
Correlations and the encoding of information in the nervous system

Stefano Panzeri^{1*}, Simon R. Schultz¹, Alessandro Treves² and Edmund T. Rolls¹

¹University of Oxford, Department of Experimental Psychology, South Parks Road, Oxford OX1 3UD, UK

²International School for Advanced Studies—Programme in Neuroscience, 34013 Trieste, Italy

Is the information transmitted by an ensemble of neurons determined solely by the number of spikes fired by each cell, or do correlations in the emission of action potentials also play a significant role? We derive a simple formula which enables this question to be answered rigorously for short time-scales. The formula quantifies the corrections to the instantaneous information rate which result from correlations in spike emission between pairs of neurons. The mutual information that the ensemble of neurons conveys about external stimuli can thus be broken down into firing rate and correlation components. This analysis provides fundamental constraints upon the nature of information coding, showing that over short time-scales correlations cannot dominate information representation, that stimulus-independent correlations may lead to synergy (where the neurons together convey more information than they would if they were considered independently), but that only certain combinations of the different sources of correlation result in significant synergy rather than in redundancy or in negligible effects. This analysis leads to a new quantification procedure which is directly applicable to simultaneous multiple neuron recordings.

Keywords: correlations; synchronization; information theory; neural code; synergy; redundancy

1. INTRODUCTION

A controversy exists over the extent to which firing rates and correlations between the responses of different cells (such as synchronization of action potentials) contribute to information representation and processing by neural ensembles. Encoding of information in the correlation between the firing of different neurons has been demonstrated in specialized sensory systems such as the retina of the salamander (Meister *et al.* 1995) and the auditory localization system of the barn owl (Carr & Konishi 1990; Carr 1993). For the mammalian cortex, the evidence is less clear. Several investigators (Vaadia *et al.* 1995; deCharms & Merzenich 1996; Murthy & Fetz 1996; Riehle *et al.* 1997; Singer *et al.* 1997; Donoghue *et al.* 1998) have presented evidence of stimulus-related changes in the correlation of firing between small populations of cortical cells. This might imply that correlations among cortical neurons reveal a substantial capability of cortical neural ensembles to code information synergistically—that is, that the ensemble of neurons provides more information than the simple sum of the contributions of the individual cells. Another possibility is that correlations actually limit, rather than improve, the information capacity of the population (Gawne & Richmond 1993; Zohary *et al.* 1994; Shadlen & Newsome 1998). A third possibility is that correlations have no major effect on the efficiency of neural codes (Golledge *et al.* 1996; Amit 1997; Rolls & Treves 1998). What is clearly needed, in order to reconcile these viewpoints, is a rigorous quantitative

methodology for addressing the roles of both correlations and rates in information encoding.

Information theory (Shannon 1948; Cover & Thomas 1991), which has found much use in recent years in the analysis of recordings from single cells (Optican & Richmond 1987; Rieke *et al.* 1996; Rolls & Treves 1998; Kitazawa *et al.* 1998), provides the basis for such an approach. Ideally, to divide the information into components indicative of the information encoding mechanisms involved, it would be desirable to take any experiment in which population activity was recorded in response to a clearly identifiable stimulus that the cells participate in encoding, and determine how many bits of information were present in the firing rates, how many in coincident firing by pairs of neurons, etc. By considering the limit of rapid information representation, in which the duration of the window of measurement is of the order of the mean inter-spike interval, it is possible to do just this. This short time-scale limit is not only a convenient approach to a complex problem, but it is also likely to be of direct relevance to information processing by the brain, as there is substantial evidence that sensory information is transmitted by neuronal activity in very short periods of time. Single unit recording studies have demonstrated that the majority of information is often transmitted in windows as short as 20–50 ms (Werner & Mountcastle 1965; Oram & Perrett 1992; Rolls & Tovéé 1994; Tovéé *et al.* 1993; Heller *et al.* 1995; Macknik & Livingstone 1998). Event-related potential studies of the human visual system (Thorpe *et al.* 1996) provide further evidence that the processing of information in a multiple-stage neural system can be extremely rapid. Finally, if one wishes to assess the information content of correlational assemblies

*Author for correspondence (stefano.panzeri@psy.ox.ac.uk).

which may last for only a few tens of milliseconds (Singer *et al.* 1997), then appropriately short measurement windows must be used.

The current work describes the use of information theory to quantify the relative contributions of both firing rates and correlations between cells to the total information conveyed.

We report an expansion of the expression for mutual information to second order in time, show that the second-order terms break down into those dependent on rate and those dependent on correlation, and demonstrate that this can form the basis of a procedure for quantifying the information conveyed by simultaneously recorded neuronal responses. We show that with pairs of cells the approach works well for time windows several hundred milliseconds long, the time range of validity decreasing approximately inversely with population size. The expansion for the information shows that in short but physiological time-scales, firing rates dominate information encoding when cell assemblies of limited size are considered, because correlations begin to contribute only with subleading terms. We find that the redundancy of coding is dependent on the specific combination of the correlation in the number of spikes fired in the time window, which in the limit of short time windows measures the probability of coincident firing and thus the extent of synchronization, and the correlation among mean response profiles to different stimuli. Furthermore, we observe that even stimulus-independent correlations may in some circumstances lead to synergistic coding.

2. METHODS

(a) *Information carried by neuronal responses and its short time-scale expansion*

Consider a period of time beginning at t_0 , of (short) duration t , in which a stimulus s is present. Let the neuronal population response during this time be described by the vector \mathbf{n} , the number of spikes fired by each cell. Typically each component will have a value of zero or one, and only rarely higher. Alternatively, we may describe the response by the firing rate vector $\mathbf{r} = \mathbf{n}/t$. In a typical cortical neurophysiology experiment we might hope to have tens of such trials with the same stimulus identification s . The stimuli considered are purely abstract: the procedures detailed in this paper are applicable to a wide variety of experimental paradigms. However, it may help to conceptualize the stimulus as, for example, that object (of some set) which is being viewed by the experimental subject; indeed, data from such an experiment are examined later in this paper. Consider the stimuli to be taken from a discrete set \mathcal{S} with S elements, each occurring with probability $p(s)$. The probability of events with response \mathbf{r} is denoted as $p(\mathbf{r})$, and the joint probability distribution as $p(s, \mathbf{r})$.

Following Shannon (1948), we can write down the mutual information provided by the responses about the whole set of stimuli as

$$I(t) = \sum_{s \in \mathcal{S}} \sum_{\mathbf{r}} p(s, \mathbf{r}) \log_2 \frac{p(s, \mathbf{r})}{p(s)p(\mathbf{r})}. \quad (1)$$

This assumes that the true probabilities $p(s, \mathbf{r})$ are available. In practice, however, we have only a small to moderate number of events from which to compute the frequency table, and as a

result a bias is introduced which must be corrected (Panzeri & Treves 1996).

Now, the information can be approximated by a power series

$$I(t) = t I_t + \frac{t^2}{2} I_u + \dots, \quad (2)$$

where I_t refers to the instantaneous information rate and I_u to the second time derivative of the information, the instantaneous information 'acceleration'. (The zeroth order, time-independent term is not included in equation (2) because it is, of course, zero, as no information can be transmitted by the neurons in a time window of null length.) For short time-scales, only the first- and second-order terms survive: higher-order terms in the series become negligible. The time derivatives of the information can be calculated by taking advantage of the short time limit, as we shall explain.

(b) *Definitions and measurement of correlations for short time-scales*

There are two kinds of correlations that influence the information. These have been previously termed 'signal' and 'noise' correlations (Gawne & Richmond 1993). They can be distinguished by separating the responses into 'signal' (the average response to each stimulus) and 'noise' (the variability of responses from the average to each stimulus). The correlations in the response variability represent the tendency of the cells to fire more (or less) than average when a particular event (e.g. a spike from another neuron) is observed in the same time window. For short time windows, this of course measures the extent of synchronization of the cells. Note that we do not assume that the trial-by-trial variability is just due to noise, and our analysis is independent of the cause of this variability, which could arise from a number of factors (Arieli *et al.* 1996). We name the correlations in the variability as 'noise correlations' only for consistency with previous literature (Gawne & Richmond 1993; Gawne *et al.* 1996). In fact, our results reported below (equation (9)), precisely quantify the contribution of correlations in the trial-by-trial variability to the information carried by the neuronal population, and show in which cases correlations in this variability contribute positively to information transmission. One way to introduce the parameters $\gamma_{ij}(s)$ quantifying noise correlation in the short time limit is in terms of the conditional firing probabilities

$$p(n_i(s) = 1 | n_j(s) = 1; s) \equiv \bar{r}_i(s) t(1 + \gamma_{ij}(s)) + O(t^2), \quad (3)$$

where $\bar{r}_i(s)$ is the mean response rate of cell i (among C cells in total) to stimulus s over all the trials in which that stimulus was present and $O(t^2)$ means that the terms neglected in equation (3) are second order in t . It has been assumed in the above that the conditional probabilities (3) scale proportionally to t ; this is the only assumption underlying our time expansion. It is a natural assumption, being violated only in the implausible case of spikes locked to one another with infinite time precision, but in any case it can be verified for any given data set.

For $i \neq j$, equation (3) gives us the probability of coincident firing: cells i and j both fire spikes in the same time period; $\gamma_{ij}(s)$ is the fraction of coincidences above (or below) that expected from uncorrelated responses, normalized to the number of coincidences in the uncorrelated case (which is $\bar{n}_i(s)\bar{n}_j(s)$, the bar denoting the average across trials belonging to stimulus s). $\gamma_{ij}(s)$ ($i \neq j$) is thus given by the following expression:

$$\gamma_{ij}(s) = \frac{\overline{n_i(s)n_j(s)}}{(\bar{n}_i(s)\bar{n}_j(s))} - 1, \quad (4)$$

and is named the ‘scaled cross-correlation density’ (Aertsen *et al.* 1989). It can vary from -1 to ∞ ; negative $\gamma_{ij}(s)$ values indicate anticorrelation, whereas positive $\gamma_{ij}(s)$ values indicate correlation. This measure of correlation is used because it has a number of advantages for short time windows over the better-known Pearson correlation coefficient for this application as described in Appendix A. (In brief, the ‘scaled cross-correlation measure’ γ_{ij} , unlike the Pearson measure, remains finite for short time windows, and in addition the mathematical derivation of our result is simpler and much more compact with this measure. We note that the particular measure used for correlations is ultimately a matter of mere notation, and it has, of course, no effect at all on the value of each of the information components presented below, as discussed in Appendix A.)

For $i=j$, equation (3) instead gives the probability of observing a spike emission by cell i , given that we have observed a different spike from the same cell i during the same time window. The ‘scaled autocorrelation coefficient’ must be measured as

$$\gamma_{ii}(s) = \frac{\overline{(n_i(s))^2} - \overline{n_i(s)}^2}{\overline{n_i(s)}^2} - 1. \quad (5)$$

Notice that in the numerator of the above equation, the subtraction of $\overline{n_i(s)}^2$ is needed to quantify correctly the number of occurrences of at least two spikes from the same cells; this was not necessary when i and j represented different cells. (This may be understood by observing that assuming independence between the spikes of the same cell is equivalent to the assumption that the cell fires according to a Poisson process. For such a process the variance of the spike count is equal to the mean. The subtraction of $\overline{n_i(s)}^2$ precisely subtracts from the variance of the spike count what is expected in the independent case.) Again, this measure was used for compactness of the results. The relation of this measure to alternative ones, such as the Fano factor, is described in Appendix A. Because we are considering a single short time window, note that these correlations are observed over repeated trials with the same stimulus.

To maintain homogeneity with respect to the noise correlation case, we chose to quantify the correlations in the signal, i.e. the correlations ν_{ij} in the mean responses of the neurons across the set of stimuli, as a signal scaled cross-correlation coefficient:

$$\nu_{ij} = \frac{\langle \overline{n_i(s)} \overline{n_j(s)} \rangle_s}{\langle \overline{n_i(s)} \rangle_s \langle \overline{n_j(s)} \rangle_s} - 1 = \frac{\langle \overline{r_i(s)} \overline{r_j(s)} \rangle_s}{\langle \overline{r_i(s)} \rangle_s \langle \overline{r_j(s)} \rangle_s} - 1. \quad (6)$$

The definition is similar to that of scaled noise correlation; the main difference is that now the average is across stimuli: $\langle (\dots) \rangle_s \equiv \sum_s p(s) (\dots)$, not across trials. As before, ν_{ij} varies between -1 and ∞ .

(c) Response probability quantification

If the firing rates conditional upon the firing of other neurons are non-divergent, as assumed in equation (3), the t expansion of response probabilities becomes essentially an expansion in the total number of spikes emitted by the population in response to a stimulus. The only responses with non-zero probabilities up to order t^k are the responses with up to k spikes from the population; the only events with non-zero probability are therefore those to second order in t with no more than two spikes emitted in total:

$$p(\mathbf{0}|s) = 1 - t \sum_{i=1}^C \overline{r_i(s)} + \frac{t^2}{2} \sum_{i=1}^C \sum_{j=1}^C \overline{r_i(s)} \overline{r_j(s)} (1 + \gamma_{ij}(s)),$$

$$p(\mathbf{e}_i|s) = t \overline{r_i(s)} - t^2 \overline{r_i(s)} \sum_{j=1}^C \overline{r_j(s)} (1 + \gamma_{ij}(s)) \quad i = 1, \dots, C,$$

$$p(\mathbf{e}_{ii}|s) = \frac{t^2}{2} \overline{r_i(s)}^2 (1 + \gamma_{ii}(s)) \quad i = 1, \dots, C,$$

$$p(\mathbf{e}_{ij}|s) = t^2 \overline{r_i(s)} \overline{r_j(s)} (1 + \gamma_{ij}(s)) \quad i, j = 1, \dots, C; \quad i < j, \quad (7)$$

where $\mathbf{0}$ is the zero response (no cells fire), \mathbf{e}_i indicates a single spike fired by cell i , and \mathbf{e}_{ij} indicates a pair of spikes fired by cells i and j . The expression (7) for the probabilities can be derived by requiring that the probabilities are normalized to unity, that the conditional firing probabilities are proportional to t , as in equation (3), and by requiring, through straightforward algebraic equations, that the firing rates and the scaled auto- and cross-correlation coefficient take the values $\overline{r_i(s)}$, $\gamma_{ii}(s)$ and $\gamma_{ij}(s)$, respectively. The information derivatives are then calculated in terms of the probabilities (equations (7)) as follows. First, one inserts the probabilities (equations (7)) into the sum over responses in equation (1). Then, for each term of the sum over responses, one uses the power expansion of the logarithm as a function of t :

$$\log_2(1 - tx) = -\frac{1}{\ln 2} \sum_{j=1}^{\infty} \frac{(tx)^j}{j}.$$

Finally, grouping together all the terms in the sum which have the same power in t , and using equation (2), one obtains the expressions for the information derivatives reported below. Note that interactions among more than two cells do not affect the second-order probabilities (and thus will not affect the information up to second order).

(d) Bias calculation

The information rate I_t and each of the three components of I_u (see equation (9) below) are affected by a systematic error when calculated from limited data samples. This is a problem general to all nonlinear functions of probabilities, but it has been noted to be of particular concern for Shannon information, from which I_t and I_u derive (Panzeri & Treves 1996). Furthermore, the bias in the n th derivative of the information has a $1/t^n$ dependence, making treatment of the bias problem even more crucial. The problem can largely be avoided by estimating the bias by a standard error propagation procedure and subtracting it from the calculated quantity. The procedure is detailed in electronic Appendix B which is provided on the Royal Society Web site.

(e) Integrate and fire simulation

Correlated spike trains were simulated using a method similar to that used by Shadlen & Newsome (1998), to which we refer for a full discussion of the advantages and limitations of the model. In brief, each cell received 300 excitatory and 300 inhibitory inputs, each a Poisson process in itself, the (possibly stimulus dependent) mean rate of which is constant across the set of inputs for any specific stimulus condition, and contributed a fixed quantity to the membrane potential. (The decay time constant of the membrane potential was 20 ms for all the simulations presented, apart from that in figure 3(c,d). To test the effects of changing the precision of synchrony, in figure 3(c,d) the membrane time constant was set to 1 ms.) When

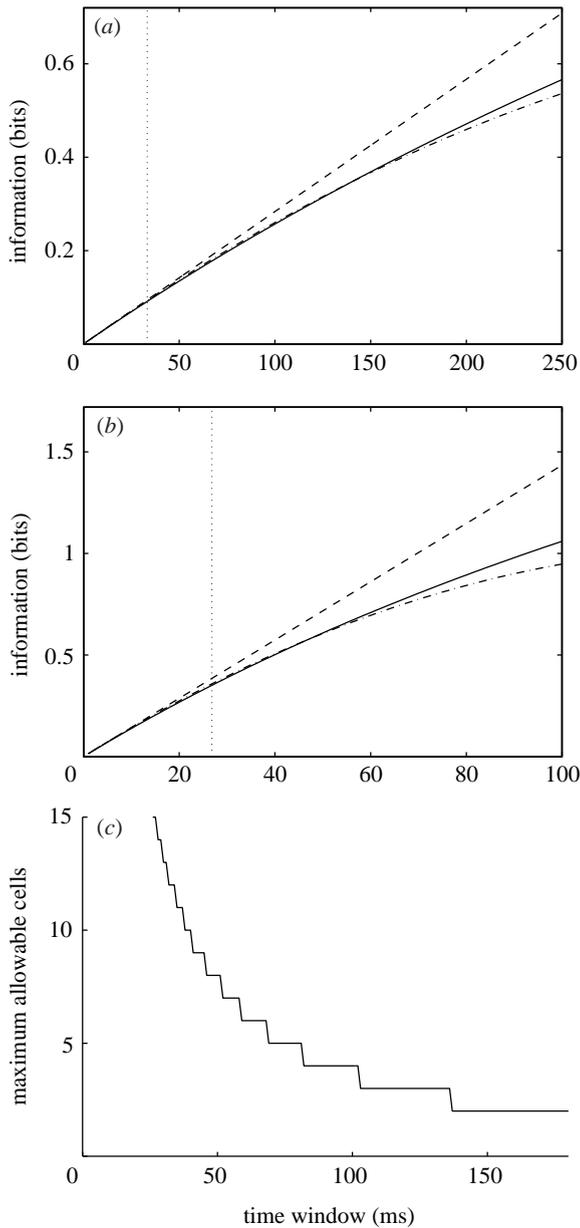


Figure 1. Range of validity of the short time-window approximation for the information. (a) The accuracy of the approximation shown for two Poisson-simulated cells with mean firing rate distributions to a set often simulated stimuli identical to those of a pair of inferior temporal cells from the data set of Booth & Rolls (1998). The true information (solid line) is compared with the first- (dashed) and second- (dot-dashed) order series approximations. The vertical dotted line indicates the mean interspike interval for the stimulus which evokes the most spikes (providing a scale of comparison for the range of validity). An extensive number of simulations with integrate and fire neurons shows that this result, pictured in the simple Poisson case, is relatively robust to neuronal firing statistics. (b) As with (a), but for five instead of two simulated Poisson cells. (c) The allowable ensemble size while limiting the error due to the approximation to 5%; extrapolation from the data set above, based on $1/C$ scaling.

the membrane potential exceeded a threshold, it was reset to a baseline value. The degree of correlation between the firing of cells was set through the proportion of inputs shared between cells (33% for figure 2(b,c), figure 3(a,b); 0% or 90% for figure 2(d-f) and figure 3(c,d)). The threshold was chosen as a

function of the membrane time constant, to guarantee the conservation of the response dynamic range, i.e. the neurons respond with approximately the same firing rate as their inputs over the range of cortical firing frequencies. (The spiking threshold was 14.5 times the magnitude of the quantal input above the baseline for a time constant of 20 ms, and it was 2.9 times the magnitude of the quantal input above the baseline for a time constant of 1 ms). Despite its simplicity, this ‘integrate and fire’ model with balanced excitation and inhibition can account for several aspects of the firing statistics observed in the responses of neurons across large regions of the cerebral cortex (Shadlen & Newsome 1998).

By fixing the proportion of shared connections, but varying the mean firing rate according to which stimulus is present, rate coding in the presence of a fixed level of correlation can be examined. The mean firing rates chosen for each stimulus were extracted from data recorded from real inferior temporal cortical cells (see below). By holding the mean rate fixed (at the global mean corresponding to the previous case) and instead varying the proportion of shared connections according to the stimulus, one can examine a pure correlational code, thus allowing a clarification and test of the present information theoretical analysis.

3. RESULTS

(a) Information in neuronal ensemble responses in short time periods

In sufficiently short time windows, two spikes at most are emitted from the population. Taking advantage of this, the response probabilities can be obtained explicitly in terms of pairwise correlations: triplet and higher-order interactions do not contribute (see §2). Here we report the result of the insertion of the response probabilities obtained in this limit into the Shannon information formula (equation (1)): exact expressions quantifying the impact of pairwise correlations on the information transmitted by groups of spiking neurons. The information depends upon both the noise correlations γ and the signal correlations ν .

In the short time-scale limit, the first (I_t) and second (I_u) information derivatives suffice to fully describe the information kinetics. The instantaneous information rate is

$$I_t = \sum_{i=1}^C \left\langle \bar{r}_i(s) \log_2 \frac{\bar{r}_i(s)}{\langle \bar{r}_i(s') \rangle_{s'}} \right\rangle_s. \quad (8)$$

Note that the expression for I_t is just a generalization to the population level (a simple sum) of the expression previously derived for single cells (Bialek *et al.* 1991; Skaggs *et al.* 1993). The expression for the instantaneous information acceleration breaks up into three terms

$$\begin{aligned} I_u = & \frac{1}{\ln 2} \sum_{i=1}^C \sum_{j=1}^C \langle \bar{r}_i(s) \rangle_s \langle \bar{r}_j(s) \rangle_s \left[\nu_{ij} + (1 + \nu_{ij}) \ln \left(\frac{1}{1 + \nu_{ij}} \right) \right] \\ & + \sum_{i=1}^C \sum_{j=1}^C [\langle \bar{r}_i(s) \bar{r}_j(s) \gamma_{ij}(s) \rangle_s] \log_2 \left(\frac{1}{1 + \nu_{ij}} \right) \\ & + \sum_{i=1}^C \sum_{j=1}^C \left\langle \bar{r}_i(s) \bar{r}_j(s) (1 + \gamma_{ij}(s)) \right. \\ & \left. \times \log_2 \left[\frac{(1 + \gamma_{ij}(s)) \langle \bar{r}_i(s') \bar{r}_j(s') \rangle_{s'}}{\langle \bar{r}_i(s') \bar{r}_j(s') (1 + \gamma_{ij}(s')) \rangle_{s'}} \right] \right\rangle_s. \end{aligned} \quad (9)$$

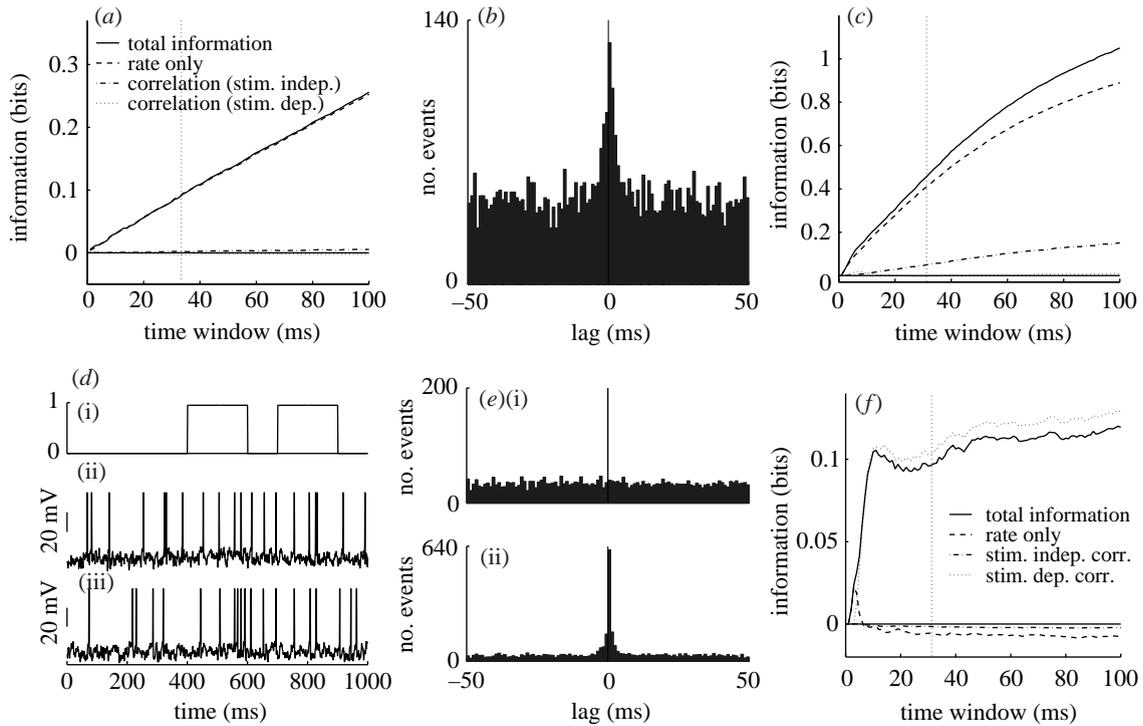


Figure 2. By simulating a quintuplet of cells, we were able to examine situations in which different information components were prominent. In each case shown, 100 trials of data were generated for each different stimulus. In the case of Poisson cells, all information is carried in the mean firing rates. This is shown in (a), in which the total information is compared with that from the rates only (the contributions of I_l and the first term of I_{ll} added together) and the additional information due to stimulus-independent correlations and stimulus-dependent correlations. (b,c) The result of an integrate and fire simulation in which common input is introduced by sharing of one-third of the connections for each pair of cells, giving the cross-correlogram shown in (b). The total information (thick solid line) and components of the information are shown versus the width of the measurement window in (c). The rate component is still dominant, but the stimulus-independent correlational component has a non-negligible effect. (d-f) A situation in which correlational information dominates: with a fixed mean firing rate, two of the five simulated cells (chosen randomly for that stimulus) increase their correlation by increasing the number of shared connections while the other two remained randomly correlated. The effect of this on cell spiking activity is shown in (d): panel (i) shows the fraction of shared connections, while panels (ii) and (iii) show the membrane potential and spike emission of the simulated cells. (e) The cross-correlograms corresponding to the low and high correlation states. The result of this is seen in (f): information due to correlations, although modest in magnitude, in this demonstration dominates the total information, as discussed in the text.

The first of these terms is all that survives if there is no noise correlation at all. Thus the rate component of the information is given by the sum of I_l (which is always greater than or equal to zero) and of the first term of I_{ll} (which is instead always less than or equal to zero). The second term is non-zero if there is some correlation in the variance to a given stimulus, even if it is independent of which stimulus is present; this term thus represents the contribution of stimulus-independent noise correlation to the information. The third component of I_{ll} represents the contribution of stimulus-modulated noise correlation, as it becomes non-zero only for stimulus-dependent correlations. We refer to these last two terms of I_{ll} together as the correlational components of the information.

The result just described shows that the short time-scale limit allows a rigorous quantification of the effect of correlations on the information conveyed by neuronal ensembles. The price to pay is the limited temporal range of applicability, as it formally requires that the mean number of spikes in the considered time window be small. What must now be addressed is the actual temporal range of applicability of the approach. The range of

validity of the (order t^2) approximation will depend on how well the information time dependence fits a quadratic approximation (e.g. for a quadratic function, the Taylor expansion to the second order would of course be exact for the entire t range). Because the assumption about the number of spikes will be broken first by the stimulus that gives the maximal response, a good scale for comparison with the range of validity is the minimum mean interspike interval to any stimulus. We studied the range of validity of the approximation, in the case of cells firing with Poisson statistics by direct calculation of the information; and in the more general case of up to four integrate and fire simulated neurons by a 'brute-force' calculation, using many trials of simulated data. Figure 1(a) shows the accuracy of the approximation for a pair of Poisson-simulated cells with mean-rate characteristics of real inferior temporal (IT) cells (Booth & Rolls 1998); the range of validity in this case would appear to be (being conservative), about six to eight times the minimum mean interspike interval. The approximation is better for smaller samples of cells. Scaling considerations suggest that the range of applicability should shrink as

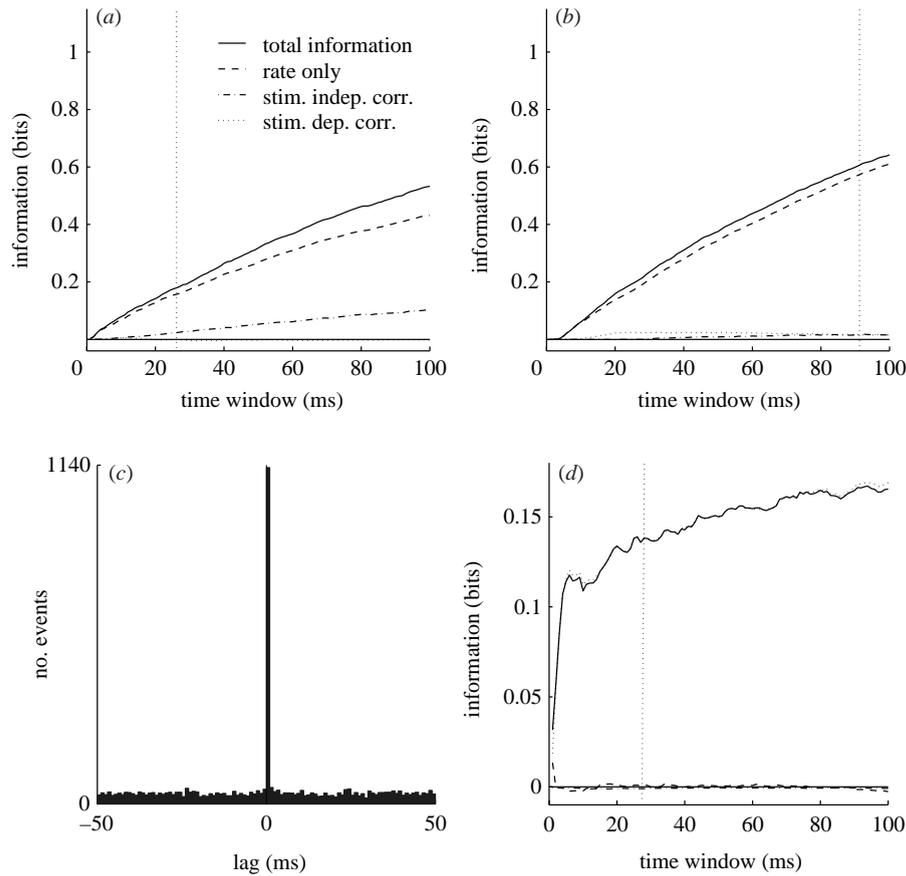


Figure 3. Illustration of the effect of changing population size (*a*), firing rates (*b*), and precision of synchrony (*c,d*). All symbols and conventions are as in figure 2. (*a*) The effect of reducing the population size from five to two cells in the integrate and fire simulation with 30% common inputs of figure 2(*b,c*). The total information is significantly reduced. In (*b*) we reduced by a factor of three the firing rates of the integrate and fire simulation with 30% common inputs of figure 2(*b,c*). The total information is reduced, and also the relative contribution of correlation becomes much less important at low rates. In (*c,d*) we simulated a correlational assembly with a constant firing rate of 20 Hz to all stimuli, and a percentage of shared connections of either 0% or 90% to different stimuli, as in figure 2(*d-f*). The membrane time constant is set to 1 ms. The cross-correlogram in the 90% shared connection state is plotted in (*c*). The information, shown in (*d*), is fully conveyed by the stimulus-dependent correlational component and, because of the increased precision of synchrony, is higher than in the case of the 20 ms membrane time constant (figure 2*f*).

the inverse of the number of cells for larger populations. This expectation is roughly confirmed by the results in figure 1(*b*), where five Poisson cells (again with typical cortical firing rates) are analysed. Integrate and fire simulations confirm that this estimate of the range of validity is relatively robust to neuronal statistics and is conserved across a wide range of response correlation values. Figure 1(*c*) shows an estimate of the allowed ensemble size versus the time window for cells with mean rates to stimuli extracted from the cortical data of Booth & Rolls (1998). We conclude that the analysis is pertinent for time-scales relevant to neuronal coding for ensembles of up to around 10–15 cells with firing rates similar to those in the inferior temporal cortex, and for even larger ensembles of cells with lower firing rates, such as medial temporal lobe cells.

(*b*) *Rate and correlation components of the information*

So what is the impact of the second-order terms on the information conveyed by the ensemble of neurons? Figure 2 provides an illustration of their effect. In the case of non-interacting cells with zero autocorrelation,

the information is carried entirely in the rate component of the information. This was exemplified by simulating a quintuplet of cells, each of which fired spikes according to a Poisson process (at any instant in time there is a given probability of firing which is constant in time throughout the experimental trial), independently of other cells, with a mean rate which was different for each of ten stimuli. The mean rates of each of the five cells to each stimulus were taken from real cells described in Booth & Rolls (1998); these real cells were the ones analysed as simultaneously recorded pairs later in the paper, facilitating comparison with the real data examined. The result is shown in figure 2(*a*), which, like the other graphs, shows separately the contributions of the three terms to the total information. The last two of these (the contributions of correlations) are negligible in the example of figure 2(*a*), because the spike trains were by design uncorrelated.

By simulating spike trains with the integrate and fire model, we were able to examine situations in which the correlational components of the information were not negligible. The first such case considered was that of a quintuplet of neurons which had a large amount of

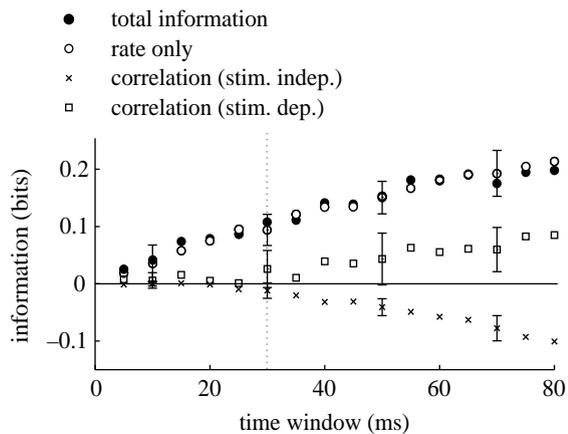


Figure 4. Application of the procedure to real neurophysiological data. The information contained in the responses of three pairs of primate inferior temporal cortical neurons about viewed objects is broken down into three components: rate, stimulus-independent noise correlation, and stimulus-modulated noise correlation. The vertical dotted line again shows for time-scale comparison the mean interspike interval to the maximally responsive stimulus.

common input (one-third of their connections were shared). These neurons fired spikes in response to a balance of excitatory and inhibitory input (see § 2), with a membrane decay time of 20 ms leading to a small amount of correlated firing, as shown by the cross-correlogram in figure 2(b). Ten stimuli again induced different mean firing rate responses from the cells. The correlation is not stimulus dependent, and therefore the third component of I_u is still zero in this case. The second component of I_u , representing the effect of stimulus-independent correlation, can, if required, again be broken up into auto- ($i = j$) and cross-correlation ($i \neq j$) parts. The elements corresponding to autocorrelation are always positive for autocorrelations γ_{ii} between -1 and 0 (as observed for the IT cells), whereas for a given (i, j) , the cross-correlation component is positive when, for positive noise correlation, the signal correlation is negative, i.e. when the cells anti-covary in their stimulus response profiles. In the case shown in figure 2(c), this term is positive, and leads to a modest increase in the total information that can be transmitted.

To model a situation where stimulus-dependent correlations conveyed information, we generated simulated data using the integrate and fire model for another quintuplet of cells which had a stimulus-dependent fraction of common input. This might correspond to a situation where cells transiently participate in different neuronal assemblies, depending on stimulus conditions. There were again ten stimuli, but this time the mean spike emission rate to each stimulus was constant at *ca.* 20 Hz, the global mean firing rate for the previous case. One of these stimuli simply resulted in independent input to each of the model cells, whereas each of the other nine stimuli resulted in an increase (to 90%) in the amount of shared input between one pair of cells (chosen at random from the ensemble such that each stimulus resulted in a different pair being correlated). The response of one such pair to changes in the amount of common input is shown

in figure 2(d). Panel (i) shows the fraction of shared connections as a function of time; panels (ii) and (iii) show the resulting membrane potentials and spike trains from the pair of neurons. During the high correlation state, there was, on average, a seven times higher probability of a coincidence in any 10 ms period than chance. This cross-correlation is also evident in the cross-correlograms shown in figure 2(e). The results are given in figure 2(f): all terms but the third of I_u are essentially zero, and information transmission, in this case, almost entirely due to stimulus-dependent correlations. The total amount of information that could be conveyed, even with this much shared input, was modest in comparison to that conveyed by rates dependent on the stimuli, at the same mean firing rate. The total information increased slightly if for the 'high correlation' state the spike trains were nearly perfectly correlated; if the 'low correlation' state corresponded not to chance, but to actual anticorrelation, then it is possible that even more information could be conveyed.

To illustrate the impact of the population size and of the overall average firing rate, we performed more simulations, again using the spike trains with the integrate and fire model with a time constant of 20 ms and 30% shared connections, as in figure 2(b,c). In figure 3(a) we tested the effects of population size by computing the information from pairs of cells instead of from five cells, by extracting pairs from the set of five cells and then averaging across pairs. The effect of reducing the size of the population from five to two cells was to reduce the total information by a factor of about two. It is, of course, expected that the amount of information depends on the number of cells in the population. It is also shown with this particular set of generated spikes that the relative amount of information in the different components is approximately similar. In figure 3(b) we show the effects of applying the methods to cells with lower firing rates. The lower firing rates were produced simply by dividing the rates by three relative to those used in figure 2(c). (This places the firing rates of the cells in the same regime as that of primate hippocampal pyramidal neurons *in vivo*, which fire at lower rates than cortical visual cells (Rolls *et al.* 1997a).) The information was computed from the responses of the population of five simulated cells. The effect of dividing the overall firing rate by three was to reduce the information by approximately two. If only the first derivative was important, we would expect the total information to decrease linearly when decreasing the rate. The sublinear decrease is mostly due to the rate component of the second derivative, which is a negative contribution, larger for higher rates and much smaller for lower rates. Similarly, the correlational component of the information is also much less important for lower rates, as evident from our analysis.

To illustrate the effects of precision of synchrony, we produced spike trains by simulating a correlational assembly with a constant firing rate of 20 Hz to all stimuli, and a percentage of shared connections of either 0% or 90% to different stimuli. In order to increase the precision of the synchrony with respect to figure 2(d-f), we decreased the membrane time constant to the value of 1 ms. This gives, in the high correlation state, a very

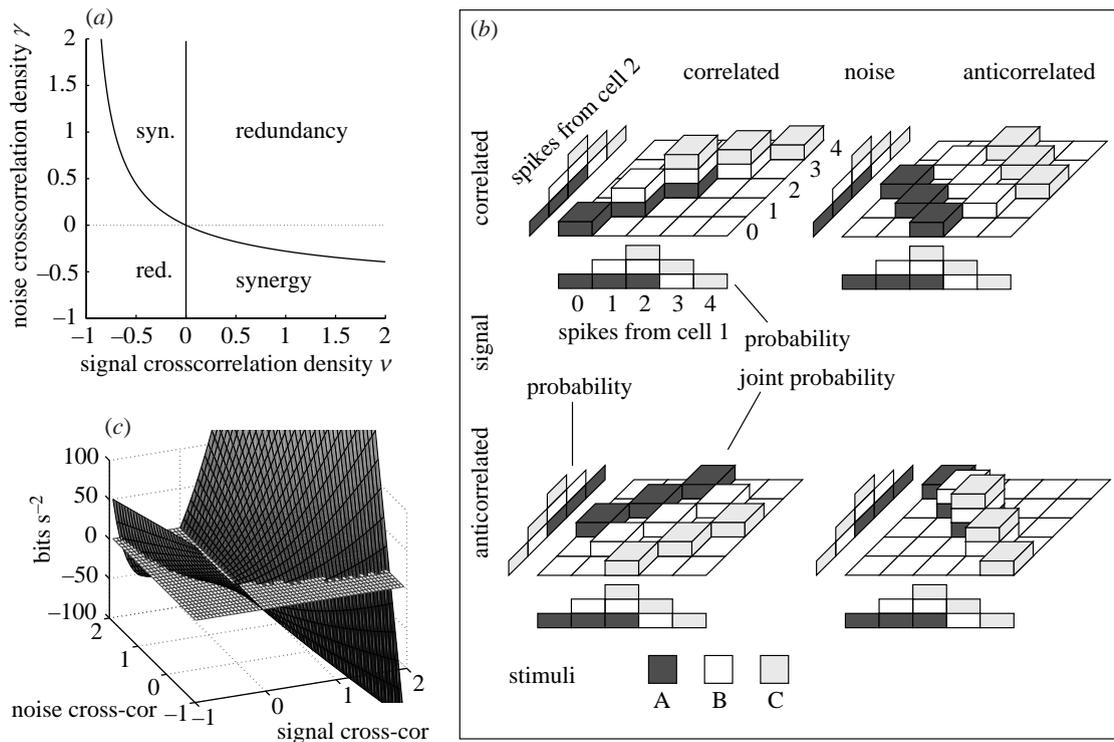


Figure 5. Redundancy and synergy. (a) Consideration of a pair of neurons, in which the noise correlation is taken to be stimulus independent, illustrates the range of possibilities. The vertical axis plots the noise correlation measured by the scaled cross-correlation density γ_{12} . The horizontal axis shows the correlation in the signal, also expressed in terms of the scaled cross-correlation density ν_{12} . If the mean responses of each neuron to the various stimuli are uncorrelated, $\nu_{12} = 0$, then (no matter what the noise correlation γ_{12}) the redundancy is exactly zero. This is the case along the vertical solid line. If the pair of neurons covary in the mean number of spikes they fire to each stimulus, then all else being equal, we have a redundant situation. If, however, there is enough anticorrelation in the noise, or vice versa there is anticorrelation in the signal and enough positive correlation in the noise, then the coding becomes synergistic. (b) Graphical depiction of the interplay between signal and noise correlation. Histograms show the probability of each cell emitting from zero to four spikes. The joint probability matrix of spike emissions from both cells is shown between the histograms for each cell. Contributions from each stimulus are labelled by shading the corresponding probability block. (c) The redundancy (divided by the square of the time window) corresponding to the above plane. Values below the zero axis indicate synergistic encoding.

precise 1 ms synchrony between the spike trains, as shown by the cross-correlogram in figure 3(c). The information shown in figure 3(d) is fully conveyed by the correlational component with no information in the firing rates, correctly reflecting the way in which the spikes were generated. The actual information with the 1 ms precision is higher by a factor of 1.5 with respect to the 20 ms membrane time constant (figure 2(f)). Thus the precision of spike synchrony has an impact on the correlational part of the information, as predicted by our analysis.

(c) Application of the method to real neuronal data

To demonstrate that the technique is applicable to real data, we applied the information component analysis to three pairs of cortical cells recorded simultaneously from the same electrode from the data set of Booth & Rolls (1998). These data came from an experiment in which the responses of neurons in the inferior temporal cortex of the macaque monkey were recorded while one of ten different objects was being viewed. The firing rates of these neurons were found by Booth & Rolls (1998) to convey information about which object was present, regardless of viewing angle. Four different views of each object were each presented five times, resulting in a total of twenty

trials for each object. The median firing rate to the best stimulus of the cells used in our analysis was 44 Hz. The same data were used to estimate realistic firing rates for all the simulations described in this paper.

Each component of the information for the cells in this data set was calculated for each pair of cells, and the standard error in the calculation of each information component was obtained by error propagation from the variances in the measurements of spike counts and coincidences over the trials. The average information encoding characteristics of the three pairs (which were qualitatively similar to each other) are shown in figure 4. In this example it is clear that the rate component is prominent in the information representation; however, the important point to gain from this result is that the information expansion provides a demonstrably practical method for analysing simultaneously recorded data with only a small number of experimental trials. The values of each component and an estimate of its error are available both for single small ensembles, and for the 'average picture' obtained by recording large numbers of such small neuronal ensembles. This encourages us to think that it will be possible to analyse simultaneously recorded data in a systematic and rigorous quantitative manner. Of course, the populations of neurons that actually act effectively together are larger

than can be studied using this, or any other presently known procedure. However, it is reasonable to assume that effects present in such large ensembles will be to some extent observable in the smaller ensembles that one can in practice record from and analyse.

DISCUSSION

(a) *Redundancy versus synergy*

A crucial point in understanding the representation of external stimuli by the activity of a population of neurons, is how the information conveyed by individual cells combines together. If many cells in the sample carry similar information, then the code is redundant, and there is not much more information in the population than that present in single cells. Another regime is where the information from different cells is independent, in which case the information increases linearly with the number of cells in the population. Yet another regime is one in which some information is available only by looking simultaneously at the responses of different neurons, with, for example, some of the information being available only from the relative timing of firing of different cells. This code is said to be synergistic, and more information is available in the population than one would obtain by the sum of the information obtained from each neuron alone. In mathematical terms, the redundancy of a code can be defined (and measured in bits) as the amount of information that would be obtained by adding the information from every cell as if each was independent minus that obtained by considering the whole neuronal ensemble (Rieke *et al.* 1996). For synergistic coding, the value of this Shannon redundancy is, of course, negative.

Equation (9) shows that overall correlations in the distribution of mean responses alone can only lead to redundancy, because all (i, j) contributions to the first term of I_H are negative, except that they are zero when there is no overall 'signal' correlation in the mean response profiles (i.e. $\nu_{ij} = 0$). To have synergistic coding of information one needs correlations in the variability of the responses (to a given stimulus), i.e. non-zero γ parameters. Even when such 'noise' correlations are independent of the stimuli, however, it is possible to have synergy. This can be demonstrated by considering the sign of the Shannon redundancy, obtained from equation (1) by subtracting the information conveyed by the population from the sum of that carried by each single cell. This shows four basic regimes of operation, illustrated in the two-cell example of figure 5. If the cells anti-covary in their response profiles to stimuli, they must have a positive noise correlation, above the boundary value depicted in figure 5(a), to obtain synergy; or if the cells do have positive signal correlation, then coincidences must be actively suppressed by a negative noise correlation stronger than the corresponding boundary value. When the signal and noise correlations have the same sign, one always obtains redundancy in the short time-scale limit. Clearly, already with pairs of cell, the interplay between correlation in the noise and correlation in the signal introduces a potential for both redundancy and synergy.

A simple example of the interplay between signal and noise correlation in a pair of cells is graphically depicted in figure 5(b). In this example there are three stimuli, A,

B and C, which occur with equal probability. The first cell emits a different mean number of spikes to each stimulus: one to A, two to B, and three to C. Let us consider cases where the second cell either fires the same mean number of spikes to each stimulus (full correlation), or alternatively has the mean responses to A and C exchanged (anticorrelation with $\nu_{12} = -1/6$). On any given occasion, noise adds zero or ± 1 spikes, with equal (one-third) probability, to the output of the cell. It is easy to check that the information carried by each cell alone is $(2/3)\log_2 3 - (4/9)$ bits, which is less than half the $H = \log_2 3$ maximum information (entropy) available in the stimulus set.

Figure 5(b) illustrates the four different situations for the second cell, depending on the way in which the mean (signal) and variance (noise) of the number of spikes it fires are related to those of the first cell. When the signal and noise are both correlated (or both anticorrelated), the joint probabilities of the numbers of spikes fired by each cell tend to bunch up along the diagonal, so that the total information from both cells is less than the sum of that obtained from each cell on its own. If, however, with correlated signal the noise is anticorrelated or if with correlated noise, the signal is anticorrelated, then the joint probabilities are more spread out. This means that the presented stimulus can be clearly identified on the basis of the joint response observed. If this spread is sufficient, it is possible to obtain synergy: the information calculated from the joint probability matrix exceeds the sum of that obtained for each cell individually. In this example, with the signal having $\nu_{12} = -1/6$ (anticorrelation), if the noise was uncorrelated between the cells, then the redundancy can be easily seen to be 0.04 bits; it is only when the noise correlation is increased above $\gamma_{12} = 0.10$ that the coding is synergistic.

Examination of equation (9) and figure 5(c) reveals the total amount of redundancy (or synergy) to be much more sensitive to noise correlation when the signal correlation is high, for example, the noise correlation leads to large redundancy if the cells are tuned to the same stimulus. If the signal correlation is small, the redundancy is close to zero, no matter how correlated the noise. This explains why the impact of noise correlation on the performance of pools of neurons has been emphasized in experiments using simple one-dimensional discrimination tasks (in which neurons from a local pool have been found to be tuned to the same stimulus (Zohary *et al.* 1994; Parker & Newsome 1998)), while noise correlation has been described as less important for groups of neurons coding for complex stimuli (Gawne & Richmond 1993; Gawne *et al.* 1996), which tend to use a more distributed encoding. This study therefore shows rigorously that correlations do not necessarily invoke redundancy, and that it is not possible in general to estimate the 'effective' number of neurons participating in encoding by simply measuring the noise correlation, as done by Zohary *et al.* (1994). Further, because the signal correlation was reported to decrease towards zero when increasing the complexity of the stimulus set used to test the neurons (Gawne *et al.* 1996; Rolls *et al.* 1997b), it predicts that noise correlation might have only a small impact in the encoding of large sets of natural stimuli, at least as far as stimulus-independent noise correlation is concerned. The

effect of small stimulus-dependent correlations in the noise is considered below, and it should not be neglected.

The possibility of synergy with constant correlations has been raised previously (Oram *et al.* 1998) as a phenomenological observation; the information expansion we have introduced places this phenomenon on a solid mathematical footing, and delineates, for short time windows, the exact boundaries of the regions of synergy and redundancy. The discussion by Oram *et al.* (1998) points out that correlations might lead to synergy when neurons are tuned to different stimuli, but it is not able to bridge between different encoding situations and predict the exact amounts of redundancy or synergy that occur, as our analysis does. The analysis of synergy presented here, unlike that of Oram *et al.* (1998), generalizes to the case of stimulus-modulated correlations: it is enough to specify the stimulus dependence of correlations and take into account also the third term of I_u .

(b) *A null hypothesis for the role of correlations in the cerebral cortex*

How might typical correlations among cells in a population scale up with the size of the population? Clearly, this is a question for experiments to address, in fact a most crucial question for those investigating correlations in neural activity. Such experiments need to be well designed, accurate and systematic. It is easy to see, from the analysis above, that with large populations even small correlations could produce extreme effects, resulting in either large redundancy or (perhaps less often) very substantial synergy. This is because, while with C cells there are C first-order terms in the information, there are obviously C^2 second-order terms, C^3 third-order ones (which depend also on three-way correlations), and so on. Depending on tiny details of the correlational structure, successive terms can affect transmitted information in both directions. Within such broad a realm of possibilities, it is then of interest to try to formulate a sort of null hypothesis, that might provide at least a reference point against which to contrast any more structured candidate theory. One example is the scaling behaviour we might expect if the correlations were not playing any special role at all in the system or area being analysed. In this 'null' hypothesis, the parameters ν_{ij} would be expected to be small, that is to deviate from zero (no correlations) only in so far as the set of stimuli used is limited (Gawne & Richmond 1993; Rolls *et al.* 1997b); similarly the stimulus-dependent noise correlation $\gamma_{ij}(s)$ would be small. The scaling behaviour corresponding to this null hypothesis can be examined by further expanding I_u as a series in these new small parameters: at times of the order of the interspike interval, second-order terms of order $-C^2\langle(\nu)^2\rangle$ (redundant) and $+C^2\langle(\gamma)^2\rangle$ (synergistic) are introduced (the angular brackets indicating the average value). If we have a large enough population of cells, and $\langle(\nu)^2\rangle$ and $\langle(\gamma)^2\rangle$ are not sufficiently small to counteract the additional C factor, these 'random' redundancy and synergy contributions will be substantial. Obviously, in this situation our expansion would start to progressively fail in quantifying the information, because higher-order terms in the t expansion become more and more important in this case. But this clearly shows that a sufficiently large population of cells, which has not been

designed to code stimuli in any particular cooperative manner, has the potential to provide large effects of redundancy or synergy, arising simply from random correlations among the firing of the different cells. This reinforces the need for systematic study of the magnitude and scaling of correlations in the cerebral cortex.

(c) *Correlational assemblies and mutual information*

If cells participate in context-dependent correlational assemblies (Singer *et al.* 1997), a significant amount of information should be found in the third component of I_u , relative to the total information, when analysing data obtained from the appropriate experiments. The challenge for the establishment of correlational theories of neural coding has thus been laid down: to demonstrate quantitatively how substantial a proportion of the information about external correlates is provided by correlations between cells, given the large amount of information that has been shown in some neural systems to be coded by rate (Rolls & Treves 1998). The second-order series expansion we have described allows precisely this to be achieved for small ensembles of cells—for a time window of 20 ms, an ensemble of about 10–15 cells which fired at a peak mean rate to a stimulus of around 50 Hz (e.g. IT cells) could be analysed; with the same time window, 25–30 cells firing at a lower peak mean rate of around 20 Hz (such as neurons from the medial temporal lobe) could be studied. Beyond this population size the information expansion can still be of use in picking up the correlational variables conveying most of the information in small subpopulations, and eliminating their relevant variables. Reduction of the response space of a large population of cells to a treatable size thus may be possible without losing salient features.

In order to test hypotheses about the role of correlations in solving the binding problem (Von der Malsburg 1995; Gray *et al.* 1992; Singer *et al.* 1997), as opposed to other solutions (Treisman & Gelade 1980; Treisman 1996), and about information coding in general (Vaadia *et al.* 1995; deCharms & Merzenich 1996), careful quantitative experimental studies of the correlations prevailing in the neural activity of different parts of the brain are needed. Data analyses based on the time-expansion approach then have the potential to elucidate the role of correlations in the encoding of information by cortical neurons.

We thank R. Baddeley, A. Renart and D. Smyth for useful discussions. We thank M. Booth for kindly providing neurophysiological data. This research was supported by an EC Marie Curie Research Training Grant ERBFMBICT972749 (S.P.), a studentship from the Oxford McDonnell–Pew Centre for Cognitive Neuroscience (S.R.S.), and by MRC Programme Grant PG8513790.

APPENDIX A. CORRELATION MEASURES

The numerical value of the information calculated is independent of the correlation measure used, and we show here how to express the results using the Pearson correlation coefficients and Fano factors instead of scaled cross-correlation coefficients. In the text, we chose to use

the scaled cross-correlation measure of equation (3) because it produces a more compact mathematical formulation of what is addressed in this paper, and has useful scaling properties as the time window becomes small.

A widespread measure for cross-correlation is the Pearson correlation coefficient $\rho_{ij}(s)$, which normalizes the number of coincidences above independence to the standard deviation of the number of coincidences expected if the cells were independent. The normalization used by the Pearson correlation coefficient quantifies the strength of correlations between neurons in a rate-independent way. However, it should be noted that the Pearson noise-correlation measure approaches zero at short time windows:

$$\rho_{ij}(s) \equiv \frac{\overline{n_i(s)n_j(s)} - \bar{n}_i(s)\bar{n}_j(s)}{\sigma_{n_i(s)}\sigma_{n_j(s)}} \\ \approx t \gamma_{ij}(s) \sqrt{r_i(s)\bar{r}_j(s)},$$

where $\sigma_{n_i(s)}$ is the standard deviation of the count of spikes emitted by cell i in response to stimulus s .

Under assumption (3), $\gamma_{ij}(s)$ remains finite as $t \rightarrow 0$; thus by using this measure we can keep the t expansion of the information explicit. This greatly increases the amount of insight obtained from the series expansion.

Similarly, an alternative to scaled autocorrelation density $\gamma_{ii}(s)$ for the measure of autocorrelations is the so called Fano factor F , that is the variance of the spike count divided by its mean (Rieke *et al.* 1996). This measure is used in neurophysiology because for the renewal process, often used as a stochastic model of neuronal firing, the variance is proportional to the mean. (Fano factors lower than unity indicate that the process is more regular than a Poisson process). F grows linearly with t for short times: $F = 1 + t \bar{r}_i(s) \gamma_{ii}(s)$. Again, we prefer $\gamma_{ii}(s)$ to the Fano factor in the information expansion because $F - 1$ approaches zero for short times.

To express the information derivatives in terms of Pearson correlation coefficients and Fano factor, it is enough to make the following simple substitutions in equation (9):

$$\gamma_{ij}(s) \rightarrow \frac{\rho_{ij}(s)}{t \sqrt{r_i(s)\bar{r}_j(s)}}, \\ \gamma_{ii}(s) \rightarrow \frac{F - 1}{t \bar{r}_i(s)}.$$

We note that the 'scaled cross-correlation measure' γ_{ij} is sensitive to the mean firing rate, as the strength of neuronal interactions might be overemphasized at low rate (Aertsen *et al.* 1989): it cannot be taken to be a linear measure of interaction strength. However, the value of the information transmitted by the the number of spikes simultaneously fired by each cell, and of each particular component, depends on the response probabilities, and not on the particular way chosen to quantify the correlations. Therefore, the particular measure used for correlations is for this application ultimately a matter of mere notation.

REFERENCES

Aertsen, A. M. H. J., Gerstein, G. L., Habib, M. K. & Palm, G. 1989 Dynamics of neuronal firing correlation: modulation of 'effective connectivity'. *J. Neurophysiol.* **61**, 900–917.

Amit, D. J. 1997 Is synchronization necessary and is it sufficient? *Behav. Brain Sci.* **20**, 683.

Arieli, A., Sterkin, A., Grinvald, A. & Aertsen, A. 1996 Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* **273**, 1868–1871.

Bialek, W., Rieke, F., de Ruyter van Steveninck, R. R. & Warland, D. 1991 Reading a neural code. *Science* **252**, 1854–1857.

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**, 1047–3211.

Carr, C. E. 1993 Processing of temporal information in the brain. *A. Rev. Neurosci.* **16**, 223–243.

Carr, C. E. & Konishi, M. 1990 A circuit for detection of interaural time differences in the brainstem of the barn owl. *J. Neurosci.* **10**, 3227–3246.

Cover, T. M. & Thomas, J. A. 1991 *Elements of information theory*. New York: John Wiley.

deCharms, R. C. & Merzenich, M. M. 1996 Primary cortical representation of sounds by the coordination of action potentials. *Nature* **381**, 610–613.

Donoghue, J. P., Sanes, J. N., Hatsopoulos, N. G. & Gaal, G. 1998 Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J. Neurophysiol.* **79**, 159–173.

Gawne, T. J. & Richmond, B. J. 1993 How independent are the messages carried by adjacent inferior temporal cortical neurons? *J. Neurosci.* **13**, 2758–2771.

Gawne, T. J., Kjaer, T. W., Hertz, J. A. & Richmond, B. J. 1996 Adjacent visual cortical complex cells share about 20% of their stimulus-related information. *Cerebr. Cort.* **6**, 482–489.

Golledge, D. R., Hildetag, C. C. & Tovée, M. J. 1996 A solution to the binding problem? *Curr. Biol.* **6**, 1092–1095.

Gray, C. M., Engel, A. K., König, P. & Singer, W. 1992 Synchronization of oscillatory neuronal responses in cat striate cortex: temporal properties. *Vis. Neurosci.* **8**, 337–347.

Heller, J., Hertz, J. A., Kjaer, T. W. & Richmond, B. J. 1995 Information flow and temporal coding in primate pattern vision. *J. Comp. Neurosci.* **2**, 175–193.

Kitazawa, S., Kimura, T. & Yin, P.-B. 1998 Cerebellar complex spikes encode both destinations and errors in arm movements. *Nature* **392**, 494–497.

Macknik, S. L. & Livingstone, M. S. 1998 Neuronal correlates of visibility and invisibility in the primate visual system. *Nature Neurosci.* **1**, 144–149.

Meister, M., Lagnado, L. & Baylor, D. A. 1995 Concerted signalling by retinal ganglion cells. *Science* **270**, 1207–1210.

Murthy, V. N. & Fetz, E. E. 1996 Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *J. Neurophysiol.* **76**, 3968–3982.

Optican, L. M. & Richmond, B. J. 1987 Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J. Neurophysiol.* **57**, 162–178.

Oram, M. W. & Perrett, D. I. 1992 Time course of neuronal responses discriminating different views of face and head. *J. Neurophysiol.* **68**, 70–84.

Oram, M. W., Földiák, P., Perrett, D. I. & Sengpiel, F. 1998 The 'Ideal Homunculus': decoding neural population signals. *Trends Neurosci.* **21**, 259–265.

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

Parker, A. J. & Newsome, W. T. 1998 Sense and the single neuron: probing the physiology of perception. *A. Rev. Neurosci.* **21**, 227–277.

- Riehle, A., Grun, S., Diesmann, M. & Aertsen, A. M. H. J. 1997 Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* **278**, 1950–1953.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R. R. & Bialek, W. 1996 *Spikes: exploring the neural code*. Cambridge, MA: MIT Press.
- Rolls, E. T. & Tové, M. J. 1994 Processing speed in the cerebral cortex and the neurophysiology of visual masking. *Proc. R. Soc. Lond. B* **257**, 9–15.
- Rolls, E. T. & Treves, A. 1998 *Neural networks and brain function*. Oxford University Press.
- Rolls, E. T., Robertson, R. G. & Georges-François, P. 1997a Spatial view cells in the primate hippocampus. *Eur. J. Neurosci.* **9**, 1789–1794.
- Rolls, E. T., Treves, A. & Tové, M. J. 1997b The representational capacity of the distributed encoding of information provided by populations of neurons in primate temporal visual cortex. *Exp. Brain Res.* **114**, 149–162.
- Shadlen, M. N. & Newsome, W. T. 1998 The variable discharge of cortical neurons: implications for connectivity, computation and coding. *J. Neurosci.* **18**, 3870–3896.
- Shannon, C. E. 1948 A mathematical theory of communication. *AT&T Bell Lab. Tech. J.* **27**, 379–423.
- Singer, W., Engel, A. K., Kreiter, A. K., Munk, M. H. J., Neuenschwander, S. & Roelfsema, P. 1997 Neuronal assemblies: necessity, signature and detectability. *Trends Cogn. Sci.* **1**, 252–261.
- Skaggs, W. E., McNaughton, B. L., Gothard, K. & Markus, E. 1993 An information theoretic approach to deciphering the hippocampal code. In *Advances in neural information processing systems*, vol. 5 (ed. S. Hanson, J. Cowan & C. Giles), pp. 1030–1037. San Mateo: Morgan Kaufmann.
- Thorpe, S., Fize, D. & Marlot, C. 1996 Speed of processing in the human visual system. *Nature* **381**, 520–522.
- Tové, M. J., Rolls, E. T., Treves, A. & Bellis, R. P. 1993 Information encoding and the response of single neurons in the primate temporal visual cortex. *J. Neurophysiol.* **70**, 640–654.
- Treisman, A. 1996 The binding problem. *Curr. Opin. Neurobiol.* **6**, 171–178.
- Treisman, A. & Gelade, G. 1980 A feature integration theory of attention. *Cogn. Psychol.* **12**, 97–136.
- Vaadia, E., Haalman, I., Abeles, M., Bergman, H., Prut, Y., Slovin, H. & Aertsen, A. 1995 Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* **373**, 515–518.
- Von der Malsburg, C. 1995 Binding in models of perception and brain function. *Curr. Opin. Neurobiol.* **5**, 520–526.
- Werner, G. & Mountcastle, V. B. 1965 Neural activity in mechanoreceptive cutaneous afferents: stimulus-response relations, Weber functions, and information transmission. *J. Neurophysiol.* **28**, 359–397.
- Zohary, E., Shadlen, M. N. & Newsome, W. T. 1994 Correlated neuronal discharge rate and its implication for psychophysical performance. *Nature* **370**, 140–143.

As this paper exceeds the maximum length normally permitted, the authors have agreed to contribute to production costs.

An electronic appendix to this paper can be found at (http://www.pubs.royalsoc.ac.uk/publish/pro_bs/rpb1423.htm).