire. Thus, the spontaneous activity has the effect of ensuring that some neurons are ready to respond fast, but they only do respond if the input is effective for the neuron.) These spikes, earlier than would have occurred by Poisson variability, can then influence other neurons, within a very few ms. These other neurons could be within the next cortical area, resulting in rapid feedforward information transmission. Of course, none of this would imply a fixed time window within which all the information must be read out from one cortical area before the next cortical area can start processing the spikes it is already receiving. Instead, the whole process within a cortical area can instead be seen as a continuous dynamical process which involves feedback inhibition between different neurons, with the neuronal population having its responses set not only by its feedforward inputs from the preceding area, but also by competition from neighboring neurons. With this dynamical process, it is found that the characteristic propagation delays from cortical area to cortical area for useful information transmission is in the order of 5–10 ms per cortical stage (Panzeri et al., 2001). This process utilizes the number of spikes received in the processing time of approximately 20 ms set by the synaptic and membrane time constants (assuming 10 ms for the synaptic and 20 ms for the membrane time constants). It is exactly this rapid processing as a result of spontaneous activity of neurons that enables even recurrent attractor networks to settle into a global basin of attraction very rapidly, in a time in the order of 1–2 time constants of the synapses (Battaglia & Treves, 1998; Rolls & Deco, 2002; Rolls & Treves, 1998; Treves, 1993). In a multilayer network modeling successive cortical areas, it has been shown that this local autoassociation circuitry implementing attractor dynamics could contribute to useful constraint satisfaction and information retrieval in each cortical area with times as short as 15–17 ms per cortical stage, which fits with the time actually taken (Panzeri et al., 2001). Thus, far from it being the case that cortical computation is so fast that there is time for only purely feedforward processing (Delorme & Thorpe, 2001; Thorpe et al., 2001; VanRullen et al., 2005; VanRullen & Thorpe, 2001b), it is instead the case that because of spontaneous firing in the cortex, recurrent processing using excitatory synapses to allow settling into a global attractor within each cortical area can be sufficiently fast to contribute usefully to cortical information processing with the time available, of approximately 15–17 ms per cortical stage.

In fact, other evidence indicates that the first spike from a cortical neuron is insufficient to account for visual recognition let alone for conscious visual perception. For example, in a study of backward masking it was found that if the 16 ms presentation of a test stimulus was followed with no delay by a mask stimulus which started 20 ms after the start of the test stimulus (a stimulus onset asynchrony, SOA, of 20 ms), then identification of which of five faces was seen was better than chance, but was far from perfect, and indeed subjects felt that they were guessing and did not report having consciously seen or processed the test stimuli. Under these conditions, cortical neurons fired for approximately 30 ms (Rolls & Tovee, 1994; Rolls et al., 1999, 1994). With an SOA of 40 ms, the neurons fired for approximately 50 ms, identification performance was better, and some consciousness of the test stimulus was evident (as shown for example by the clarity ratings). Thus, reasonable performance at the task required cortical neurons to be firing for approximately 50 ms, with performance above chance but not good if the neurons were firing for 30 ms. Thus, adequate performance at the task required more than just the first spike, but instead the number of spikes that could occur in a time of cortical firing of 30–50 ms. As shown in the present paper, much more information is available in time windows of 20–50 ms than is just available from one spike, and it must be these extra spikes that contribute to the better recognition performance, and also to the conscious awareness of the stimuli

(Rolls, 2003, 2005a; Rolls & Tovee, 1994; Rolls et al., 1999, 1994). Further evidence for the inadequacy of the first spike alone is that behavioral accuracy in image identification was confirmed to increase with SOAs of 40–60 ms in visual masking (Bacon-Mace et al., 2005).

We have shown elsewhere that most of the information available from the spiking activity of neurons is present in the number of spikes, and not in stimulus-dependent synchrony between the spike firing times of different neurons (Franco et al., 2004; Rolls et al., 2004, 2003b). This is the case even when two objects must be discriminated in a complex natural scene, and processes such as binding and segmentation are required (Aggelopoulos et al., 2005). Thus, on the basis of the evidence presented in this paper, and elsewhere, we suggest that the information is represented by a population of neurons in the number of spikes occurring in a short time window of for example 20 ms, and that the relative spike timing or synchrony of the different neurons carries little additional information (less than 5%). Moreover, we have indicated here and elsewhere how this information could be used very effectively and rapidly by populations of integrate-and-fire neurons with spontaneous activity.

We have thus seen how cortical information could be fast in a dynamical system allowing the number of spikes in approximately 20 ms to be utilized by neurons that receive a large number of inputs, and integrate over the currents injected into them by the afferent inputs. It is a further issue about how the cortical visual system could use processing occurring in approximately 15 ms per cortical stage to implement invariant object recognition. We have proposed for this a hierarchical series of networks with feedforward convergence from stage to stage, and competition implemented within each stage, and training of the feedforward synaptic connections using an associative synaptic modification rule with a short term memory trace to allow the network to learn invariant representations of objects, and shown that such a network could learn invariant representations (Elliffe, Rolls, & Stringer, 2002; Rolls, 1992; Rolls & Deco, 2002; Rolls & Milward, 2000; Wallis & Rolls, 1997). Within this system, the type of information transmission described here would be sufficiently fast, and the dynamical processing within each area to implement the competition through feedback inhibitory neurons sufficiently fast, to account for the speed of operation of the whole system, of approximately 15–17 ms per cortical area (Deco & Rolls, 2004; Rolls & Deco, 2002). The operation from stage to stage would be feedforward (with recurrent processing within each stage), with insufficient time given the evidence from backward masking just described for information to travel from the primary visual cortex (with a latency of approximately 40–50 ms) to the inferior temporal visual cortex (with a latency of approximately 90–100 ms), and then back again to V1 to influence further processing (Rolls, 2003, 2005a; Rolls & Tovee, 1994; Rolls et al., 1999, 1994). Instead, the top-down feedback processing is implicated in top-down attentional effects, and would take longer to implement in the dynamical system (Deco & Rolls, 2004, 2005a, 2005b, 2002).

The linear increase in the information with the number of neurons (Fig. 4c) is evidence that the neurons convey independent information. Further evidence for this is that the same population of neurons was shown to have weakly ergodic encoding, that is the single cell sparseness a^S has the same value as the population sparseness a^P (Franco, Rolls, Aggelopoulos, & Jerez, 2007). This occurs if the responses of the neurons to the set of stimuli are uncorrelated, that is each neuron is independently tuned to the set of stimuli (Lehky, Sejnowski, & Desimone, 2005). The linear increase in information with the number of stimuli and the weak ergodicity thus both provide evidence that the neuronal responses are uncorrelated, and this is potentially an important conclusion about the encoding of stimuli by these neurons. This linear increase indicates little redundancy between the spike counts from different neurons, and indicates that this is an efficient code, for the fact that the information increases linearly with the number of neurons indicates that the number of stimuli that can be encoded increases exponentially with the number of neurons. The spike count code described here is thus a powerful code. Moreover, much of the information in the type of code described here can be decoded by neurons that use a process as simple as producing a synaptically weighted sum of all their inputs in a short time (i.e. calculating a dot product between their input firing rate vector and their synaptic weight vector) (Franco et al., 2004; Rolls et al., 1997a).

In conclusion, the results described in this paper add to our developing understanding of how information is encoded in the visual system by showing that in the inferior temporal visual cortex considerable information is available from the first spike to arrive in response to a stimulus though this is especially when the first spike of a population of neurons is considered; that the order in which the spikes arrive from the different neurons does not appear to add significant information to that available from knowing that a spike has arrived from some but not other neurons; and that more information is available if all the spikes in a short time window are taken into account. The results thus provide evidence that is consistent with the number of spikes from different neurons in a short time being the way in which the major part of the information about which stimulus has been shown is present in neuronal activity in the inferior temporal visual cortex. It will be of interest if these approaches can be extended to other cortical areas.

Acknowledgments

This research was supported by a Wellcome Trust and Medical Research Council Grants to Professor E.T. Rolls. We acknowledge helpful discussions with Dr. Alessandro Treves (Trieste). L.F. acknowledges support through a Ramon y Cajal Fellowship from the Spanish Ministry of

Education and Science. L.F. and J.M.J. acknowledge support from Grant CICYT-TIN2005-02984 (FEDER).

References

Aggelopoulos, N. C., Franco, L., & Rolls, E. T. (2005). Object perception in natural scenes: encoding by inferior temporal cortex simultaneously recorded neurons. *Journal of Neurophysiology, 93*, 1342–1357.

Bacon-Mace, N., Mace, M. J., Fabre-Thorpe, M., & Thorpe, S. J. (2005). The time course of visual processing: backward masking and natural scene categorisation. *Vision Research, 45*, 1459–1469.

Baddeley, R. J., Abbott, L. F., Booth, M. J. A., Sengpiel, F., Freeman, T., Wakeman, E. A., et al. (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society of London B, 264*, 1775–1783.

Battaglia, F. P., & Treves, A. (1998). Stable and rapid recurrent processing in realistic auto-associative memories. *Neural Computation, 10*, 431–450.

Baylis, G. C., Rolls, E. T., & Leonard, C. M. (1987). Functional subdivisions of the temporal lobe neocortex. *Journal of Neuroscience, 7*, 330–342.

Booth, M. C. A., & Rolls, E. T. (1998). View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex. *Cerebral Cortex, 8*, 510–523.

Deco, G., & Rolls, E. T. (2003). Attention and working memory: a dynamical model of neuronal activity in the prefrontal cortex. *European Journal of Neuroscience, 18*, 2374–2390.

Deco, G., & Rolls, E. T. (2004). A neurodynamical cortical model of visual attention and invariant object recognition. *Vision Research, 44*, 621–644.

Deco, G., & Rolls, E. T. (2005a). Attention, short-term memory, and action selection: a unifying theory. *Progress in Neurobiology, 76*, 236–256.

Deco, G., & Rolls, E. T. (2005b). Neurodynamics of biased competition and co-operation for attention: a model with spiking neurons. *Journal of Neurophysiology, 94*, 295–313.

Deco, G., & Rolls, E. T. (2005c). Synaptic and spiking dynamics underlying reward reversal in orbitofrontal cortex. *Cerebral Cortex, 15*, 15–30.

Deco, G., & Rolls, E. T. (2006). Decision-making and Weber's Law: a neurophysiological model. *European Journal of Neuroscience, 24*, 901–916.

Deco, G., Rolls, E. T., & Horwitz, B. (2004). "What" and "where" in visual working memory: a computational neurodynamical perspective for integrating fMRI and single-neuron data. *Journal of cognitive neuroscience, 16*, 683–701.

Delorme, A., & Thorpe, S. J. (2001). Face identification using one spike per neuron: resistance to image degradations. *Neural Networks, 14*, 795–803.

Desimone, R. (1991). Face-selective cells in the temporal cortex of monkeys. *Journal of Cognitive Neuroscience, 3*, 1–8.

Elliffe, M. C. M., Rolls, E. T., & Stringer, S. M. (2002). Invariant recognition of feature combinations in the visual system. *Biological Cybernetics, 86*, 59–71.

Feigenbaum, J. D., & Rolls, E. T. (1991). Allocentric and egocentric spatial information processing in the hippocampal formation of the behaving primate. *Psychobiology, 19*, 21–40.

Franco, L., Rolls, E. T., Aggelopoulos, N. C., & Jerez, J. M. (2007). Neuronal selectivity, population sparseness, and ergodicity in the inferior temporal visual cortex.

Franco, L., Rolls, E. T., Aggelopoulos, N. C., & Treves, A. (2004). The use of decoding to analyze the contribution to the information of the correlations between the firing of simultaneously recorded neurons. *Experimental Brain Research, 155*, 370–384.

Judge, S. J., Richmond, B. J., & Chu, F. C. (1980). Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Research, 20*(6), 535–538.

Keysers, C., & Perrett, D. I. (2002). Visual masking and RSVP reveal neural competition. *Trends in cognitive sciences, 6*, 120–125.

Lehky, S. R., Sejnowski, T. J., & Desimone, R. (2005). Selectivity and sparseness in the responses of striate complex cells. *Vision Research, 45*, 57–73.

Nowak, L. G., & Bullier, J. (1997). The timing of information transfer in the visual system. In K. S. Rockland, J. H. Kaas, & A. Peters (Eds.). *Extrastriate visual cortex in primates* (12, pp. 205–241). New York: Plenum Press.

Panzeri, S., Rolls, E. T., Battaglia, F., & Lavis, R. (2001). Speed of feedforward and recurrent processing in multilayer networks of integrate-and-fire neurons. *Network: Computation in Neural Systems, 12*, 423–440.

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

Perrett, D. I., Rolls, E. T., & Caan, W. (1982). Visual neurons responsive to faces in the monkey temporal cortex. *Experimental Brain Research, 47*(3), 329–342.

Rolls, E. T. (1992). Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. *Philosophical Transactions of the Royal Society of London B, 335*, 11–21.

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

Rolls, E. T. (2003). Consciousness absent and present: a neurophysiological exploration. *Progress in Brain Research, 144*, 95–106.

Rolls, E. T. (2005a). Consciousness absent or present: a neurophysiological exploration of masking. In H. Ogmen & B. G. Breitmeyer (Eds.), *The First half second: The microgenesis and temporal dynamics of unconscious and conscious visual processes* (pp. 89–108). Cambridge, MA: MIT Press, chap. 106.

Rolls, E. T. (2005b). *Emotion explained*. Oxford: Oxford University Press.

Rolls, E. T., Aggelopoulos, N. C., Franco, L., & Treves, A. (2004). Information encoding in the inferior temporal cortex: contributions of the firing rates and correlations between the firing of neurons. *Biological Cybernetics, 90*, 19–32.

Rolls, E. T., Aggelopoulos, N. C., & Zheng, F. (2003a). The receptive fields of inferior temporal cortex neurons in natural scenes. *Journal of Neuroscience, 23*, 339–348.

Rolls, E. T., & Deco, G. (2002). *Computational neuroscience of vision*. Oxford: Oxford University Press.

Rolls, E. T., Franco, L., Aggelopoulos, N. C., & Reece, S. (2003b). An information theoretic approach to the contributions of the firing rates and correlations between the firing of neurons. *Journal of Neurophysiology, 89*, 2810–2822.

Rolls, E. T., & Milward, T. (2000). A model of invariant object recognition in the visual system: learning rules, activation functions, lateral inhibition, and information-based performance measures. *Neural Computation, 12*, 2547–2572.

Rolls, E. T., & Tovee, M. J. (1994). Processing speed in the cerebral cortex and the neurophysiology of visual masking. *Proceedings of the Royal Society of London B, 257*, 9–15.

Rolls, E. T., & Tovee, M. J. (1995). Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *Journal of Neurophysiology, 73*, 713–726.

Rolls, E. T., Tovee, M. J., & Panzeri, S. (1999). The neurophysiology of backward visual masking: information analysis. *Journal of Cognitive Neuroscience, 11*, 335–346.

Rolls, E. T., Tovee, M. J., Purcell, D. G., Stewart, A. L., & Azzopardi, P. (1994). The responses of neurons in the temporal cortex of primates, and face identification and detection. *Experimental Brain Research, 101*(3), 473–484.

Rolls, E. T., & Treves, A. (1998). *Neural networks and brain function*. Oxford: Oxford University Press.

Rolls, E. T., Treves, A., & Tovee, M. J. (1997a). The representational capacity of the distributed encoding of information provided by

populations of neurons in the primate temporal visual cortex. *Experimental Brain Research, 114*, 177–185.

Rolls, E. T., Treves, A., Tovee, M. J., & Panzeri, S. (1997b). Information in the neuronal representation of individual stimuli in the primate temporal visual cortex. *Journal of Computational Neuroscience, 4*, 309–333.

Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience, 19*, 109–139.

Thorpe, S., Delorme, A., & Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural Networks, 14*, 715–725.

Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature, 381*, 520–522.

Tovee, M. J., & Rolls, E. T. (1995). Information encoding in short firing rate epochs by single neurons in the primate temporal visual cortex. *Visual Cognition, 2*(1), 35–58.

Tovee, M. J., Rolls, E. T., & Azzopardi, P. (1994). Translation invariance in the responses to faces of single neurons in the temporal visual cortical areas of the alert macaque. *Journal of Neurophysiology, 72*(3), 1049–1060.

Tovee, M. J., Rolls, E. T., Treves, A., & Bellis, R. P. (1993). Information encoding and the responses of single neurons in the primate temporal visual cortex. *Journal of Neurophysiology, 70*(2), 640–654.

Treves, A. (1993). Mean-field analysis of neuronal spike dynamics. *Network, 4*, 259–284.

Treves, A., Panzeri, S., Rolls, E. T., Booth, M., & Wakeman, E. A. (1999). Firing rate distributions and efficiency of information transmission of inferior temporal cortex neurons to natural visual stimuli. *Neural Computation, 11*, 611–641.

VanRullen, R., Guyonneau, R., & Thorpe, S. J. (2005). Spike times make sense. *Trends in neurosciences, 28*, 1–4.

VanRullen, R., & Thorpe, S. J. (2001a). Is it a bird? Is it a plane? Ultra-rapid visual categorisation of natural and artifactual objects. *Perception, 30*, 655–668.

VanRullen, R., & Thorpe, S. J. (2001b). Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. *Neural Computation, 13*, 1255–1283.

Wallis, G., & Rolls, E. T. (1997). Invariant face and object recognition in the visual system. *Progress in Neurobiology, 51*, 167–194.