Exp Brain Res (2004) 155: 370–384
DOI 10.1007/s00221-003-1737-5

RESEARCH ARTICLES

Leonardo Franco · Edmund T. Rolls · Nikolaos C. Aggelopoulos · Alessandro Treves

The use of decoding to analyze the contribution to the information of the correlations between the firing of simultaneously recorded neurons

Received: 10 December 2002 / Accepted: 25 September 2003 / Published online: 13 January 2004

Abstract A new decoding method is described that enables the information that is encoded by simultaneously recorded neurons to be measured. The algorithm measures the information that is contained not only in the number of spikes from each neuron, but also in the cross-correlations between the neuronal firing including stimulus-dependent synchronization effects. The approach enables the effects of interactions between the 'signal' and 'noise' correlations to be identified and measured, as well as those from stimulus-dependent cross-correlations. The approach provides an estimate of the statistical significance of the stimulus-dependent synchronization information, as well as enabling its magnitude to be compared with the magnitude of the spike-count related information, and also whether these two contributions are additive or redundant. The algorithm operates even with limited numbers of trials. The algorithm is validated by simulation. It was then used to analyze neuronal data from the primate inferior temporal visual cortex. The main conclusions from experiments with two to four simultaneously recorded neurons were that almost all of the information was available in the spike counts of the neurons; that this Rate information included on average very little redundancy arising from stimulus-independent correlation effects; and that stimulus-dependent cross-correlation effects (i.e. stimulus-dependent synchronization) contribute very little to the encoding of information in the inferior temporal visual cortex about which object or face has been presented.

L. Franco · 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
Tel.: +44-1865-271348
Fax: +44-1865-310447

A. Treves
SISSA-Cognitive Neuroscience Sector,
Via Beirut 2–4,
34014 Trieste, Italy

Keywords Synchronization · Cross-correlation · Inferior temporal visual cortex · Temporal coding · Redundancy

Introduction

To analyze how neurons encode information about stimuli or other events, it is useful to apply information theory, because this allows the contributions of different possible factors (such as the number of spikes vs the relative timing of spikes from different cells) to be measured quantitatively and with the same metric (Shannon 1948; Cover and Thomas 1991; Rolls and Deco 2002). Simultaneously recorded neurons sometimes show cross-correlations in their firing, that is the firing of one cell is systematically related to the firing of the other cell. One example of this is neuronal response synchronization. The cross-correlation, to be defined below, shows the time difference between the cells at which the systematic relation appears. A significant peak or trough in the cross-correlation function could reveal a synaptic connection from one cell to the other, or a common input to each of the cells, or any of a considerable number of other possibilities. If the synchronization occurred for only some of the stimuli, then the presence of the significant cross-correlation for only those stimuli could provide additional evidence separate from any information in the firing rate of the neurons about which stimulus had been shown. Information theory in principle provides a way of quantitatively assessing the relative contributions from these two types of encoding, by expressing what can be learned from each type of encoding in the same units, bits of information.

When applying information theory to the responses of two or more simultaneously recorded neurons, the number of possible combinations of the relative times of the spikes of the different cells becomes very large. That is, the dimensionality of the space which must be filled adequately with real neurophysiological data to obtain reliable estimates of the information becomes so large that the information estimates become unreliable, and in fact are biased upwards. Even bias correction measures

(Panzeri and Treves 1996; Treves and Panzeri 1995) cannot completely correct for this amount of undersampling. In this situation the dimensionality of the space in which the neuronal responses are measured must be reduced. A recent approach to this issue has been to simply count the number of spikes in a single short time window from the simultaneously recorded cells, and to use these spike counts to estimate the information that is contributed by different factors, including factors such as synchronization of the spikes of different cells. This is the approach taken by Panzeri et al. (1999), Rolls et al. (2003a), and Rolls et al. (2003b). However, this approach is inherently limited to a small number of cells and only one to two spikes from each cell, because otherwise the Taylor expansion used breaks down (Bezzi et al. 2002). Moreover, the Taylor expansion approach we used measured the information from stimulus-dependent co-modulation of firing rates and not from synchronization that is not reflected in co-modulation, and when the approach was extended to measure synchronization effects more directly (Panzeri et al. 2001), it was still limited to small numbers of spikes and small numbers of neurons, and needed large numbers of trials.

In this paper we develop a new approach to the measurement of the information conveyed by simultaneously recorded cells. The approach can be applied in principle to any number of cells, from a small number to very large numbers, and the time period in which the information is measured is not limited, and measures stimulus-dependent synchronization effects even when they are not reflected in co-modulation of firing rates. The method uses a decoding procedure to estimate from the neuronal data the probability that each stimulus in a set was shown on a given trial, and then measures the mutual information between the estimated stimulus and the actual stimulus that was shown on that trial (Rolls et al. 1997). The advantage of the decoding approach is that the dimensionality of the space over which the information is measured is reduced to depend on the number of stimuli used, and not on the number of stimuli, the number of cells, and the number of responses for each stimulus that are required by the direct method without decoding. In practice, twice as many trials as there are stimuli are required for the decoding approach (Rolls et al. 1997). The decoding approach used here is developed from the method described by Rolls et al. (1997), which was developed to measure the information contained in the firing rates of multiple neurons. The method is extended here to incorporate information present in the cross-correlations between the neurons. We not only describe the approach here, but show how it can be used to separate stimulus-independent and stimulus-dependent effects con-

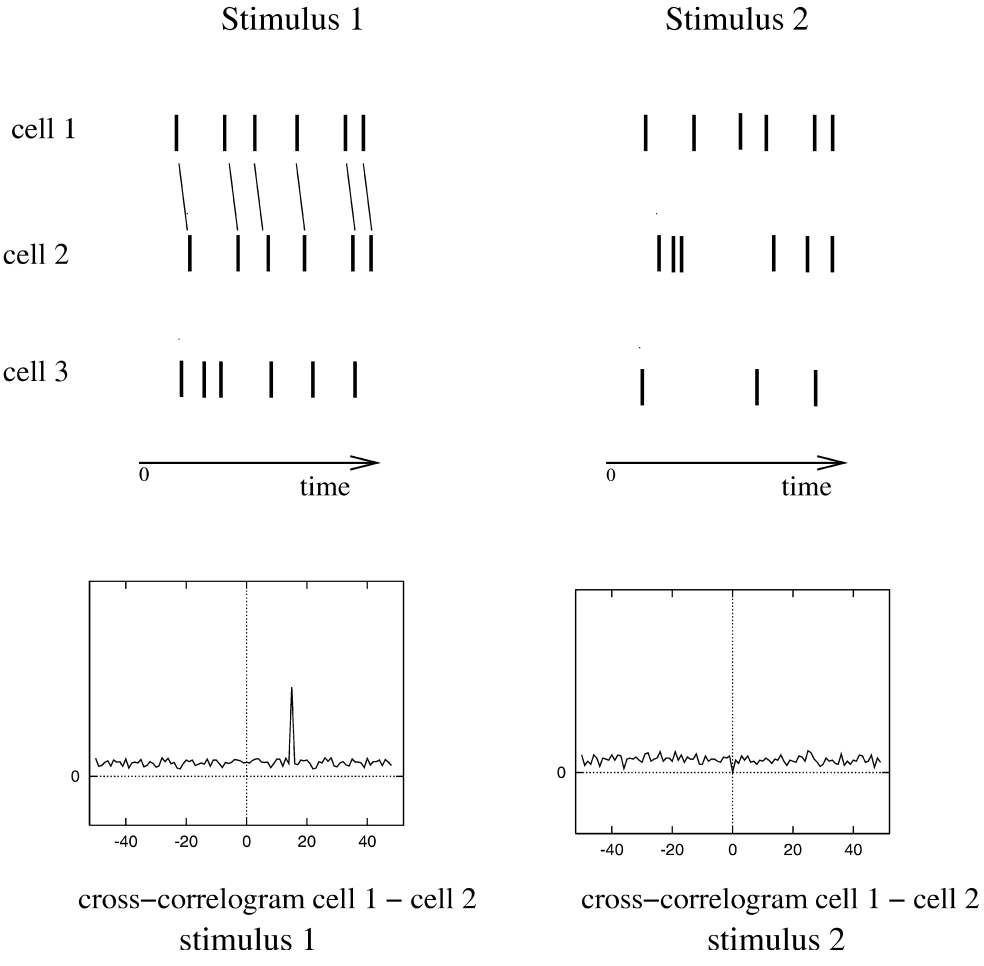

Fig. 1 Illustration of the information that could be carried by spike trains. The responses of three cells to two different stimuli are shown on one trial. Cell 3 reflects which stimulus was shown in the number of spikes produced, and this can be measured as spike count or rate information. Cells 1 and 2 have no spike count or rate information, because the number of spikes is not different for the two stimuli. Cells 1 and 2 do show some synchronization, reflected in the cross-correlogram, that is stimulus dependent, as the synchronization is present only when stimulus 1 is shown. The contribution of this effect is measured as the stimulus-dependent synchronization information

tributed by the correlations between the neurons, show how the statistical significance of these different contributions can be measured, validate the method with simulated neuronal data produced to contain information contributed in different ways by the spikes of the different neurons, and evaluate the power efficiency of the method. We then show how the approach can be applied to real neuronal data from the macaque inferior temporal visual cortex. For the neurophysiology, the stimuli were a set of eight objects, faces, and scenes presented while the monkey performed a visual discrimination task. If synchronization was being used to bind the parts of each object into the correct spatial relationship to other parts, this might be expected to be revealed by stimulus-dependent cross-correlations in the firing of simultaneously recorded groups of two to four cells using multiple single-neuron microelectrodes.

There have been previous analyses of the contributions of stimulus-independent covariations between neurons (Gawne and Richmond 1993; Abbott and Dayan 1999; Lee et al. 1998; Sompolinsky et al. 2001; Reich et al. 2001), and the current paper provides new ways to estimate their contributions from the limited numbers of trials of data that are frequently available in neurophysiological experiments. Moreover, we note that there is current great interest in the possible contribution to neural encoding of stimulus-dependent correlations including synchronization (Singer 2000), and it is in this context that we have developed the new approach described here. Decoding methods for neuronal data have been studied by a number of investigators (Pouget et al. 1998; Salinas and Abbott 1994), and have been applied to the measurement of information from neuronal populations (Robertson et al. 1999; Nirenberg et al. 2001; Hatsopoulos et al. 1998; Oram et al. 1998, 2001; Dan et al. 1998; Panzeri et al. 1998, 1999; Rolls et al. 1997), but this is the first paper that uses decoding methods to estimate stimulus-dependent contributions from cross-correlations between the firing of neurons.

Materials and methods

The information measurement algorithm

Figure 1 illustrates how synchronization occurring only for some of the stimuli could be used to encode information about which stimulus was presented. In the figure the spike trains of three neurons are shown after the presentation of two different stimuli on one trial. As shown by the cross-correlogram in the lower part of the figure, the responses of cell 1 and cell 2 are synchronized when stimulus 1 is presented, as whenever a spike from cell 1 is emitted, another spike from cell 2 is emitted after a short time lag. In contrast, when stimulus 2 is presented, synchronization effects do not appear. Thus, based on a measure of the synchrony between the responses of cells 1 and 2, it is possible to obtain some information about what stimulus has been presented. The contribution of this effect is measured as the stimulus-dependent synchronization information. Cells 1 and 2 have no information about what stimulus was presented from the number of spikes, as the same number is found for both stimuli. Cell 3 carries information in the spike count in the time window (which is also called the firing rate) about what stimulus was presented. (Cell 3 emits six spikes for stimulus 1 and three spikes for stimulus 2.)

The example shown in Fig. 1 is for the neuronal responses on a single trial. Given that the neuronal responses are variable from trial to trial, we need a method to quantify the information that is gained from a single trial of spike data in the context of the measured variability in the responses of all of the cells, including how the cells' responses covary in a way which may be partly stimulus-dependent, and may include synchronization effects. The direct approach is to apply the Shannon mutual information measure (Shannon 1948; Cover and Thomas 1991):

$$I(s, \mathbf{r}) = \sum_{s \in S} \sum_{\mathbf{r}} P(s, \mathbf{r}) \log_2 \frac{P(s, \mathbf{r})}{P(s)P(\mathbf{r})}, \quad (1)$$

where $P(s, \mathbf{r})$ is a probability table embodying a relationship between the variable s (here, the stimulus) and $\mathbf{r}$ (a vector where each element is the firing rate of one neuron).

However, because the probability table of the relation between the neuronal responses and the stimuli, $P(s, \mathbf{r})$ is so large (given that there may be many stimuli, and that the response space which has to include spike timing is very large), in practice it is difficult to obtain a sufficient number of trials for every stimulus to generate the probability table accurately, at least with data from mammals in which the experiment cannot usually be continued for many hours of recording from a whole population of cells. To circumvent this undersampling problem, Rolls et al. (1997) developed a decoding procedure, in which an estimate (or guess) of which stimulus (called s') was shown on a given trial is made from a comparison of the neuronal responses on that trial with the responses made to the whole set of stimuli on other trials. One then obtains a conjoint probability table $P(s, s')$, and then the mutual information based on probability estimation (PE) decoding (I_p) between the estimated stimuli s' and the actual stimuli s that were shown can be measured:

$$\langle I_p \rangle = \sum_{s \in S} \sum_{s' \in S} \mathrm{P}(s, s') \log_2 \frac{\mathrm{P}(s, s')}{\mathrm{P}(s)\mathrm{P}(s')} \quad (2)$$

$$\langle I_p \rangle = \sum_{s \in S} \mathrm{P}(s) \sum_{s' \in S} \mathrm{P}(s'|s) \log_2 \frac{\mathrm{P}(s'|s)}{\mathrm{P}(s')}. \quad (3)$$

These measurements are in the low dimensional space of the number of stimuli, and therefore the number of trials of data needed for each stimulus is of the order of the number of stimuli, which is feasible in experiments. In practice, it is found that for accurate information estimates with the decoding approach, the number of trials for each stimulus should be at least twice the number of stimuli (as shown in Fig. 8).

Decoding procedures

The nature of the decoding procedure is illustrated in Fig. 2. The left part of the diagram shows the average firing rate (or equivalently spike count) responses of each of three cells (labelled as Rate Cell 1, 2, 3) to a set of three stimuli. The last row (labelled Response single trial) shows the data that might be obtained from a single trial and from which the stimulus that was shown (St. ?) must be estimated or decoded, using the average values across trials shown in the top part of the table, and the probability distribution of these values. The decoding step essentially compares the vector of responses on trial St. ? with the average response vectors obtained previously to each stimulus. This decoding can be as simple as measuring the correlation, or dot (inner) product, between the test trial vector of responses and the response vectors to each of the stimuli. This procedure is very neuronally plausible, in that the dot product between an input vector of neuronal activity and the synaptic response vector on a single neuron (which might represent the

average incoming activity previously to that stimulus) is the simplest operation that it is conceived that neurons might perform (Rolls and Deco 2002; Rolls and Treves 1998). Other decoding procedures include a Bayesian procedure based on a Gaussian or Poisson assumption of the spike count distributions as described in detail by Rolls et al. (1997). The Gaussian one is what is used throughout this paper, and it is described below. The new step taken in this paper is to introduce into the table data(s, $\mathbf{r}$) shown in the upper part of Fig. 2 new columns, shown on the right of the diagram, containing a measure of the cross-correlation (averaged across trials in the upper part of the table) for some pairs of cells (labelled as Corrln Cells 1–2 and 2–3). The decoding procedure can then take account of any cross-correlations between pairs of cells, and thus measure any contributions to the information from the population of cells that arise from cross-correlations between the neuronal responses. If these cross-correlations are stimulus-dependent, then their positive contribution to the information encoded can be measured. This is the new concept for information measurement from neuronal populations introduced in this paper. We describe below how the cross-correlation information can be introduced into the table, and then how the information analysis algorithm can be used to measure the contribution of different factors in the neuronal responses to the information that the population encodes.

Further details of the decoding procedures are as follows (see also Rolls et al. 1997). 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. 4) in a cross-validation procedure described by Rolls et al. (1997).

$$P(s'|\mathbf{r}) = \frac{P(\mathbf{r}|s')P(s')}{P(\mathbf{r})}. \tag{4}$$

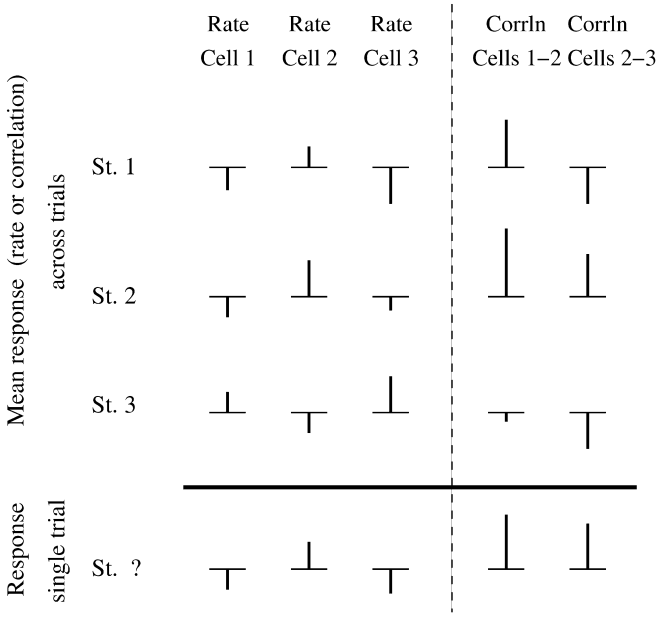

Fig. 2 The left part of the diagram shows the average firing rate (or equivalently spike count) responses of each of three cells (labelled as Rate Cell 1, 2, 3) to a set of three stimuli. The right two columns show a measure of the cross-correlation (averaged across trials) for some pairs of cells (labelled as Corrln Cells 1–2 and 2–3). The last row (labelled Response single trial) shows the data that might be obtained from a single trial and from which the stimulus that was shown (St. ?) must be estimated or decoded, using the average values across trials shown in the top part of the table. From the responses on the single trial, the most probable decoded stimulus is stimulus 2, based on the values of both the rates and the cross-correlations

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:

$$P(\mathbf{r}) = \sum_{s'} P(\mathbf{r}|s')P(s'). \tag{5}$$

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 $P(s')\Pi_c P(r_c|s')$, where r_c is the firing rate of cell c. We note that $P(r_c|s')$ is derived from the responses of cell 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 $P(r_c|s')$.[1] 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',

$$P(s'|s) = \sum_{\mathbf{r}\in \text{test}} P(s'|\mathbf{r})P(\mathbf{r}|s). \tag{6}$$

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 $P^R{}_N(s, s')$ 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 in Eq. 3.

We also generated a second (Frequency) table $P^F{}_N(s, s^p)$ from the fraction of times an actual stimulus s elicited a response that led to a predicted (single most likely) stimulus s^p. From this probability table the mutual information (I_{ml}) measure based on maximum likelihood decoding was calculated with Eq. 7:

$$\langle I_{ml}\rangle = \sum_{s\in S}\sum_{s^p\in S} \mathrm{P}(s, s^p)\log_2 \frac{\mathrm{P}(s, s^p)}{\mathrm{P}(s)\mathrm{P}(s^p)}. \tag{7}$$

A detailed comparison of maximum likelihood and probability decoding is provided by Rolls et al. (1997), but we note here that probability estimate decoding is more regularized (see below) and therefore may be safer to use when investigating the effect on the information of the number of cells. For this reason, the results

[1] When using the Gaussian, the probabilities of 0, 1, 2, 3, etc., spikes were estimated as follows for each stimulus. The probability of zero spikes was obtained directly by the proportion of trials that had 0 spikes. The mean and standard deviation of the positive part of the Gaussian were computed from the remaining spike counts. We note that because the spike counts are approximately Poisson distributed, the variance increases in proportion to (and equals) the mean. A consequence of this is that the mean is located one standard deviation above zero (assuming that the Poisson is a good fit). Thus any inaccuracies due to truncating the fitted Gaussian below 0 are small, because only a small fraction of the data lie more than one standard deviation below the mean. In practice, this truncated Gaussian was chosen over the Poisson distribution (with an additional weight at r_c=0), because we have found previously (Rolls et al. 1997) and with the present data set that with our neuronal populations the Gaussian fit produces slightly higher values for both percentage correct and information. However, the fact that the performance with the Poisson fit was almost as good as the truncated Gaussian fit indicates that truncation per se is probably not a major issue. The indication that the Poisson fit does not work quite as well as the Gaussian fit with our data reflects the fact that the variability of the spike counts does not fit a Poisson distribution perfectly, and the spike count distributions can be fitted better using the two parameters provided by the Gaussian fitting procedure. The theoretical advantages of using different types of decoding for the spike counts are considered in the "Discussion."

described in this paper were obtained with probability estimation (PE) decoding. The maximum likelihood decoding does give an immediate measure of the percentage correct.

Another approach to decoding is the dot product (DP) algorithm which computes the normalized dot products between the current firing vector **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 s^p) for the probability table $P(s, s^p)$. (It can also be used to provide percentage correct measures.)

We note that any decoding procedure can be used in conjunction with information estimates both from the full probability table (to produce I_p), and from the most likely estimated stimulus for each trial (to produce I_{ml}).

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. (1997) 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. (1997) 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 twice 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. 1997). We further note that the bias correction term becomes very small when more than ten cells are included in the analysis (Rolls et al. 1997).

We note that if Bayesian decoding is used an assumption is that the joint probability distribution of the spike count responses of the cells is approximated by the product of the separate probability distributions for each cell. This approximation holds if the distributions are independent, and may be less exact if there are correlations between the neurons' responses. In practice this is not a limitation of the method in that the level of correlations found in practice produce only a relatively small distortion of the probability values used to compute the information, partly because these probability values are normalized before being used, reducing the distortion especially when relatively few (e.g. 20) trials of data per stimulus are used. We also note that trying to estimate the parameters for the joint probability distribution would require a very large number of trials of data (Gill et al. 1981). We note that the approximation in any case does not apply to dot product decoding, and the fact that qualitatively similar results are obtained with both types of decoding is consistent with the hypothesis that the Bayesian decoding works satisfactorily, as considered further in the "Discussion".

Response quantification

The data from the neuronal activity that was entered into the table data(*s,* **r**) shown in the upper part of Fig. 2 was as follows.

From the response of each cell *c* to each stimulus, we extracted a single mean spike count in a fixed time window (or firing rate, $\mathbf{r}_c$ expressed in spikes per second). From these spike counts, the algorithm measured the information in the firing rates, and in any co-modulation of the firing rates of neurons.

We also introduced a measure of the synchronization between pairs of neurons into the data table so that the information available in the synchronization could be measured. The measure of the synchronization that was introduced into the table data(*s,* **r**) on each trial was the value of the Pearson cross-correlation coefficient calculated for that trial at the appropriate lag for cell pairs that had significant cross-correlations. This value of this Pearson cross-correlation coefficient for a single trial was calculated from pairs of spike trains on a single trial by forming for each cell a vector of 0's and 1's, the 1's representing the time of occurrence of spikes with a temporal resolution of 1 ms. Resulting values within the range 1 to −1 were shifted to obtained positive values. An advantage of basing the measure of synchronization on the Pearson cross-correlation coefficient is that it measures the amount of synchronization between a pair of neurons independently of the firing rate of the neurons. The lag at which the cross-correlation measure was computed for every single trial, and whether there was a significant cross-correlation between neuron pairs, was identified from the location of the peak in the cross-correlogram taken across all trials. The cross-correlogram was calculated by, for every spike that occurred in one neuron, incrementing the bins of a histogram that corresponded to the lag times of each of the spikes that occurred for the other neuron. The raw cross-correlogram was corrected by subtracting the 'shift predictor' cross-correlogram (which was produced by random re-pairings of the trials), to produce the corrected cross-correlogram. It was normalized to be in the range ±1. Examples of the cross-correlograms calculated across all stimuli (and used to define the appropriate lag), and for each stimulus, are shown in Fig. 4.

The exact measure of the synchronization of the firing of each pair of cells that is used is not critical for the approach, and indeed similar results are obtained if the sum of the three synchronization values described above centered at the appropriate lag in the cross-correlogram is used. The number of these values that are taken can be altered to detect the timing precision within which spikes are counted as being synchronized or not. In practice, we used a precision of ±1 ms in this paper. Different datasets might benefit from different measures of the co-variation (Aertsen et al. 1989; Brody 1999; Konig 1994). The cross-correlation values entered into the table data(*s,* **r**) were also scaled so that the maximum correlation range from any cell pair had the same value as the maximum spike count range in the table. The rationale for this was that the spike counts and correlations could in principle contribute on an equal basis, and in practice it was found that the algorithm was little affected by the exact ratio of this scaling. The decoding procedure was applied to the full table of spike rates and correlation measures shown in Fig. 2. The correlation measures obtained in this investigation from the neurophysiological recordings had an approximately Poisson distribution (as also observed by Hatsopoulos et al. 1998), which was approximated by the truncated Gaussian distribution described under "Decoding procedures" because the way this was calculated fit the distribution well, and could also be efficiently applied to the spike rate values in the table.

Although not used by the decoding algorithm, we define below the terms 'signal' and 'noise' correlations, as they are useful in understanding the encoding of information by groups of simultaneously recorded cells, and the algorithm we describe measures the influences of these correlations.

The correlations in the mean responses of the neurons across the set of stimuli (sometimes called 'signal' correlations) ν

ν_{ij} can be thought of as the degree of similarity in the mean response profiles (averaged across trials) of the cells i and j to different stimuli. ν_{ij} is sometimes called the 'signal' correlation (Gawne and Richmond 1993; Shadlen and Newsome 1994, 1998). It is defined by:

$$\nu_{ij} = \frac{\langle \bar{r}_i(s)\bar{r}_j(s)\rangle_s}{\langle \bar{r}_i(s)\rangle_s \langle \bar{r}_j(s)\rangle_s} - 1, \tag{8}$$

where $\bar{r}_i(s)$ is the mean rate of response of cell i to stimulus s over all the trials in which that stimulus was present. It can vary from –1 to

∞ ($< \ldots >_s$ indicates the ensemble average over the s stimuli). The similarity of the mean response profiles can also be measured by the Pearson correlation coefficient, r.

The correlations in the neuronal response variability from the average to each stimulus (sometimes called 'noise' correlations) γ:

$$\gamma_{ij}(s) = \frac{\overline{r_i(s) r_j(s)}}{(\bar{r}_i(s)\bar{r}_j(s))} - 1. \quad (9)$$

This has been called the 'noise' correlation (Gawne and Richmond 1993; Shadlen and Newsome 1994, 1998) because it reflects the trial by trial co-variation in the responses of the neurons, and is also called the 'scaled cross-correlation density' (Aertsen et al. 1989; Panzeri et al. 1999). It can vary from –1 to ∞; negative values of $\gamma_{ij}(s)$ indicate anticorrelation, whereas positive values of $\gamma_{ij}(s)$ indicate correlation.

Generation of test data

Simulations of neuronal activity with defined types of information in the spike trains were made to provide data to evaluate the information measurement algorithm, and to demonstrate how it can be used to identify the contributions of different factors. The spike trains were generated with populations of integrate-and-fire neurons using a procedure similar to that of Shadlen and Newsome (1998) and as described by Rolls et al. (2003b). Each cell (in a population which was generally 10, but was altered as described below for individual tests) received 300 excitatory and 300 inhibitory inputs, each a Poisson process in itself, whose (possibly stimulus-dependent) mean rate was constant across the set of inputs for any specific stimulus condition, and contributed a fixed quantity to the membrane potential. When the membrane potential exceeded a threshold, it was reset to a baseline value, and a spike was emitted. Common inputs were provided when appropriate by connecting 50% of the inputs of the cells to the same input source. When stimulus-dependent correlations were required, these were generated by providing a percentage of shared connections within the receiving population of either 0% or 90% for different stimuli.

In cases where the neurons had different firing rates to each of the stimuli, we set the probability distribution of the firing rates to the different stimuli to be exponential, as this is an approximation to what real neurons exhibit (Treves et al. 1999; Baddeley et al. 1997). That is, for real neurons, the firing rate is high for a few stimuli, and increasingly low for further stimuli (as illustrated in Fig. 6).

Neurophysiological method

The responses of single neurons in the temporal cortical visual areas were measured to a set of ten 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 et al. 1997, 2003; Booth and Rolls 1998). The stimuli included $S=8$ images of objects, faces, natural scenes of the type that produce differential responses from inferior temporal cortex neurons, and examples of which have been illustrated previously (Rolls and Tovee 1995). The set of stimuli were shown once in random order, then 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 (though not orthogonally; see Rolls and Tovee 1995; Rolls and Deco 2002) to the set of stimuli used. 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 (USA) Discovery software. All procedures, including preparative and subsequent ones, were carried out in accordance with the NIH Principles of laboratory animal care (NIH publication No. 86–23, revised 1985), the guidelines of The Society for Neuroscience, and were licenced under the UK Animals (Scientific Procedures) Act, 1986. The sites of the neuronal recordings included in this investigation were in the cortex in area TE (Rolls et al. 2003a).

Results

Tests with simulated data

Using simulated data we tested the operation of the algorithm in different key cases in which different factors contributed to the information available in the neuronal responses.

Information in the stimulus-dependent correlations

The spike trains were generated by an integrate-and-fire simulation in which we simulated a correlational assembly with a constant firing rate of 20 spikes/s to all stimuli, and a percentage of shared connections of either 0% or 90% for different stimuli. There were ten cells in the assembly, and for each of the ten stimuli one pair of cells had common connections. In this case there were 30 trials for each stimulus.

In Fig. 3, the results of applying the information analysis are shown. There were 45 cross-correlograms between the cell pairs ($0.5\boldsymbol{n}_c(\boldsymbol{n}_c-1)$ where $\boldsymbol{n}_c$ is the number of cells), but many of them were close to zero and not significant. The ten cell pairs with the most significant cross-correlations were provided for the decoding algorithm to use, and for this analysis, the rate entries in the table shown in Fig. 2 were not used. Examples of the cross-correlograms for one cell pair are shown in Fig. 4, where it is possible to see that the responses of the cells are correlated for stimulus number 6 and that there are no significant correlations for the rest of the stimuli. Selecting a subset of the cross-correlations can be performed if there are a very large number of them and may help the decoding algorithm by reducing the noise contributed by low values. However, this is not an essential step, and indeed similar results are obtained if this selection of cross-correlation measures is not performed. Figure 3 shows that the information grows sublinearly with the number of cross-correlation pairs used by the algorithm from the table data(s, $\mathbf{r}$) shown in Fig. 2. The sublinear increase reflects the fact that the ten cross-correlation values do not contribute totally independently to the information measured by the algorithm. Figure 4 also shows that shuffling the trials within a stimulus results in no measured information (dashed line), showing that the stimulus-dependent cross-correlation information detected by the algorithm (solid line) does reflect the correlations

between cell pairs that are produced because the spikes are correlated within individual trials.

To provide a statistical measure of the significance of the information derived from stimulus-dependent correlations that are detected by the algorithm, we used the trial shuffling in a Monte Carlo procedure to measure the mean and the standard deviation of the information that could arise by chance from random pairings of spikes from different trials. The mean information and ±2 standard deviations of these values are shown by the dashed line in Fig. 3 (left). From this type of display, the stimulus-dependent information measured by the algorithm (solid line) is assessed as being significant if it is more than for example two times this standard deviation of the measure from the mean value obtained with the shuffling (dashed line). (The values obtained by the Monte Carlo procedure are approximately normally distributed.) This statistical analysis provides a useful check when only few trials of data are available so that the cross-correlation measures may be noisy. The mean of the value obtained in the Monte Carlo procedure is subtracted from the raw information detected. The power efficiency of the method is shown later (in Fig. 8). To check that the algorithm can accurately measure the information that is related to stimulus-dependent synchronization, we show in Fig. 3 (right) that when a defined amount of information is present in the synchronization while there is no information in the rates, then all of this information (2 bits in this case) is detected by the algorithm. [In this test data, there were four cells, four stimuli, and each cell gave four spikes to each stimulus. From the ten possible pairs of cells, four pairs had synchronization that was fully correlated (to within ±1 ms) for one of the stimuli.] The algorithm correctly detected no information in the spike counts, and no stimulus-independent contributions to the amount of information measured. The total information measured was two bits, all from the stimulus-dependent synchronization.

Information and redundancy arising from correlated response profiles (i.e. positive 'signal' correlations) and with no 'noise' correlations

Here we treat the case of cells with correlations in the firing rate response profiles across the set of stimuli, i.e. with positive signal correlations (measured by ν in Panzeri et al. 1999; Rolls et al. 2003b). (If all the neurons had the same tuning to the set of stimuli, and the cells were noise free, the cells would be completely redundant. If there is some noise, the cells would be to some extent redundant. The noise referred to here is the trial-by-trial variability, called the 'noise' correlation (and measured by γ in Panzeri et al. 1999; Rolls et al. 2003b), and in this subsection the noise we consider is just random noise which is uncorrelated across the cells.

We analyzed a case where the ten simulated neurons had Poisson spike trains, and their response profiles had a mean correlation of ν=0.15, r=0.85. The information available from only the spike counts is shown (that is, there are no entries in the cross-correlation columns of the table data(s, **r**) shown in Fig. 2). The redundancy is reflected in the less than linear increase in the information as a function of the number of cells, and compares with the linear increase shown in the same figure for cells with no correlation in their response profiles.

As shown in Fig. 5, shuffling the trials (within each stimulus) does not influence the information, and this

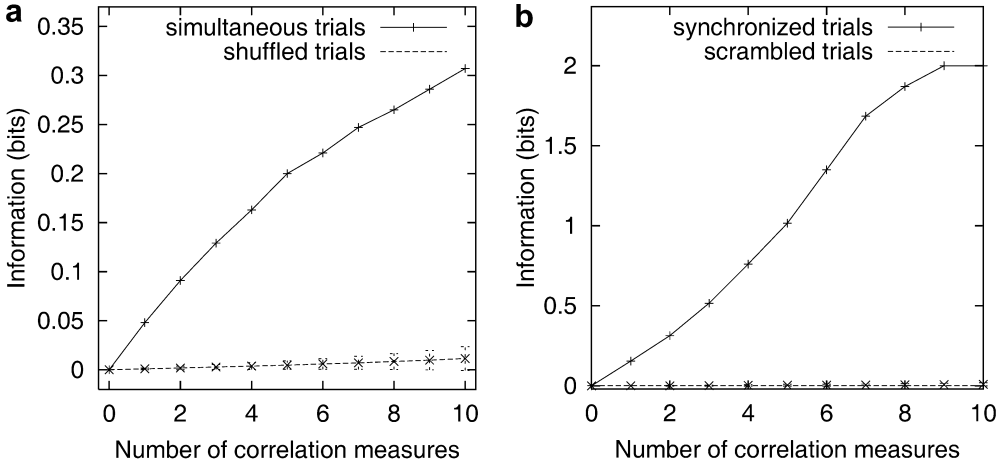

Fig. 3 *Left* The values of the information available from stimulus-dependent synchronization between neurons. The information available from up to ten selected cross-correlations between the responses of ten neurons is shown. The only information plotted is that from the cross-correlations, with no contributions assessed from the numbers of spikes. The dashed line shows the results of a control in which the spike data from the neurons is randomly shuffled between different trials for each stimulus before the information analysis, so that any stimulus-dependent synchronization effects will be lost. The error bars show ±2 standard deviations of the information estimates obtained on different shufflings in a Monte Carlo procedure used to estimate the variability of the information estimate. *Right* The values of the information from stimulus-dependent synchronization between neurons when there were 2 bits of information available from this in a test data set, and no information from the firing rates about which stimulus was presented. In this test data, there were four cells, four stimuli, and each cell gave four spikes to each stimulus. From the ten possible pairs of cells, four pairs had synchronization that was fully correlated (to within ±1 ms) for one of the stimuli

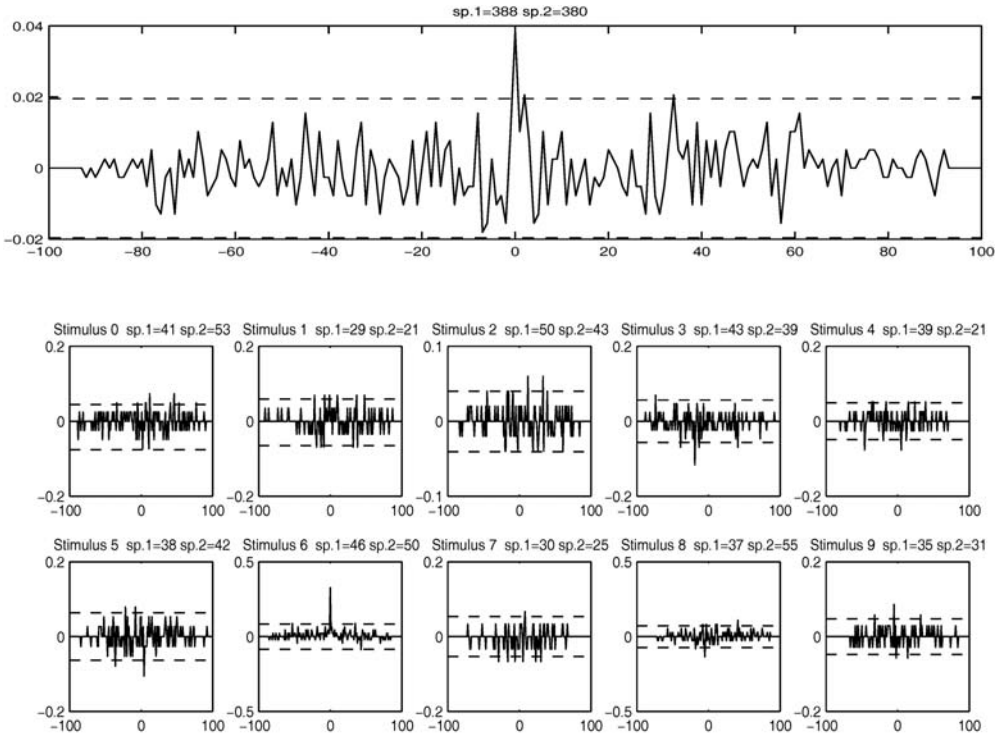

Fig. 4 The cross-correlograms from one pair of cells to the different stimuli used to generate the data analyzed in Fig. 3. Above we show the cross-correlogram calculated across all stimuli. Below we show the cross-correlograms for each stimulus. A significant cross-correlation was available between the responses of this pair of cells for stimulus 6 at approximately 0 ms lag, and the value of the cross-correlation at this lag was used in the measure of the magnitude of the cross-correlation by the decoding algorithm. The dashed lines show the 95% confidence intervals of the cross-correlation estimate. The cross-correlograms are noisy because we chose to use a limited number of trials for each stimulus (25) in order to illustrate the application of the approach to real data, for which the number of trials of data available may be limited

reflects the zero 'noise' correlation, that is that the neurons have independent variability in their spike trains. The lack of effect of the shuffling on the information estimated with the algorithm is the important evidence it provides that the spike trains of the different neurons are independent, i.e. that the 'noise' correlation measuring whether neurons have co-variability, whether it is stimulus-dependent or not, is zero. Such a case might arise when cells have no common input.

Information and redundancy arising from correlated response profiles (i.e. positive 'signal' correlations) with positive 'noise' correlations

We analyzed a case where the simulated integrate-and-fire cells share common input, generating correlated noise (that is trial-by-trial variability) for all stimuli. The correlated noise might arise from common inputs, and is stimulus-independent. (The spike trains were generated with integrate-and-fire neurons using a procedure similar to that of Shadlen and Newsome (1998) and as described by Rolls et al. (2003b) and in "Materials and methods." The common input was added by connecting 50% of the inputs of both cells to the same input source. In no case in this paper is there a ceiling effect to the information that can be extracted by the algorithm, in that the information available does not approach the maximum that would be needed to code for the ten stimuli, i.e. 3.32 bits (Rolls et al. 1997).

In Fig. 6 the values of the information available from the rates of ten simulated neurons about what stimuli have been presented are shown when the response profiles of the cells are highly correlated (with a mean Pearson correlation coefficient across pairs of cells of r=0.85). In this case, the information does not grow linearly with the number of cells, due to the redundancy in the profiles. However, in this case the information that is available from the ten cells is not as great as in Fig. 5 and this reduction is due to the correlated noise arising from the 50% of common inputs. This statement is confirmed by the fact that shuffling the trials (within a stimulus) as shown in Fig. 6 raises the curve to that expected with independent firing (which was shown in Fig. 5). The interactions that occur between the 'signal' and 'noise' correlations are described in the "Discussion" with the help of Fig. 9.

Information and redundancy arising from anti-correlated response profiles (i.e. negative 'signal' correlations) with positive 'noise' correlations

In Fig. 7 the values of the information available from the rates of ten simulated neurons about what stimuli have been presented are shown, in a case where the response profiles of the cells are tuned to different stimuli, with an average 'signal' correlation of ν=−0.018 (r=−0.09). The correlated noise arises from 50% of common inputs, and is stimulus-independent. In this case the information increases approximately linearly with the number of cells (because the anti-correlation is low, and indeed cannot be made on average to have a large negative value with a set of more than a very few stimuli). However, in this case shuffling the trials (which removes the effect of the correlated noise) results in less information being

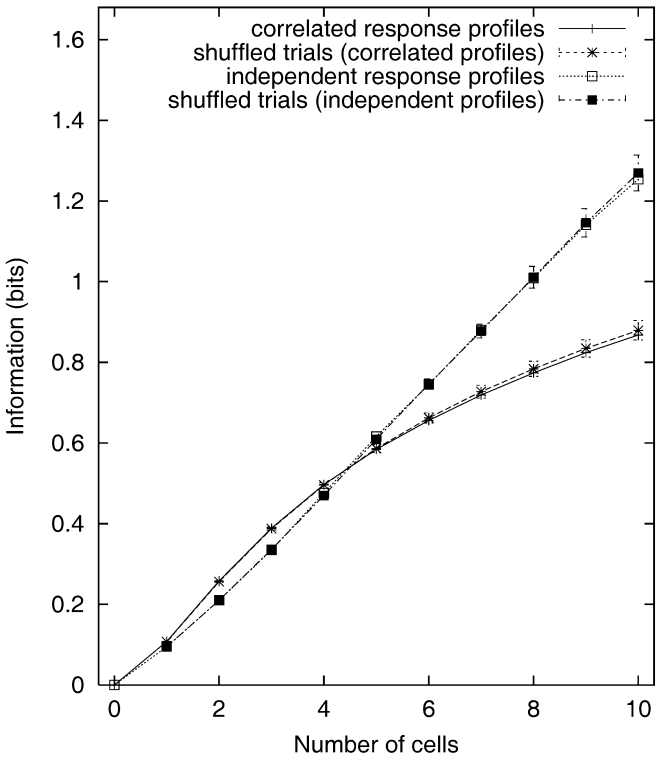

Fig. 5 The effect of correlated response profiles of neurons, with uncorrelated variability on a trial-by-trial basis (i.e. with positive 'signal' correlation and with zero 'noise' correlation). The information available from only the spike counts is shown. The redundancy is reflected in the less than linear increase in the information as a function of the number of cells. Shuffling the trials (within each stimulus) does not influence the information reflecting the zero 'noise' correlation. The mean Pearson correlation between the response profiles was r=0.85 (ν=0.15). Also shown (dashed line) is the case when the response profiles, i.e. 'signal' correlations, are approximately zero (r=0.009, ν=0.002)

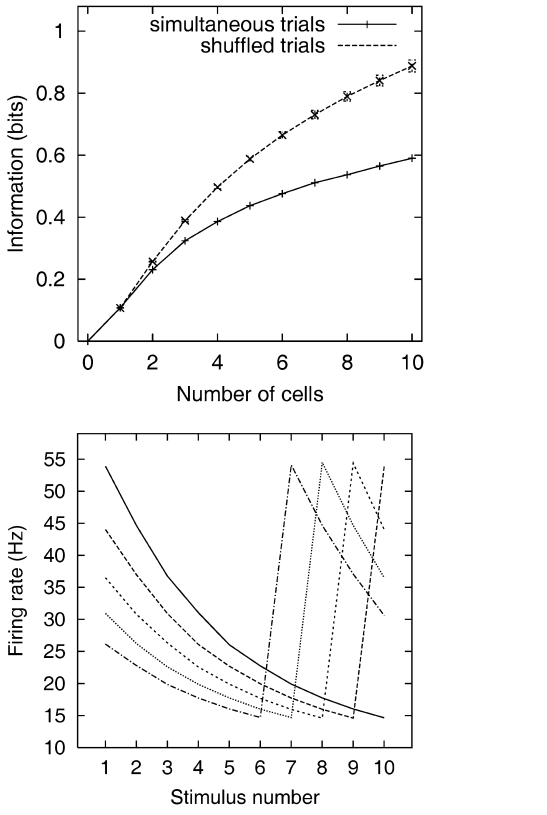

Fig. 6 *Above* The effect of correlated response profiles of neurons, with correlated variability on a trial-by-trial basis (i.e. with positive 'signal' correlation and positive 'noise' correlation). The information available from only the spike counts is shown. The redundancy is reflected in the less than linear increase in the information as a function of the number of cells, as in Fig. 5. Shuffling the trials (within each stimulus) increases the information, showing that the simultaneously recorded unshuffled spike count data reflect in addition a second type of redundancy, due to the positive 'noise' correlations (when they occur with positive 'signal' correlations). The mean Pearson correlation between the response profiles was r=0.85 (ν=0.15). Below: the response profiles of neurons 1, 3, 5, etc. to the ten discrete stimuli. The correlated noise was introduced by providing the ten neurons with 50% of common input

extracted. The gain of information, in this case approximately 0.06 bits, in the unshuffled condition is thus synergy which arises when the 'signal' and 'noise' correlations have the opposite sign (see "Discussion" and Panzeri et al. 1999; Rolls et al. 2003b). This case is important because it underlines the point that when the 'signal' correlations are low, the information carried by the cells is almost independent, even when there is a high 'noise' correlation.

Power efficiency and accuracy of the information measurement algorithm

To analyze the accuracy and power efficiency (in terms of the number of trials of data needed) of the measurement of information using decoding procedures as described in this paper, we performed simulations for cases in which the information measures provided by the approach could be compared with the exact value of the information computed directly from Eq. 1 and knowledge of the exact probability distributions of the neuronal responses to each stimulus. Figure 8(top) shows an example of information from four simulated Poisson cells responding to four different stimuli with firing rates and firing rate distributions to the different stimuli similar to those of cells in inferior temporal visual cortex (Treves et al. 1999). In this example, and other simulated cases, the decoding methods give reasonably accurate estimates of the true information, with a loss of information of approximately 10% for maximum likelihood (ML) and 15–20% for full probability estimation (PE) decoding combined with a Gaussian or Poisson fit of the responses. We note that decoding approaches inherently can only approach the true information (Pouget et al. 1998; Cover and Thomas 1991; Rolls et al. 1997), and that therefore the values obtained are very satisfactory. It has also been suggested that decoding procedures are likely to become more accurate when increasing numbers of cells and longer times are considered (Panzeri et al. 1999), although in those cases the direct measurement of the information becomes computationally very long. For the case of maximum

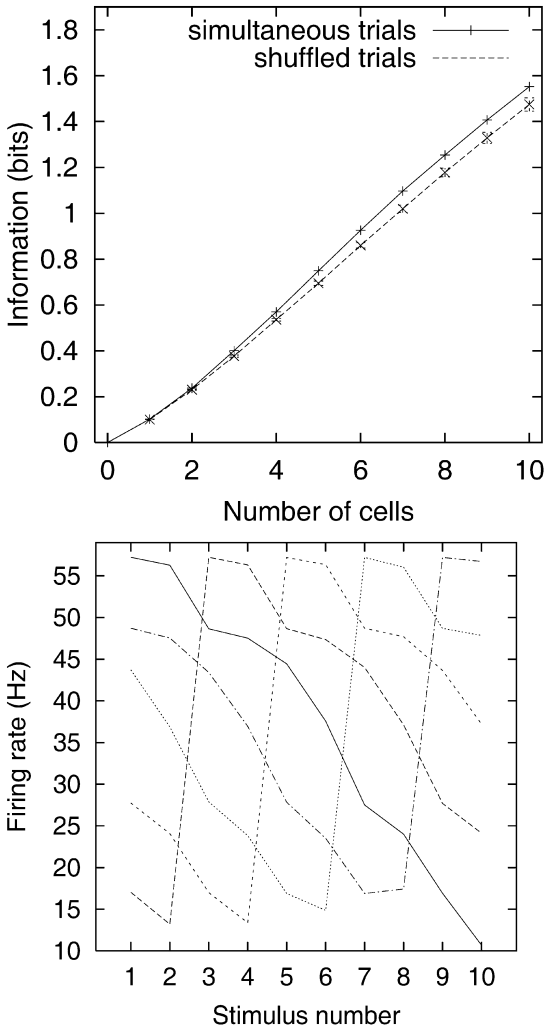

Fig. 7 *Above* The effect of anti-correlated response profiles of neurons, with correlated variability on a trial-by-trial basis (i.e. with positive 'signal' correlation and positive 'noise' correlation). The information available from only the spike counts is shown. Because the anti-correlation in the response profiles is small (Pearson correlation r=−0.09, ν=−0.018), the information increases approximately linearly with the number of neurons. Shuffling the trials (within each stimulus) decreases the information, showing that the simultaneously recorded unshuffled spike count data reflect some synergy, due to the positive 'noise' correlations (when they occur with negative 'signal' correlations). *Below*: the response profiles of neurons 1, 3, 5, etc., to the ten discrete stimuli. The correlated noise was introduced by providing the ten neurons with 50% of common input

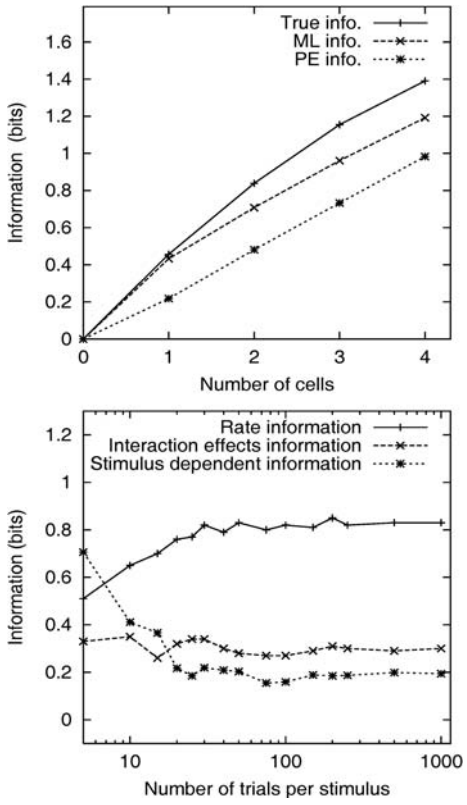

Fig. 8 *Above* The amount of information extracted by the decoding procedures, when using a Gaussian fit of the responses combined with maximum likelihood (ML) or with full probability estimation (PE) methods, compared to the exact information value. In the simulation there were four Poisson simulated cells firing to four different stimuli. *Below* The efficiency of the decoding procedure as a function of the number of trial for each stimulus shown for the different cases analyzed in the paper. For the analyses shown, there were ten cells and ten stimuli

likelihood decoding, conditions for an optimal decoding were analytically derived by Samengo (2002). We note that the efficiency of the approach for measuring the information in the stimulus-dependent cross-correlations will be similar to that shown in Fig. 8, as the values of the cross-correlations have an approximately Poisson distribution.

Figure 8(bottom) shows how the estimates by the algorithm of the information arising in different ways from spike trains depend on the number of trials of data for each stimulus. In the case shown, there were ten stimuli, and ten cells. For all cases, the information estimate has settled down when the number of trials is approximately 20 trials per stimulus, that is when the number of trials is twice the number of stimuli. In the case of the information from the firing rates, the information is underestimated when the number of trials is less than twice the number of stimuli, and this may be due to the probability tables generated from the responses available to each stimulus being inaccurate due to undersampling.

The stimulus-dependent cross-correlation information is overestimated a little when the number of trials is less than twice the number of stimuli, and this may be because the correlation measures have different variability than assumed by the bias correction procedures. The redundancy arising from interactions between the signal and noise correlations illustrated in Figs. 6 and 7, which are detected by the changes produced by trial shuffling within a stimulus, were little affected by altering the number of trials. This is because these interaction effects are measured just by the difference that the shuffling produces in an estimate made of the Rate information.

Application to the representation of information in the inferior temporal visual cortex

We tested the method with real neuronal data as follows, to show how the algorithm can be applied to real neuronal data. In an example of one of the experiments (bj293) simultaneous recordings were made from a group of four neurons in the inferior temporal cortex to a set of eight stimuli of the type known to elicit responses in some neurons in this region. The neurons were recorded on three microelectrodes within the same cortical area but separated by up to 2 mm. There were 16 trials for each stimulus, and the data from a 200-ms epoch starting 100 ms after the presentation of each stimulus were analyzed. The results of the analysis are shown in Fig. 10. The Rate information (i.e. based only on the firing rate counts) is shown above. The information rises almost linearly with the number of cells, and shuffling the trials increases the information by the small amount of 0.03 bits, indicating a small amount of redundancy arising from the stimulus-independent correlation term of the information, which reflects interactions between the signal correlations ν and the noise correlations γ as will be explained in the "Discussion" using Fig. 9.

Figure 10 (lower) shows the analysis from the same experiment based only on the correlations between pairs of neurons, to measure any information in stimulus-dependent synchronization. The cross-correlation values used were for a lag of zero, because there was some evidence for the presence of cross-correlations at this lag from the cross-correlograms. The cross-correlations from all possible cell pairs (6) were used. The information is shown as a function of the number of cross-correlations included in the analysis. The solid line shows the data without trial shuffling. The facts that the information after shuffling was lower, and was less than two standard deviations from the unshuffled values, indicate that the stimulus-dependent cross-correlation (synchronization) information is not significantly different from what could arise by chance pairings of trials. Further, the magnitude of the information available from the cross-correlations was small, approximately 0.078 bits minus the 0.052 bits which the Monte Carlo trial shuffling shows could arise by chance, that is 0.026 bits. This is small in relation to the information available from the firing rates, shown above, of 0.341 bits.

The measures discussed so far from the set of four real cells in experiment bj293 used for Fig. 10 are of the information available from the Rates, and separately that available from the cross-correlations between pairs of cells. As these two terms could be redundant, we also

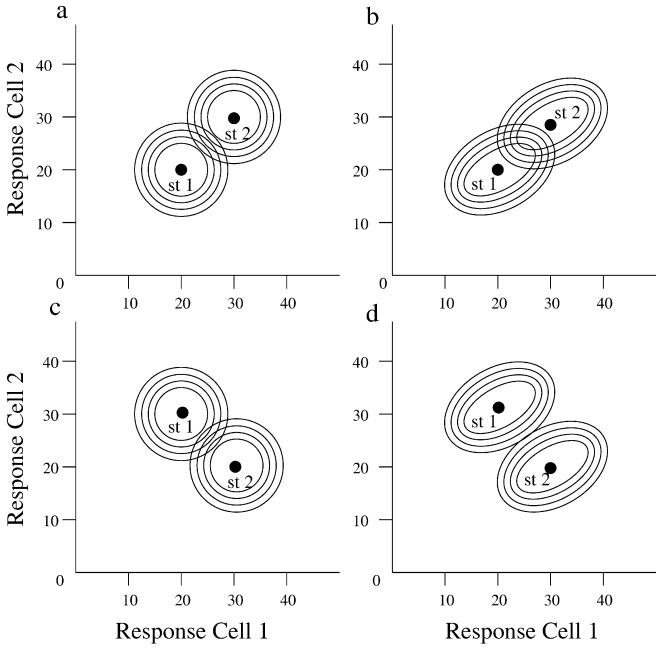

Fig. 9 The effects on the information available from the spike counts of correlations between the response profiles of the cells ('signal' correlations) and of covariations in the firing rate spike count variability ('noise' correlations). The mean response of each cell to each stimulus is shown by the filled circle, and the contour lines show the probabilities of obtaining particular firing rates on individual trials. (This figure reflects in part earlier work of Oram et al. (1998.)

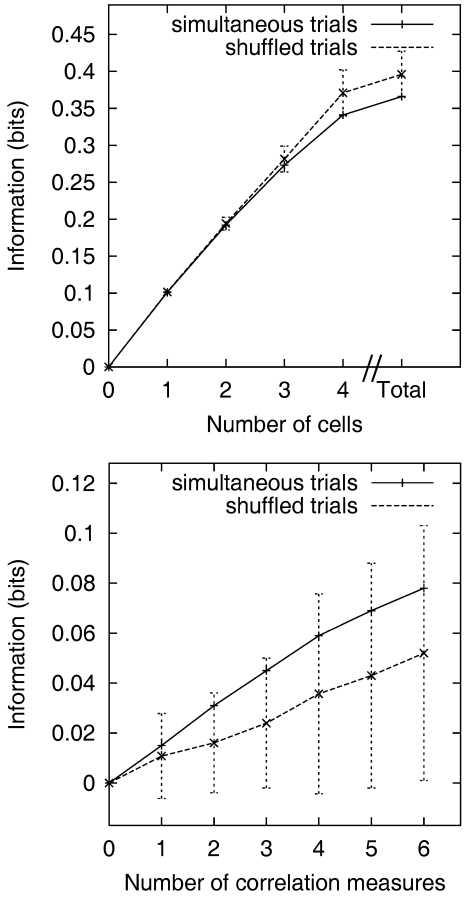

Fig. 10 Information analysis on a set of four simultaneously recorded neurons. The rate information is calculated from the spike number of the cells (top) while the information in the cross-correlation for zero lag is plotted in the graph below where the information is that extracted from the correlations between the six pairs that can be formed with the four neurons analyzed. The Total Information shown in the upper graph is that measured when both the rate and the cross-correlation data are used together in the algorithm (see text)

show in Fig. 10 (upper) the Total Information available from both the firing rates and the cross-correlations when both are present in the table data(*s*, **r**) from which the information is measured. The values for this Total Information are shown on the right of the upper part of the figure. In this case the Total Information had a value that is approximately the sum of the Rate Information plus the Stimulus-Dependent Information from the cross-correlations, indicating that in this case these two terms provide almost independent contributions to the total information available.

We applied the algorithm described in this paper to a population of 46 neurons recorded in simultaneous sets of 2–4 from the inferior temporal visual cortex. The stimulus set analyzed consisted of eight visual stimuli of objects and faces of the type known to be effective for activating neurons in the inferior temporal visual cortex (Rolls 2000). Each neuron was shown to respond to some but not others of the set of stimuli. The response latencies were typically 90–100 ms, and the period analyzed was 100–300 ms poststimulus onset. Because each stimulus is composed of a number of parts arranged in the correct spatial configuration, synchronization between the spikes of different neurons might be used to bind together the parts (Singer 2000), and it is therefore of importance to measure whether any extra information about which stimulus was being shown was available from stimulus-dependent synchronization. In addition, a previous analysis we performed (on a different population of neurons) of how information increases with the number of cells in a population did not measure any redundancy arising from stimulus independent correlations between the numbers of spikes obtained from the different cells on a trial-by-trial basis because the cells were not simultaneously recorded (Rolls et al. 1997), and it was possible to measure the degree of redundancy arising in this way from each set of two to four simultaneously recorded cells in the new dataset.

The results are shown in Table 1. In each experiment (e.g. bj185) simultaneous recordings were made from the number of cells indicated, with each cell providing when measured separately at least 0.025 bits of information about the stimulus set. The average number of trials for each stimulus was 16, so that the criterion for operation of the algorithm, two times the total number of stimuli, was met. Across the set of experiments, the mean information from the firing rates (i.e. spike counts) was 0.142 ±0.088 bits. As shown in the last column of Table 1, the mean value of the stimulus-independent correlation information included in this Rate term was −0.006 ±0.014. The negative value for this term on average indicates that it is redundancy which arises from interactions between the signal and noise correlations as will be shown in Fig. 9b, though as shown in Table 1 for some cases this term can be positive (see Fig. 9d). The finding that this value is small, approximately 4.1% of the total information on average, indicates that the cells encode information almost independently (Rolls et al. 1997, 2003a).

Table 1 The contributions (in bits) of the different components to the information extracted by a decoding algorithm in 200 ms from 20 sets of simultaneously recorded inferior temporal cortex neurons when shown eight stimuli effective for the cells

Experiment	Number of cells	Rate information from the spike counts	Stimulus dependent information from cross correlation	Total information	Stimulus independent information from correlations included in the rate
bj185	2	0.084	−0.028	0.085	0.007
bj207	2	0.056	−0.001	0.058	−0.003
bj213	2	0.108	0.000	0.108	−0.002
bj215	3	0.085	0.015	0.094	0.003
bj220	3	0.142	0.015	0.152	−0.026
bj229	2	0.066	0.003	0.073	0.001
bj278	3	0.185	0.000	0.188	−0.019
bj280	3	0.312	−0.006	0.312	−0.005
bj283	2	0.183	0.000	0.183	−0.028
bj285	2	0.034	0.001	0.034	−0.001
bj287	3	0.120	−0.017	0.121	0.010
bj288	3	0.153	0.001	0.163	0.006
bj290	4	0.037	0.013	0.047	−0.016
bj291	2	0.121	0.000	0.121	0.009
bj292	2	0.182	0.000	0.182	0.009
bj292b	4	0.213	0.005	0.216	−0.018
bj293	4	0.341	0.026	0.366	−0.030
Mean SD	2.70	0.142±0.088	0.002±0.012	0.147±0.090	−0.006±0.014
Percentage	-	96.8%	1.1%	100%	−4.1%

The column of Table 1 showing the stimulus-dependent information from the cross-correlations indicates that the contribution of this term was on average small (1.1%), and indeed for no experiment was statistically significant. (The values of the stimulus-dependent cross-correlation information shown are corrected by shuffling for any random pairings of spike trains in the set of trials that might provide by chance measurable information, and because of this correction the values for any individual experiment can appear as small negative numbers.) The total information is that measured when all contributions from both the firing rate and cross-correlations are included. (This value is approximately, but not exactly, the sum of the Rate and Stimulus-dependent cross-correlation terms shown in Table 1. The Total Information is not exactly the sum mainly because these two components need not be independent.)

Discussion

The algorithm described here allows the information that can arise in different ways from the spike trains of simultaneously recorded neurons to be extracted and identified as follows. One type of information is the stimulus-dependent cross-correlation information, which can arise from stimulus-dependent correlations between the spike trains. This is detected in the algorithm by using only the correlation columns in the table data(*s,* **r**) shown in Fig. 2. Shuffling the trials (within stimuli) allows a (typically small) correction to be applied, and for the standard deviation of the estimate to be measured, as shown in Fig. 4. This is the only stimulus-dependent cross-correlation information that is available from the spike trains. All the other cross-correlation information is stimulus-independent.

The Rate information, that is the amount of information in the firing rates and excluding any stimulus-dependent synchronization effects, arises and is identified, in several ways. We note that the Rate information is estimated using just the rate columns of table data(*s,* **r**) shown in Fig. 2. We also note further that to obtain an estimate of the Total Information in the spike trains, the algorithm must be run on the whole table data(*s,* **r**) shown in Fig. 2, and may not be the sum of the Rate and Stimulus-Dependent Synchronization Information, as they could encode the same information, and be redundant. The first factor affecting the Rate information is, as shown in Fig. 5, any redundancy that is related to the correlation between the response profiles of the different neurons (i.e. positive or negative 'signal' correlations) and which is not related to an interaction with the noise correlations. This effect is detected by a sublinear increase in the information with the number of neurons, and a lack of effect of shuffling the trials. The second factor contributing to the Rate information is, as shown in Figs. 6 and 7, an interaction between the 'signal' correlations and the 'noise' correlations. The magnitude of this interaction is detected by the size of the effect produced by shuffling spike count measures between the trials within each stimulus in table data(*s,* **r**) shown in Fig. 2.

The ways in which interactions between the signal and noise correlations contribute redundancy or synergy to the Rate Information are clarified with reference to Fig. 9. The issue we deal with is not stimulus-dependent synchrony effects, but instead how the variability of the spike counts interacts with the response profiles of the cells. The left part of Fig. 9 shows the case where the degree of variability of the spike counts of two cells is uncorrelated. The trial by trial variability, measured by the noise correlation, is zero, and is indicated by the circular contour lines surrounding the mean responses of each cell to each stimulus. In the upper case, the response profiles of the two cells (measured by the 'signal' correlations) are correlated, in that when cell 1 has a large response, so does cell 2. The lower left case shows anticorrelated response profiles, in that when cell 1 fires fast (in response to stimulus 2), cell 2 fires slowly. In both cases on the left of the diagram there is some overlap in the response profiles, that is some uncertainty about which stimulus is present on some trials. We show on the right the effect of correlated noise (produced for example by common input to the two cells). When the response profiles are correlated (upper right), the positive 'noise' correlation increases the uncertainty, that is, decreases the information about which stimulus was shown. The interaction between the signal and noise correlations in this case increases the redundancy of the encoding by the two cells. In contrast, when the response profiles are anticorrelated (lower right), the positive 'noise' correlation decreases the uncertainty, that is, increases the information about which stimulus was shown. In this second case, because the signal and noise correlations have opposite signs, there is synergy (the opposite of redundancy) produced by the interaction. This simple case demonstrates that interactions between the signal and noise correlations can produce synergy (which is produced if the signs of the signal and noise correlations are opposite) or redundancy (if the signs of the signal and noise correlations are the same). We have demonstrated above how the decoding approach to information measurement we describe in this paper is able to diagnose these interaction effects, which are revealed by the procedure of shuffling the trials within a stimulus in the part of the table data(*s,* **r**) shown in Fig. 2 that contains the spike count (i.e. firing rate) information.

We now compare the current approach with earlier approaches to measuring the total information available in the spike trains of simultaneously recorded neurons, including the information available from effects such as stimulus-dependent synchronization. First, the original Taylor expansion (Rolls et al. 2003a, 2003b; Panzeri et al. 1999) used an estimate of the stimulus-dependent synchrony effects that was effectively an estimate of the trial by trial co-variability of the spike trains of different neurons, rather than a direct measure of the synchronization. In the algorithm described here, we decided in practice to utilize a direct and standard measure of the synchronization between pairs of neurons for the correla-

tion columns of table data(*s,* **r**) shown in Fig. 2, because we believe this to be most relevant, and synchronization is a key current issue in understanding the neural encoding of information. We note that it is a minor alteration to the current decoding algorithm to replace the correlation columns of table data(*s,* **r**) shown in Fig. 2 with other measures of the co-variability of neuronal spike trains, including a measure analogous to the scaled cross-correlation density (Panzeri et al. 1999) computed for each trial. We note that a spike time version of the Taylor expansion algorithm has been described by Panzeri and Schultz (2001), though this algorithm suffers greatly from the limitations described next. Third, the Taylor expansion approach (Bezzi et al. 2002; Rolls et al. 2003a, 2003b; Panzeri et al. 1999) is severely limited in the numbers of spikes with which it can deal, as it can deal with only up to two spikes on every trial from any neuron (which limits greatly the length of the time epochs of neuronal firing to which it can be applied), and is also limited in the number of cells to which it can be applied. In contrast, the present approach works better with large numbers of cells (and indeed may be a little limited in its application to a small number of cells, although in practice we have shown that the number can be as few as two or three cells, see Fig. 8a), and is unlimited in the total time epoch that is considered. The current approach is relatively efficient in the number of trials of data needed, in that as shown in Fig. 8b the algorithm operates well with as few as twice the number of trials as there are stimuli. The algorithm we describe here should have many applications, as it can in principle be applied to datasets with very large numbers of simultaneously recorded neurons, and with temporal windows that can be as long as are of interest. Another approach (Hatsopoulos et al. 1998) used a direct method to compute the information, and a disadvantage of this approach is that to avoid the limited sampling bias problem, many trials of data are needed. Another approach (Oram et al. 2001) is to use a neural network to estimate the probability distributions, but this regularizes the data in a particularly complex and data-dependent way that is difficult to interpret and can sometimes leads to paradoxical results (Panzeri and Treves 1996). The decoding approach introduced in this paper has the advantages that it can operate with relatively limited numbers of trials, and has no limitations on the number of cells or spikes that can be analyzed.

While one of the main contributions of this paper has been to introduce measures of spike synchrony between neurons into the decoding algorithm, it is useful to discuss the efficiency of different decoding algorithms, in the context of how they are applied to the spike count distributions which are part of what has to be decoded. The information measurement procedure we describe can of course be used with any decoding algorithm. (We note that no decoding procedure is likely to be perfect, and that different decoding methods may be appropriate for different types of data; Robertson et al. 1999.) First, the Bayesian decoding used with for example Gaussian and Poisson fits to the spike count probability distributions has the potential weakness that it assumes that the spike counts of the different neurons are independent. However, Wu et al. (2001) show that for example when maximum likelihood estimation is used with Bayesian decoding, then correlations do not strongly affect the accuracy of the information measurement, particularly as the number of elements in the decoded vector becomes large ≈100. Second, we note that dot product decoding does not have this difficulty associated with Bayesian decoding, and moreover is very biologically plausible, in that the simplest model of a neuron holds that it performs a dot product between its input spike activity vector and its synaptic weight vector. Dot product decoding is not as efficient as maximum likelihood Bayesian decoding (Rolls et al. 1997), but the latter is concerned with measuring all the information that potentially may be present, without any constraints on the decoding process. The dot product decoding does at least allow the information obtained as a function of the number of neurons in the population, and including or not the synchrony between neurons, to be estimated. Third, we introduce in this paper the synchronicity measure of neuronal firing to estimate the information available from stimulus-dependent correlations between the neurons. However, we also introduce a method for measuring the stimulus-independent correlations that are present in the spike counts of the simultaneously recorded neurons. This method involves comparing the information before and after shuffling of the trials within each stimulus. The effect of the shuffling is to remove trial-by-trial covariability in the spike counts of the different neurons (measured by γ_{ij}), and thus is able to remove the effects of stimulus-independent information, which only is non-zero if γ_{ij} is non-zero (Rolls et al. 2003b; Panzeri et al. 1999). Given that both the unshuffled and the shuffled information estimates have been through the decoding process, both are likely to be underestimates of the true information. However, the effect of the shuffling procedure is to measure how much, relatively, the decoded information measures are influenced by stimulus-independent correlations. Provided that it is remembered that this estimate is relative to the decoded values, the procedure is very useful, as established by the analyses of the quantitatively simulated datasets shown in Figs. 5, 6, and 7.

When applied to real data from 17 experiments each with 2–4 simultaneously recorded inferior temporal cortex neurons (see Table 1), the main conclusions were that almost all of the information was available in the spike counts of the neurons, that this Rate information included on average very little redundancy arising from stimulus-independent correlation effects, and that stimulus-dependent cross-correlation effects (i.e. stimulus-dependent synchronization) contribute very little to the encoding of information in the inferior temporal visual cortex about which object or face has been presented. Although conceivably over very large numbers of neurons, or under different stimulus presentation conditions, or in different brain areas, stimulus-dependent synchronization of spikes from different neurons might statistically provide

some evidence about which stimulus was presented, we note that in no individual experiment in Table 1 was the stimulus-dependent information statistically significant; in all cases it was small in relation to the information available from the spike counts of the neurons, and there was little evidence that the stimulus-dependent synchronization information was orthogonal to the information available from the spike counts.

Acknowledgements This research was supported by the Medical Research Council, grant PG9826105, by the Human Frontier Science Program, by the MRC Interdisciplinary Research Centre for Cognitive Neuroscience, and by the Wellcome Trust.

References

Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. Neural Comput 11:91–101

Aertsen AMHJ, Gerstein GL, Habib MK, Palm G (1989) Dynamics of neuronal firing correlation: modulation of 'effective connectivity'. J Neurophysiol 61:900–917

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 Biol Sci 264:1775–1783

Bezzi M, Diamond ME, Treves A (2002) Redundancy and synergy arising from pairwise correlations in neuronal ensembles. J Comput Neurosci 12:165–174

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

Brody CD (1999) Correlations without synchrony. Neural Comput 11:1537–1551

Cover TM, Thomas JA (1991) Elements of information theory. Wiley, New York

Dan L, Alonso JM, Usrey WM, Reid RC (1998) Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. Nature Neurosci 6:501–507

Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? J Neurosci 13:2758–2771

Gill PE, Murray W, Wright MH (eds) (1981) Practical optimization. Academic Press, London

Hatsopoulos NG, Ojakangas CL, Paninski L, Donoghue JP (1998) Information about movement direction obtained by synchronous activity of motor cortical neurons. Proc Natl Acad Sci U S A 95:15706–15711

Konig P (1994) A method for the quantification of synchrony and oscillatory properties of neuronal activity. J Neurosci Methods 54:31–37

Lee D, Port NL, Kruse W, Georgopoulos AP (1998) Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. J Neurosci 18:1161–1170

Nirenberg S, Carcieri SM, Jacobs AL, Latham PE (2001) Retinal ganglion cells act largely as independent encoders. Nature 411:698–701

Oram MW, Foldiak P, Perrett DI, Sengpiel F (1998) The 'ideal homunculus': decoding neural population signals. Trends Neurosci 21:259–265

Oram MW, Hatsopoulos NG, Richmond BJ, Donoghue JP (2001) Excess synchrony in motor cortical neurons provides direction information that is redundant with the information from coarse temporal response measures. J Neurophysiol 86:1700–1716

Panzeri S, Schultz SR (2001) A unified approach to the study of temporal, correlational and rate coding. Neural Comput 13:1311–1349

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

Panzeri S, Treves A, Schultz S, Rolls ET (1998) Decoding population responses in short time epochs. In: Niklasson L, Boden M, Ziemke T (eds) Proceedings of the 8th International Conference on Artificial Neural Networks. Springer-Verlag, London, pp 965–972

Panzeri S, Schultz SR, Treves A, Rolls ET (1999) Correlations and the encoding of information in the nervous system. Proc R Soc Lond B Biol Sci 266:1001–1012

Panzeri S, Treves A, Schultz S, Rolls ET (1999) On decoding the responses of a population of neurons from short time epochs. Neural Comput 11:1553–1577

Panzeri S, Petersen RS, Schultz SR, Lebedev M, Diamond ME (2001) The role of spike timing in the coding of stimulus location in rat somatosensory cortex. Neuron 29:769–777

Pouget A, Zhang K, Deneve S, Latham P (1998) Statistically efficient estimation using population coding. Neural Comput 10:373–401

Reich DS, Mechler F, Victor JD (2001) Independent and redundant information in nearby cortical neurons. Science 294:2566–2568

Robertson RG, Rolls ET, Georges-François P, Panzeri S (1999) Head direction cells in the primate pre-subiculum. Hippocampus 9:206–219

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, Deco G (2002) Computational neuroscience of vision. 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, Treves A (1998) Neural networks and brain function. Oxford University Press, Oxford

Rolls ET, Treves A, Tovee MJ (1997) The representational capacity of the distributed encoding of information provided by populations of neurons in the primate temporal visual cortex. Exp Brain Res 114:149–162

Rolls ET, Aggelopoulos NC, Franco L, Treves A (2003a) Information encoding in the inferior temporal visual cortex: contributions of the firing rates and the correlations between the firing of neurons. Biol Cybern (in press)

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

Salinas E, Abbott LF (1994) Vector reconstruction from firing rates. J Comput Neurosci 1:89–107

Samengo I (2002) Information loss in an optimal maximum likelihood decoding. Neural Comput 14:771–778

Shadlen M, Newsome W (1994) Is there a signal in the noise? Curr Opin Neurobiol 5:248–250

Shadlen M, Newsome W (1998) The variable discharge of cortical neurons: implications for connectivity, computation and coding. J Neurosci 18:3870–3896

Shannon CE (1948) A mathematical theory of communication. AT&T Bell Labs Tech J 27:379–423

Singer W (2000) Response synchronisation: a universal coding strategy for the definition of relations. In: Gazzaniga M (ed) The new cognitive neurosciences, 2nd edn, chap 23. MIT Press, Cambridge, MA, pp 325–338

Sompolinsky H, Yoon H, Kang K, Shamir M (2001) Population coding in neuronal systems with correlated noise. Phys Rev E64

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

Wu S, Nakahara H, Amari S (2001) Population coding with correlation and an unfaithful model. Neural Comput 13:775–797