

How Independent Are the Messages Carried by Adjacent Inferior Temporal Cortical Neurons?

Timothy J. Gawne and Barry J. Richmond

Laboratory of Neuropsychology, National Institute of Mental Health, Bethesda, MD 20892

There are at least three possibilities for encoding information in a small area of cortex. First, neurons could have identical characteristics, thus conveying redundant information; second, neurons could give different responses to the same stimuli, thus conveying independent information; or third, neurons could cooperate with each other to encode more information jointly than they do separately, that is, synergistically. We recorded from 28 pairs of neurons in inferior temporal cortex of behaving rhesus monkeys. Each pair was recorded from a single microelectrode. Both the magnitude and the temporal modulation of the responses were quantified. We separated the responses into signal (average response to each stimulus) and noise (deviation of each response from the average). Linear regression showed that an average of only 18.7% of the magnitude of the signal carried by one neuron could be predicted from the magnitude of the other, and only 22.0% could be predicted by including the temporal modulation. For the noise, the figures were 5.5% and 6.3%, respectively, even less than for the signal. Information theoretic analysis shows that the pairs of neurons we studied carried an average of 20% redundant information. However, even this relatively small amount of redundancy places a severe upper limit on the information that can be transmitted by a neuronal pool. A pool of neurons for which each pair is mutually redundant to extent y can only carry a maximum of $1/y$, here five times, as much information as one neuron alone. Information theoretic analysis gave no evidence for the presence of information as a function of both neurons considered together, that is, synergistic codes. Cross-correlation showed that at least 61% of the neuronal pairs shared connections in some manner. Given these shared connections, if adjacent neurons had had identical characteristics, then the noise on the outputs of these neurons would have been highly correlated, and it would not be possible to separate the signal and noise. The severe impact of correlated noise and information redundancy leads us to propose that the processing carried out by these neurons evolved both to provide a rich description of many stim-

ulus properties and simultaneously to minimize the redundancy in a local group of neurons. These two principles appear to be a major constraint on the organization of inferior temporal, and possibly all, cortex.

[Key words: information, temporal modulation, coding, signal, noise, vision]

Previous studies have shown that single neurons in the visual system have a limited capacity to carry information (McClurkin et al., 1991b). A single neuron by itself does not have the capacity to explain perception or behavior, which must be a property of the simultaneous functioning of many neurons. The spread of fibers terminating in a local area suggests that adjacent neurons, here defined as neurons recorded from the same microelectrode, receive many common inputs. This organization could result in adjacent neurons processing data in the same or very similar ways, but the actual relationship of the anatomical organization to the messages being carried is not known. Without understanding how the messages carried by adjacent cortical neurons are related, it is not possible to determine how information is processed and transmitted in the nervous system.

The physiological relations between neurons in a local group have most often been studied through the use of the cross-correlation technique (Eggermont et al., 1983; Melssen and Epping, 1987; Krüger and Aiple, 1988). Cross-correlation techniques use the relationship between the times of occurrence of individual spikes from pairs of neurons to define effective connectivities (Melssen and Epping, 1987). However, Melssen and Epping (1987) point out that unique anatomical patterns of connections cannot be determined using these techniques. In addition, cross-correlation techniques are not designed to explore whether the information carried by pairs of neurons is related.

Considering that adjacent neurons probably share common afferents, one might expect them to contain similar information in their responses. If they carry the same message with uncorrelated noise, then the signal-to-noise ratio would increase with the square root of the number of neurons. If the number of neurons is relatively large, then the message would be resistant to the loss of one or a few neurons. However, if the noise were correlated, then averaging the responses of many neurons would not improve the signal-to-noise ratio, and the only advantage in having many neurons would be resistance to cell loss.

Another possibility is that neurons could perform quite different operations on the messages coming in from their afferents. Then adjacent neurons could carry messages that were independent or nearly so. Each message would have a relatively low signal-to-noise ratio, but the total amount of information would

Received Sept. 17, 1992; revised Nov. 30, 1992; accepted Dec. 28, 1992.

We thank Dr. Bruce Smith and the staff of the Research Services Branch for their work on the data acquisition hardware, Emad Eskandar for his valuable help in the laboratory, and Dr. Karen Pettigrew of the Statistics and Mathematics branch of the NIMH and Dr. John Hertz of NORDITA and NIMH for their valuable advice on the analysis and on the manuscript. This work was supported in part by Air Force Office of Scientific Research Grant AFOSR-ISSA-88-0073.

Correspondence should be addressed to Dr. Barry J. Richmond, Laboratory of Neuropsychology, National Institute of Mental Health, Building 9, Room 1E104, 9000 Rockville Pike, Bethesda, MD 20892.

0270-6474/93/132758-14\$05.00/0

be equal to the sum of that carried by the individual neurons. This scheme would preserve a complete description of many aspects of a stimulus in a relatively small volume of cortex. In addition, the noise would be uncorrelated even if the neurons shared all inputs.

Finally, it is possible that neurons could carry messages that are present only in their responses considered jointly, that is, synergistically. For example, the difference between the responses of a pair of neurons might be very precise. In this case it would not be possible to interpret fully the response of a single neuron without knowing the responses of the neighboring neurons as well.

These different encoding schemes imply different roles for adjacent neurons, and for the nature of information processing in the nervous system. To determine which, if any, of these encoding schemes is used by neuronal groups in inferior temporal (IT) cortex, we studied visually elicited responses of pairs of neurons recorded from a single microelectrode. Correlating the responses of pairs of neurons would confound the effects of signal and noise. Therefore, we separated the neuronal responses into a signal component (average response to each stimulus) and a noise component (difference from the average response). This allows the signal and noise to be dealt with separately.

Information is always calculated from a specific stimulus set. For the results of any experiment to be interpreted as generic, it is important to use a rich and varied stimulus set. With only a few stimuli, we would not have been able to determine whether two neurons generally transmitted the same message (i.e., carry redundant information) or not (i.e., carry independent information). For example, if we had used only two stimuli, the responses of any pair of adjacent neurons would have been perfectly correlated, because two points can always be connected with a straight line, and thus the results would have been biased. Therefore, we used 32 orthogonal stimuli based on the Walsh transform (Ahmed and Rao, 1975; Harmuth, 1977; Richmond et al., 1987). Because the firing of neurons in IT cortex depends upon the behavioral conditions, the stimuli were presented during a simple match-to-sample memory task. In this task each stimulus can take on three different behavioral meanings: sample, match, and nonmatch. Thus, our full stimulus set consisted of 96 unique combinations.

Using linear regression and information theory, we have shown that the signals carried by adjacent neurons in IT cortex are largely uncorrelated and independent, and that the noise is more uncorrelated than the signal. As we pointed out above, there is a great potential for correlated noise to appear in the responses of adjacent neurons, which would seriously interfere with efficient information processing. The need to avoid the consequences of correlated noise and redundancy of information provides strong pressure favoring selection of connection weights that keep the responses of adjacent neurons nearly independent both for signal and noise.

Some of these data have appeared in an abstract (Gawne and Richmond, 1990).

Materials and Methods

Under surgical anesthesia, using standard sterile surgical techniques, two rhesus monkeys (*Macaca mulatta*) were prepared for chronic single-unit recording, with a stainless-steel cylinder implanted on the top of the head in the stereotaxic plane (AP +15, L 21) to allow penetration to the IT cortex (Richmond et al., 1983). A coil of Teflon-coated stainless-steel wire was implanted under Tenon's capsule of one eye, which allowed eye position to be continuously and accurately monitored via




Figure 1. Stimulus set: the low-resolution 4×4 set of Walsh patterns used as visual stimuli. Both this and the opposite contrasts were used, for a total of 32 patterns. Each stimulus was 2.3° square, and all stimuli were centered at the point of fixation.

the magnetic field/search coil technique (Robinson, 1963; Judge et al., 1980). These monkeys were trained to perform a variation of a standard fixation task, in which the fixation target was turned off when the stimulus was presented, thus avoiding any interference between the stimulus and the fixation point (Wurtz, 1969; Richmond et al., 1983).

Experimental paradigm. The stimulus set consisted of 16 Walsh patterns and the 16 contrast-reversed ones for a total of 32 visual stimuli (Fig. 1) (Ahmed and Rao, 1975; Richmond et al., 1987). These stimuli were used in a nonmatch-to-sample task (Fig. 2) (Eskandar et al., 1992). The monkeys initiated trials by grasping a metal bar and fixating a small black spot subtending 0.15° . After 350 msec, the fixation point disappeared and a pattern, the sample, subtending 2.5° , appeared in the center of the visual field for 352 msec. Following the sample pattern there was a delay period of 550 msec, during which the fixation point reappeared. After the delay period, the fixation point was turned off again, and a test pattern appeared for 352 msec. If the test stimulus was different from the sample stimulus, that is, a nonmatch, the monkey was required to release the bar within 550 msec after its disappearance to obtain a liquid reward. If the test stimulus was the same as the sample stimulus, that is, a match, the monkey was required to continue holding the bar for 550 msec until a third nonmatching stimulus appeared. In this case the monkey was required to release the bar during the presentation of this third stimulus to obtain its reward. The monkey was required to maintain fixation within 1.0° of the fixation point throughout the duration of a trial; otherwise the trial was aborted. At the completion of a trial there was an interval of at least 750 msec before the next sample stimulus appeared. The type of trial, match or nonmatch, was chosen randomly, as was the pairing of nonmatch with sample stimuli. This ensured that the monkey could not anticipate the type of trial and that there were nearly equal numbers of trials for each stimulus in each context. The monkeys were trained until they performed correctly at least 85% of the time.

The stimuli are specified by two factors, visual pattern and behavioral condition. The visual patterns were the 32 Walsh patterns, and the behavioral conditions consisted of sample, match, and nonmatch. Most of the time we shall regard each combination of pattern and condition as constituting a unique stimulus, but at times we shall examine one factor regardless of the value of the other factor, for example, the effect of Walsh pattern 1 regardless of the condition, or the effect of the match regardless of the visual pattern. This will always be explicitly noted.

Multiunit separation. To separate units in the multiunit recordings, we exploited the fact that spikes from different cells generally have different shapes. To do this we used a modification of a technique first developed by Abeles and Goldstein (1977). In this technique, a set of

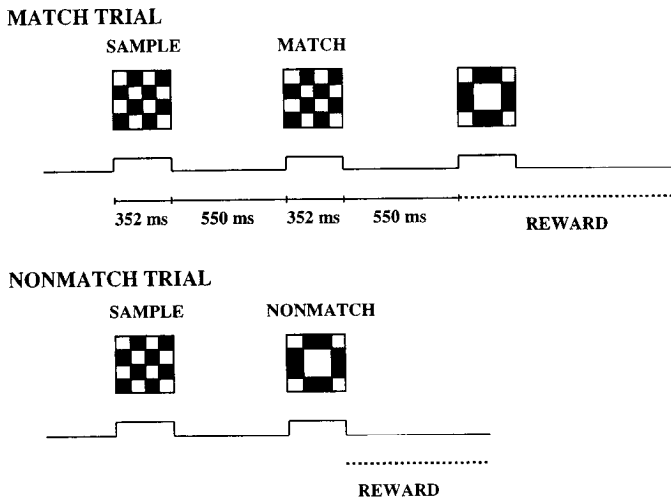


Figure 2. Schematic diagram of the nonmatch-to-sample task. In this task, a sample picture is followed by a test picture. If the second picture does not match the first, the monkey must release a bar to be rewarded. If the second picture matches the first, then the monkey must wait for the appearance of a third nonmatching picture and then release the bar in order to be rewarded.

waveforms known as the principal components are first calculated from the waveforms of a set of spikes recorded from all of the neurons. The principal components form a basis for the decomposition of the spike waveforms, giving a set of coefficients for each waveform (Ahmed and Rao, 1975). This is analogous to the Fourier transform, where a signal is broken down into a series of sine waves, and can then be completely characterized by the coefficients for the different sine waves. However, unlike Fourier analysis, the shape of the principal components is not fixed in advance, but is a function of the data themselves. The coefficients of the first principal component account for the greatest variance possible with a linear measure. The coefficients of each succeeding principal component account for the maximum variance possible for the residual, with the restriction that it is orthogonal to the others. The coefficients of one principal component are uncorrelated with the coefficients of any other principal component. Thus, the principal components provide a very efficient representation of the data. For our needs in spike separation, the critical property is that the coefficients of the first two principal components of the spikes account for the largest amount of the variance possible with two waveforms. Two principal components were used because a two-dimensional plot can be easily and rapidly displayed on a computer screen.

In implementing the Abeles and Goldstein technique, the signal from the microelectrode was filtered to pass frequencies only between 500 Hz and 10 kHz, and was then sampled by a digital signal processing system (Communications, Automation, and Control DSP32C with DBCS 5339 Analog to Digital converter, Allentown, PA). This system connects to an IBM-PC compatible computer, which handled the display and user interface. The sampling rate of the A/D converter was set at 20,000 samples per second. Spikes were detected by a simple voltage threshold, and 20 data points were used for each individual spike. When stable neurons appeared on the oscilloscope display, the data from 100 spikes were collected, and then the principal components were calculated from these data. The first two principal components were loaded into the digital signal processing system. Thereafter, whenever a spike was detected it was plotted as a single point with the coefficient of the first principal component on the vertical axis and the coefficient of the second principal component on the horizontal axis.

The points from different neurons fell into different clusters, and clusters could be manually selected by drawing a box around the points. By associating different clusters with different neuronal spikes, the signals from different neurons were identified. The position of the electrode was adjusted to optimize the separation of multiple units, and only well-isolated clusters were used. Overlapping spikes cannot be detected with this technique. However, the system reset immediately after the amplitude on a spike went below threshold, so the time during which the

system did not respond (the dead time) was approximately equal to the time that the spike waveform was above threshold, less than 1 msec. The time of spike occurrence was recorded to the nearest millisecond by our laboratory computer.

Response quantification. To quantify a neuron's response, we first convolved the spike train with a Gaussian kernel, $\sigma = 20$ msec. The resulting spike density function provides a low-pass-filtered estimate of the probability of spike occurrence over time, 3 dB cutoff 7 Hz (Ahmed and Rao, 1975; Silverman, 1986; Richmond et al., 1987). Changing the σ changes the trade-off between resolution in time and resolution in amplitude. A larger σ loses the ability to resolve rapidly changing events but, with averaging over a larger time, is better able to resolve the absolute levels of a signal. We tried several different values of σ ranging from 5 to 20 msec on data from each of several neurons, and found that over this range there were no significant differences among any of our analytic measures.

Our analysis was restricted to an interval of 256 msec duration starting 90 msec after the stimulus onset. The end of this interval coincided with the stimulus being turned off. The spike densities were sampled every 4 msec over a 256-msec-long window, giving a 64-dimensional vector representation of the response. These vectors were decomposed into the principal components, which incorporate both the magnitude and temporal pattern of the responses (Richmond and Optican, 1987). The principal components were calculated from the spike densities of all of the individual responses to all of the stimuli. The principal components were calculated for each neuron independently. We will refer to the first principal component as ϕ_0 , the second as ϕ_1 , and so on. Details of the principal component method are given in the previous section.

We represented the data in this experiment with the coefficients of the first five principal components, which together accounted for $93.3 \pm 0.2\%$ of the variance of the data from all of our cells. We restricted ourselves to the first few principal components because they account for most of the variance of the response, and because it is currently not possible to perform our analyses on more than a few principal components. In addition, previous work has shown that only the first few principal components carry significant amounts of information in the response of a single cell in IT cortex (Optican and Richmond, 1987). As seen before, the coefficient of the first principal component was always highly correlated with the number of spikes in the response (Richmond and Optican, 1987; Richmond et al., 1990; McClurkin et al., 1991a).

Throughout this paper we will be using linear regression, both univariate (i.e., involving one independent variable) and multivariate (i.e., involving several independent variables). We will report the multiple R^2 values from the regression because R^2 shows how much of the variance in the data is accounted for by the regression model. For the univariate case, this is equivalent to the square of the correlation coefficient (Draper and Smith, 1980). Although R^2 does not indicate whether the correlation is positive or negative, for the purposes of this study the sign of the correlation was not a central issue.

The mean R^2 for the correlation of the coefficients of the first principal component versus the spike count for all 56 cells was 0.88 ± 0.013 . The multiple R^2 gives the variance accounted for by the linear regression; therefore, 88% of the variance of the spike count can be determined by knowing the value of the coefficient of the first principal component. The principal components are orthogonal, and the coefficients of any one principal component are uncorrelated with the coefficients of any other. The first five principal components account for the magnitude and most of the temporal pattern of a neuron's responses.

If the data had been analyzed in the time domain, the responses would have had one dimension for each millisecond of the response, in this case 256 msec. An impractically large number of trials would have been required to analyze data of such high dimensionality directly. Reducing the dimensionality of the representation reduces the number of trials required to achieve a given level of power to a practical level. We used the principal components because of their orthogonality and efficiency in accounting for the variance in the data.

ANOVA. To determine which aspects of the neuronal response were affected by the two factors, pattern and condition, a two-way analysis of variance (ANOVA) was applied to the coefficients of the first three principal components, each taken alone. We accepted $p < 0.05$ as significant. For the pattern factor, there were 32 patterns in the stimulus set, consisting of the 16 Walsh patterns and their contrast-reversed mates. For the condition factor, there were three behavioral conditions:

sample, match, and nonmatch. The Brown–Forsythe correction was used to compensate for unequal variances.

Separation of signal and noise in neuronal responses. To estimate the signal in the response of a neuron to a particular stimulus condition, we used the average of the coefficients of the principal components of all of the responses in that stimulus condition. To estimate the noise, we subtracted the estimated signal for the class from the response to each individual stimulus presentation for each principal component. The noise was not partitioned into the different stimulus conditions, but was analyzed as a single set across all stimulus conditions. Thus, for each stimulus class the signal was represented by the coefficients of several average principal components, and for the noise we had a set of principal component coefficients for each individual response. Since the principal components are a linear transform, this separation of signal and noise could have been done in the time domain. However, as mentioned previously, an impractically large number of trials would have been required to perform this analysis without first reducing the dimensionality of the signals to a manageable level. The use of principal components achieves this with maximum efficiency (Ahmed and Rao, 1975).

Correlations of signal and noise across neurons. One straightforward method to determine whether two neurons have the same response characteristics is to perform a linear regression of the response of one versus the response of the other. Although this is a useful measure, the results must be interpreted carefully. Consider two neurons that have identical response characteristics, but have nearly uncorrelated noise. Depending on how much noise contaminates the signal, linear regression could show that the responses of the neurons are nearly uncorrelated. However, if one took the average response for each stimulus condition across many repetitions, and performed the correlation on the averages, the correlation would be extremely high. Using the averaged principal component coefficients as estimates of the neuronal signals allows us to determine if two neurons have the same response characteristics even in the presence of noise.

We calculated the linear regressions for both the average responses and the noise responses for the first principal component, that is, ϕ_0 from cell A correlated with ϕ_0 from cell B. Because the higher principal components are typically noisier, the correlation of each principal component from one neuron with each one from the other neuron may be unreliable for other than the first one. In a simple case, the first principal component from one neuron might have corresponded with the second principal component from an adjacent neuron. By using canonical correlation, we were able to examine the relation of the first principal component from one neuron against all of the components from the second neuron at once (Morrison, 1967). In this way, we could identify whether the amount of variance accounted for was higher when several principal components were considered together.

The accuracy of an estimate is a function of the number of samples. To study the effects of sample size on our results, we performed two controls. First we took the data from a single neuron, resampled the data with replacement to create two new data sets, and calculated the correlation between these two data sets. Any deviation from a correlation of 1.0 is due to the variability in the data. Second, we took the data from each pair of neurons and shuffled the stimulus conditions randomly relative to the responses, so that any relationship between the average responses of the two neurons to any stimulus would have been broken. Any apparent correlation in this shuffled data must have arisen from the variations that are not related to the stimuli. Any correlations above this level can be considered significant. Note that our estimate of the noise correlation is much less sensitive to the sample size because it is computed directly from the differences from the average across all of the trials, typically over 1000.

Information calculations. We used information theory to search for relations in our data that cannot be found using linear regression techniques. For example, the relationship between adjacent neurons could be a function of the two responses taken together. One possible code of this type would occur if the signal from a single neuron was quite variable, but the difference between the responses of two adjacent neurons was very precisely determined by the stimulus. This type of information can be detected only by examining the responses of the neuronal pairs together, and we will refer to this situation as a cooperative, or synergistic, code. Regressing the response of one neuron against the responses of the other will not identify such a relation. However, information theory can be used to detect such relationships.

The probability of correctly identifying the identity of a stimulus by

knowing the neuron's response can be quantified by information theory. Transmitted information can be calculated from the probabilities of occurrence of stimulus–response pairs using

$$T(S; R) = \sum_j p(s_j) \sum_k p(r_k | s_j) \log \frac{p(r_k | s_j)}{p(r_k)}$$

where $T(S; R)$ is the average amount of information transmitted about the stimuli by the responses, $p(s_j)$ is the probability of occurrence of stimulus s_j , $p(r_k)$ is the probability of occurrence of response r_k , and $p(r_k | s_j)$ is the probability of occurrence of response r_k given that stimulus s_j was presented (Shannon, 1948; Abramson, 1963; Optican and Richmond, 1987; Optican et al., 1991). If the probability of a correct guess is higher than it would be had a stimulus been selected at random, then the neuron has carried information about the stimulus. Since information theory depends only on probabilities of occurrences, this part of the calculation is model-free. Using information theory, any non-random relationship, linear or nonlinear, can be quantified regardless of its underlying mechanism.

To study how the information carried in the responses of adjacent neurons is related, we first calculated the information for each neuron separately, and then for the pair taken together, that is, jointly. Because information theory takes the noise into account, this calculation was carried out with the coefficients of the principal components from the individual responses. If neurons are independent, then the joint information should be equal to the sum of the information calculated separately. If the neurons carry information that is a function of both responses taken simultaneously, that is, synergistically, then the joint information should be greater than the information calculated by adding the information for each neuron taken separately.


At present, we cannot calculate the amount of transmitted information carried by more than five principal components taken simultaneously. Therefore, for the joint calculation, we used only two principal components from each neuron, for a total of four components.

Cross-correlograms. To compare our results with those obtained through cross-correlation techniques, we calculated the cross-correlogram between each pair of adjacent neurons. For every spike that occurred for one neuron, we added to a histogram the relative times of occurrence, or lag, of all the spikes that occurred for the other neuron, ± 50 msec. Because our technique of multiunit analysis could not detect simultaneous action potentials, we did not calculate the cross-correlogram for zero lag (± 0.5 msec). We calculated a control cross-correlogram by randomly shuffling the neuronal responses for each stimulus before performing the computations. The actual and control cross-correlograms diverged only from the shuffled at delays centered around zero. We chose an interval (always ± 50 msec) that was just greater than the widest interval for which the two histograms diverged. We applied the Wilcoxon signed rank test to these data over the same interval for all 28 data sets.

Results

The neurons in this study were recorded from the same two monkeys, and in the same areas of IT cortex, as in a previous study (Eskandar et al., 1992). We searched until we isolated one unit that responded clearly to the visual stimuli. We required that our multiunit analyzer differentiate at least one other well-isolated unit. We did not require that the second unit be clearly driven by the stimuli. We bypassed single neurons. Although we did not deliberately search for nonvisual units, we encountered 23 locations with differentiable units that were bypassed because none was clearly visually responsive to our stimulus set. In addition, we isolated four pairs of units that were determined to be nonvisual by inspection as the experiment was underway, at which point the recording was terminated and the search for new units continued. Thus, it seems that neurons that responded to our stimulus set were located in patches. We gathered data from 28 pairs of neurons, 14 from each monkey. For each distinct stimulus condition, there were an average of 16.2 ± 2.2 trials (range, 5–58) for all of the neurons in this study. Either the stimulus pattern, behavioral context, or both elicited

Figure 3. Results from the two-way ANOVA on the two factors of pattern (*P*) and behavioral condition (*C*), and the interaction between the two (*I*). The ANOVA was performed on the first three principal components separately (ϕ_0 , ϕ_1 , and ϕ_2), matching a previous study (Eskandar et al., 1992). The numbers are the percentages of the cells that showed significance at the $p < 0.05$ level. The *open bars* are for all the cells in the study ($n = 56$). The *shaded bars* are the results obtained using only the neuron in each pair that was most strongly modulated ($n = 28$).



significant response modulation for 55 of 56 (98.2%) of the neurons in this study, for at least one of the first three principal components of the response (ANOVA, Fig. 3).

In 89% (50 of 56) of the neurons, the coefficients of the second or third principal component showed significant stimulus dependence (ANOVA). This result is in accord with an analysis done on the first three principal components in a similar study in IT cortex that used the identical paradigm (Eskandar et al., 1992). Thus, both the strength and the temporal modulation of the stimulus-dependent responses of the great majority of these neurons were affected by both stimulus pattern and behavioral condition.

When the data from a pair of neurons were examined, the two neurons sometimes had very similar changes in firing rate as the stimulus pattern was changed (Fig. 4A). Had only these two patterns been used, it might have been concluded that the neurons were performing essentially the same function. However, the same two neurons had very different changes in firing

rate when two other patterns were presented (Fig. 4B), suggesting that these two neurons had different functions. Examples where two neurons appeared to be responding to the stimuli either similarly or differently can be easily culled from the whole data set (Fig. 5). To determine the extent to which the responses of adjacent neurons were actually independent, we examined the correlations between neurons with linear regression, and their degree of independence with information theory.

Correlations for signal and noise

In general, the response strength of one neuron showed very little correlation with the response strength of its simultaneously recorded neighbor. The correlation, that is, multiple R^2 , of the coefficients of the first principal component of one cell versus the coefficients of the first principal component of the other cell for all patterns and all conditions was 0.057 ± 0.017 . The multiple R^2 is a direct measure of the amount of variance in one cell's responses that can be accounted for by the second

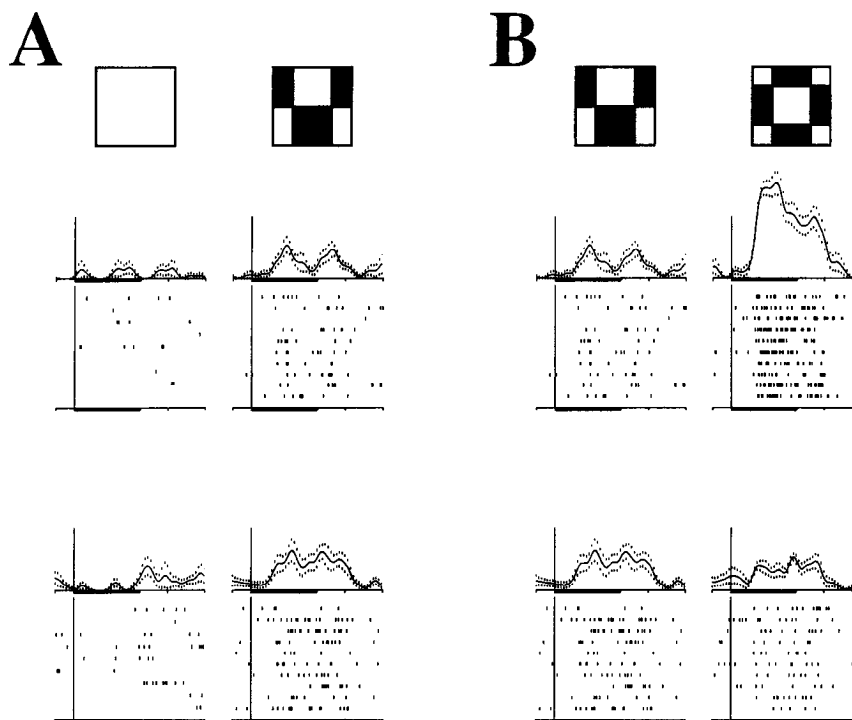


Figure 4. Similarities and differences in responses from two adjacent neurons. The continuous *waveforms* above the rasters are the spike densities with standard errors, an estimate of the probability that there will be a spike during that millisecond. The *dark bars* under each response denote the time that the stimuli were on. These data were taken from stimuli presented in the sample condition. The *vertical bar* indicates stimulus onset and, in the spike density diagrams, its height corresponds to a probability of 0.1 that a spike will occur in a 1 msec interval, or an instantaneous firing rate of 100 spikes/sec. *A* shows an example where the two neurons responded in a nearly identical manner to two different patterns, and *B* shows an example from the same pair of neurons where they responded differently.

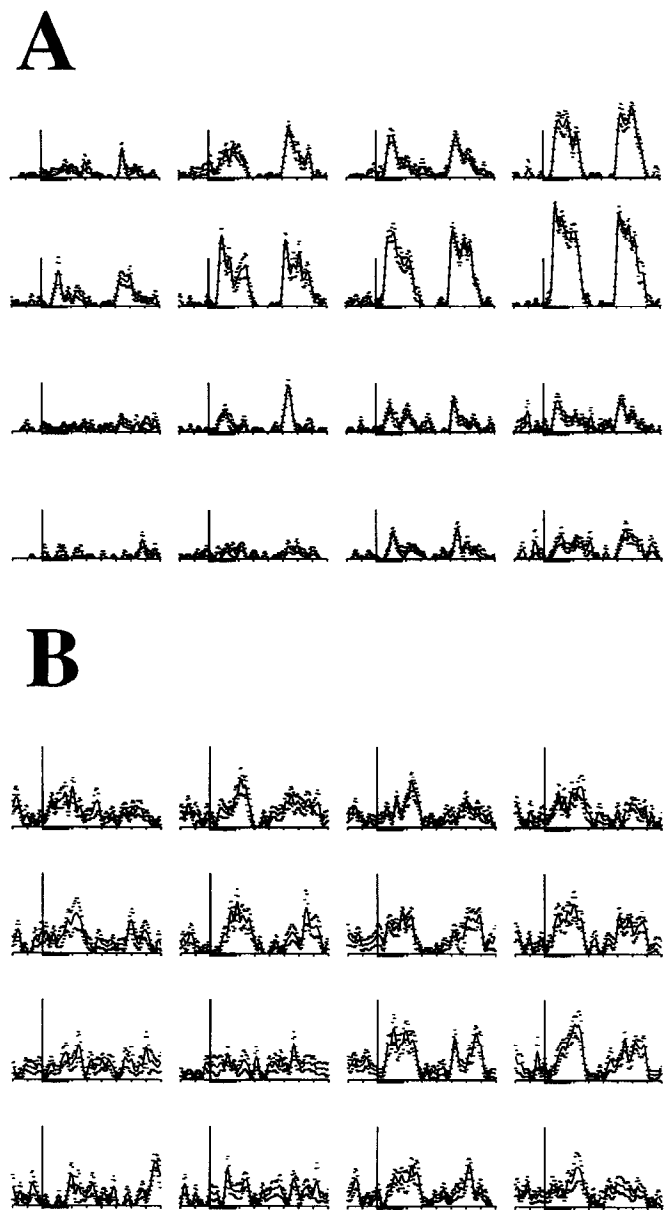
cell's responses; thus, on average only 5.7% of the variance of one cell's response strength can be accounted for by examining the response strength of the adjacent cell. As described in Materials and Methods, we separated the responses into a signal (the average value of the responses for all repetitions of a particular stimulus) and noise (the difference on a trial-by-trial basis between the response of a neuron and the average value for each particular pattern and stimulus condition). The ratio of the variance of the signal to the variance of the noise for the response strength as represented by the first principal component was 0.21 ± 0.02 (\pm SE).

For the pair of cells illustrated in Figures 4 and 5, linear regression for the first principal component of the signal from one neuron versus the first principal component of the signal from the other showed that 25% of the signal from one neuron could be accounted for by knowing the signal from the other; that is, the signals show little correlation. The multiple R^2 between the average values of the coefficients of the first principal component of the signal from one neuron versus the coefficients of the first principal component of the signal from the adjacent neuron for all 28 pairs is 0.187 ± 0.057 (range, 0.00–0.45; Fig. 6).

To better interpret these results, we estimated the effect of the number of trials on their reliability (see Materials and Methods). To control for the finite number of trials per stimulus, we resampled the data with replacement twice for each neuron to produce two data sets that should have had a multiple R^2 of 1.0. However, the average multiple R^2 for this control was 0.65 ± 0.08 , which, although less than 1.0, is substantially greater than the value of 0.187 obtained from the actual data. The control R^2 was greater than the actual R^2 for every pair of neurons in this study (nonparametric sign test, $p < 0.0001$). Thus, the general lack of correlation between the responses of adjacent neurons in our results is not a consequence of estimating the average response to a stimulus from a small number of trials. For another control, we shuffled the responses for both members of the pair of neurons, and then correlated them. The multiple R^2 averaged 0.012 ± 0.003 , which is substantially below the unshuffled correlation above ($p < 0.001$). The multiple R^2 is therefore not a consequence of having studied neurons with completely uncorrelated responses.

The multiple R^2 for the coefficients of the first five principal components from one cell versus the coefficients of the first principal component for the other cell in each pair was 0.22 ± 0.03 ; that is, the coefficients of the first five principal components of one neuron can on average predict at most 22% of the variance of the signal represented by the coefficients of the first principal component of an adjacent neuron. Therefore, there is no linear mapping between different aspects of the response of one cell that can be used to predict the response of one of the components in the adjacent cell.

The noise correlations were uniformly small; only $5.5 \pm 0.8\%$ (range, 0.0–22.0) of the noise for the coefficients of the first principal component from one neuron could be predicted from the noise on the coefficients of the first principal component of the adjacent neuron (Fig. 6). Similarly, using the coefficients of all five principal components from one cell to predict the coefficient of the first principal component from the other, only $6.3 \pm 1.0\%$ of the noise variance on one cell could be accounted for by the noise on the adjacent cell. Thus, for either the response magnitude or the temporal modulation, the noise in the responses of adjacent neurons shows little correlation.



walsh353

Figure 5. Illustration of the lack of correlation of responses in two neurons. Here we show the responses to the 16 positive contrast stimuli for the sample case. Both cells responded strongly to the visual stimuli, and had similar baseline firing rates. The first cell (*A*) tended to fire in a graded manner, with only a few stimuli eliciting the maximum firing rate. However, while the second cell (*B*) was also strongly affected by this stimulus set, it responded with the same level to many stimuli. Linear regression showed that only 25% of the variance of the average response to each stimulus could be predicted by using the average response of the second neuron. Thus, while not completely uncorrelated, the response properties of these two neurons are clearly different. The R^2 for the noise, which is simply the difference from the average on a trial-by-trial basis, was 0.045, almost completely uncorrelated. Fiducial marks are the same as in Figure 4.

The number of trials used in calculating the noise correlation was 2069 ± 285 (range, 633–7509). Thus, there is almost no uncertainty in these estimates. Correlating the outputs of two simulated neurons that have uncorrelated noise, using the same

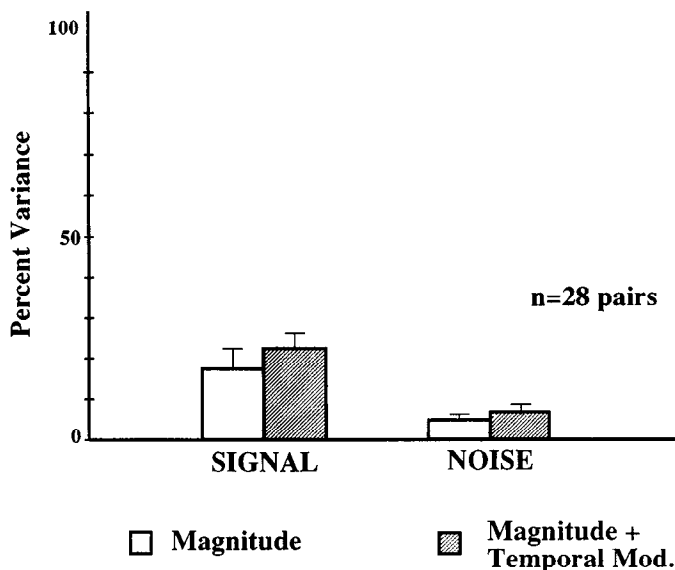


Figure 6. Variance in the response strength of one neuron accounted for by the response of the other. The open bars show the amount of variance accounted for in the magnitude of the response of one neuron using only the magnitude of the response of the other. The shaded bars show the amount of variance accounted for in the magnitude of the response of one neuron using the first five principal components from the other cell (see Materials and Methods). The analysis was done for both the signal and the noise separately.

number of trials and amount of variance as for our actual data, results in a multiple R^2 of 0.0022, that is, virtually zero.

We calculated the correlation between the coefficients of the first principal components from adjacent neurons for each behavioral condition separately (Fig. 7). There were no significant differences between the correlations for any pair of conditions. The greatest differences were between the match and nonmatch

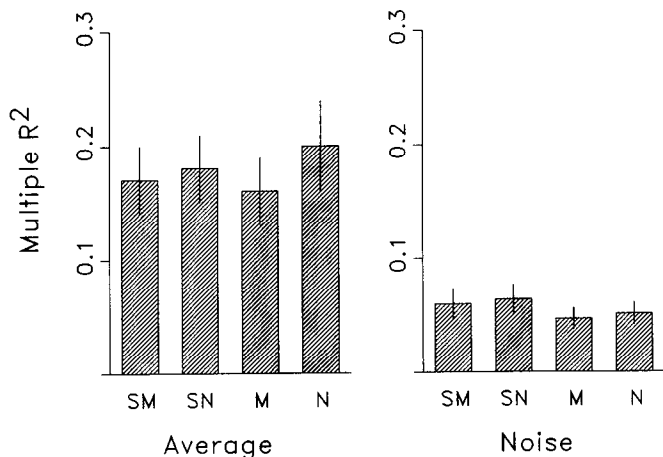


Figure 7. Graph of the correlations between the average and the noise under the different behavioral conditions. SM stands for the sample before the match condition, and SN for the sample before the nonmatch condition. The sample condition was partitioned in this way in order to have equal numbers for comparing the different conditions, and so that the error bars between the different conditions may be compared directly. M is the matching condition, and N the nonmatching condition. While there was a tendency for the correlation of the signal to be higher in the nonmatch than the matching condition, and for the correlation of the noise to be lower in the match than the sample condition, these differences were not statistically significant.

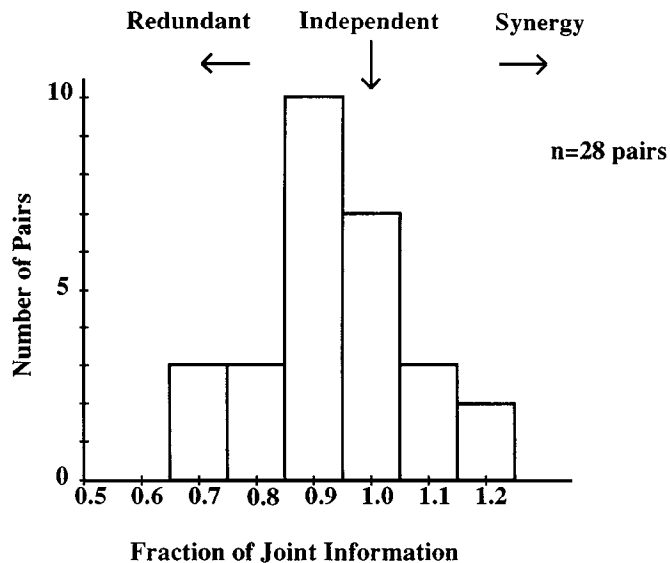


Figure 8. The fraction of the joint transmitted information that is accounted for by the sum of the information transmitted by each neuron separately. The vertical arrow indicates the point of exact independence. Synergistic codes would cause the data to fall to the right of the arrow, and redundant codes to the left. The 0.5 level is the point at which two neurons would be if they were redundant and had correlated noise. (Data below the 0.5 level are precluded, because the joint information cannot be less than the information for any single neuron.) The results show that the information carried by adjacent neurons is closer to being independent than redundant, and do not indicate either very strong or widespread occurrence of synergistic codes.

conditions for the signal (t test, $p = 0.21$), and between the sample before the nonmatch and the match conditions for the noise (t test, $p = 0.13$). Thus, the behavioral condition did not significantly affect the degree of correlation between either the signal or the noise of adjacent neurons. It was not possible to determine the correlation separately across the three different behavioral conditions, because three data points are insufficient to perform a meaningful linear regression.

Information

To compare our results with those from previous studies (Eskandar et al., 1992), we calculated the information transmitted by each neuron individually using the coefficients of the first three principal components. The mean information transmitted about a stimulus for all 56 cells was 0.23 ± 0.04 (\pm SE) bits. Using only the coefficient of the first principal component, the mean transmitted information was 0.14 ± 0.04 bits, which is only 61% of the preceding value. Because our criterion for acquiring data was that only one cell of a pair need be visually responsive, we averaged the transmitted information using only data from the one cell of each pair with the highest transmitted information. This yielded values of 0.29 ± 0.05 bits using the coefficients of the first three principal components, and 0.20 ± 0.04 bits for just the coefficient of the first principal component, which is 67% of the value for three. These results are similar to those from a previous study performed on single neurons in IT cortex (Eskandar et al., 1992).


We calculated the information transmitted by each cell separately about the stimulus set, and by both cells of the pair taken together. To check for joint, or synergistic, codes, we calculated the percentage of the joint transmitted information accounted

Redundant Example

T_{A+B} : 1.01 Bits

T_A : 0.72 Bits

T_B : 0.67 Bits



for by the sum of the transmitted information calculated for each neuron separately (Fig. 8). If the joint information consistently exceeded the sum of the information for each pair separately, this would be evidence for a synergistic code. Of the 28 pairs, the joint information exceeded the sum of the information from each neuron separately by a maximum of 20.1%. The average of the joint information was 92.0% for all of the pairs. Thus, we did not find evidence for either very strong or widespread synergistic codes.

Simulated neuronal pairs

We tested the abilities of linear regression and information theory to detect differences in coding using simulations in which

we controlled the response characteristics. Figures 9–11 show three simulated pairs of neurons that use three different strategies for encoding information about a stimulus: redundant with independent noise, independent, and synergistic. There are 16 different stimulus conditions, and each stimulus is presented for 20 trials. Each simulated neuron responds with a burst of white noise, and the average response to each stimulus is graphed. In these examples, only the mean firing rate was modulated by the stimulus. The transmitted information was calculated for the response magnitude of each cell separately, T_A and T_B , and also for the responses of both cells jointly, T_{A+B} , and then compared.

Because our spike generation process was always the same, the cross-correlations between neuronal pairs are identical for

Independent Example

T_{A+B} : 1.51 Bits

T_A : 0.72 Bits

T_B : 0.79 Bits




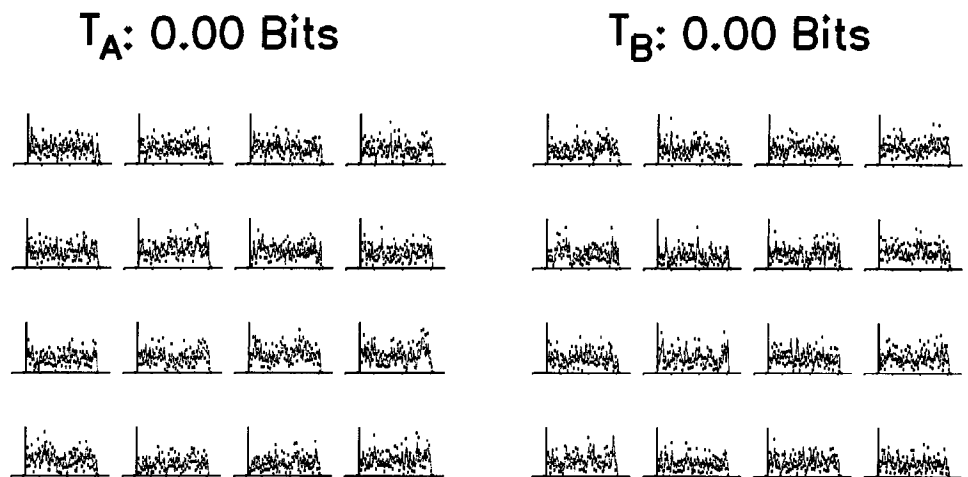
Figure 9. Simulated neuronal pairs with redundant information. The two panels show the results from two simulated neurons. The simulation used a number of trials and a level of noise similar to those in our actual data, and all of the same analyses were performed. There were only 16 stimuli in this simulation, and 20 trials for each stimulus. Each simulated neuron responded to a stimulus with a burst of white noise, whose strength was proportional to the number of the stimulus. The response characteristics of the two neurons were identical, although the noise was independent. This resulted in a multiple R^2 between the average responses of 0.87, and 0.043 for the noise. The information calculated for neuron *A* (left) alone was 0.72 bits, and for neuron *B* (right) alone was 0.67 bits. The information calculated from a joint code, that used the responses from neurons *A* and *B* simultaneously, was 1.01 bits.

Figure 10. Simulated neuronal pairs with independent information. This figure shows the results of a simulation that is nearly identical to that in Figure 9, except that the response characteristics of the two neurons are independent. This resulted in a multiple R^2 of only 0.0127 for the average of the responses, and 0.0007 for the noise. The actual joint information is 1.51 bits, and the prediction assuming independence is also 1.51 bits.

Synergistic Example

 T_{A+B} : 0.64 Bits

Figure 11. Simulated neuronal pairs with synergistic information. In this simulation, the response strength of the first neuron varies randomly irrespective of the stimulus. However, the response of the second neuron will be the same as the response of the first neuron for some stimuli, and either larger or smaller for other stimuli. Thus, the response of one neuron in isolation can yield no information about what the stimulus was, but the responses of both neurons can yield considerable information. The information is zero from either alone, but it is 0.64 bits when calculated from the joint responses of neurons *A* and *B*. In this case the response of one neuron cannot be interpreted without using the responses of its neighbor.



all three cases. For the example here, the use of white noise as the response meant that the cross-correlograms were always perfectly flat, independent of the functional relationship between the two simulated neurons. At the other extreme, we could have constrained the spike generation process to produce only a spike on even 5 msec boundaries, but left everything else unchanged. In this case the correlogram would show strong peaks at 0 msec, ± 5 msec, ± 10 msec, and so on, but the functional relationships between the mean firing rates of the two neurons would have been completely unaffected. Thus, while cross-correlation is quite sensitive to the details of the mutual spike generation process, it is of limited utility in determining how the messages carried by pairs of neurons are related, at least for the time scales that we examined in this study.

In the simulation in Figure 9, the response strength was a function of the stimulus, and both simulated neurons carried the same message. The noise was completely uncorrelated between the two neurons. The multiple R^2 between the average responses was 0.87, and the multiple R^2 between the noise was 0.043. The multiple R^2 was not 1.0, because the finite sample size prevents us from knowing the true average response for a particular stimulus. In this example, the correlation between the average response for one neuron and the average response calculated from data from the same neuron randomly resampled is 0.9 for the first cell versus itself (multiple R^2), indicating that the lack of perfect correlation was indeed due to the effect of noise on the means calculated from a finite sample size.

In the simulation in Figure 10, the response strength was a function of the stimulus, but the two simulated neurons carried independent messages. The multiple R^2 between the average responses was 0.0127, and the multiple R^2 between the noise was 0.0007. Therefore the amount of information transmitted should be the sum of the information transmitted by each neuron separately for a finite noisy sample $T_A + T_B \approx T_{A+B}$.

It is also possible that adjacent neurons cooperate in transmitting information about a stimulus. In the simulation in Fig-

ure 11, the response strength of the first neuron varies randomly irrespective of the stimulus. However, the response of the first neuron will be the same as the response of the second neuron for some stimuli, and either larger or smaller for other stimuli. Thus, the average responses of both neurons are identical and do not depend upon the stimuli. Although the response of one neuron in isolation can yield no information about what the stimulus was, the responses of both neurons considered jointly can nevertheless yield considerable information, which information theory can detect. Thus, $T_{A+B} \gg T_A + T_B$.

Cross-correlograms

Seventeen of the 28 pairs of neurons had response cross-correlograms that were significantly different from the control ($p < 0.05$, Wilcoxon signed rank test). Figure 12 shows two cross-correlograms, one significant and one not. Due to the limitations of our multiunit analysis technique, the cross-correlation for zero lag (the interval within 0.5 msec of zero) was not calculated. The multiple R^2 for the cells with significant cross-correlograms was 0.18 ± 0.14 , and 0.11 ± 0.14 for the cells that did not have significant cross-correlograms. This trend was not significant (t test, $p = 0.105$). In the examples shown in Figure 12, the values of R^2 for the coefficients of the first principal components for both pairs of neurons were similarly low, 0.15 for the nonsignificant one and 0.22 for the significant one, showing that the cross-correlogram and the correlation of the average values of the first principal component measure different aspects of the data and need not be related.

Discussion

As a step toward understanding the relationships among neurons within local groups, we recorded from adjacent pairs of neurons (i.e., recorded from the same microelectrode) in IT cortex. Because the neurons were close, we expected that they might share many common inputs, and the cross-correlogram data support

this expectation. However, this does not necessarily imply that they carry the same information.

We report two major findings. First, the signals carried by two adjacent neurons in IT cortex are nearly uncorrelated and almost independent. Second, the noise carried by adjacent neurons is even more uncorrelated than the signal. However, even the small amount of redundancy we find severely constrains the amount of information that can be carried by a pool of neurons. These results lead us to a new concept about how small, local groups of cortical neurons are organized; namely, their organization is constrained by the need to avoid correlated signals and noise.

Independence and redundancy

At any given instant, the amount of information about a visual stimulus is limited to that carried by the optic tracts, on the order of 10^6 fibers from each eye (Bruesch and Arey, 1942). This is because changing the representation of a signal cannot increase the absolute information content, although the new representation may be more convenient for further processing. Previous work indicates that the information-carrying capacity of fibers in the optic tract is not radically greater than for single visually responsive neurons found in several cortical areas (McClurkin et al., 1991b). Therefore, the large numbers of neurons in cortex dealing with the visual process carry in aggregate no more information than the relatively few fibers in the optic tract.

This constraint on the information capacity of visual cortical areas is extremely broad, and says nothing about how redundancy might be manifested. The simplest situation would be to have large populations of neurons carrying identical messages, with the signals from perhaps thousands of neurons averaged to produce a single highly precise message. However, this strategy has problems. For one thing, it makes extremely inefficient use of neurons. As the number of neurons in a population increases, the signal-to-noise ratio rises as the square root of the number of neurons. Furthermore, this presupposes that the noise on different neurons is independent. If neurons had the same or nearly the same response properties and had many of the same inputs, then it would be impossible for them to have uncorrelated signals and noise. When averaging across multiple noisy elements, there will be no improvement to the extent that the noise is identical.

We suggest that it might be so inefficient to have many neurons doing essentially the same processing that it is precluded by design. Because there are many more neurons in visual cortical areas than there are fibers in the optic tract, at some level the representation of a visual stimulus in cortex must be massively redundant. This redundancy could be satisfied by small overlaps in the response properties of an extremely heterogeneous population of neurons, such as we have seen here.

The amount of redundant information that we found seems at first glance to be low. However, finding that the information carried by adjacent neurons is not completely independent places an upper bound on the amount of information that can be represented by a pool of neurons. This upper bound is simply $1/\gamma$ times the amount of information transmitted by a single neuron, where γ is the fraction of the amount of information that is common between any two neurons (Fig. 13). If cortical neurons were constrained such that all pairs were redundant to the degree found here, it would be impossible to represent more than a few bits of information no matter how many neurons were used. Because we can recognize a large number of distinct

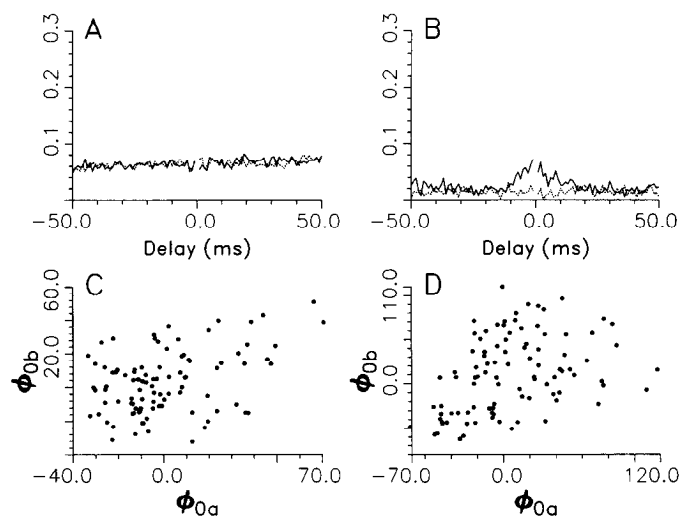


Figure 12. Cross-correlograms. In *A* and *B*, the dark line represents the cross-correlation between the spikes from two neurons, that is, the probability that, given a spike on one neuron, there would be a spike on the other neuron at that delay. The dashed line represents a control where the order of the trials for each stimulus class was shuffled before calculating the cross-correlation. The Wilcoxon signed rank test was used to judge the significance of the difference between these two lines. The cross-correlogram in *B* is significant ($p < 0.05$), and the one in *A* is not. *C* and *D* show the scatterplots of the average for each stimulus condition (signal) for the first principal components for the corresponding pair of neurons in *A* and *B*. Linear regression of the first principal component between the neurons of each pair gave a multiple R^2 of 0.15 in *C* and 0.22 in *D*. Thus, these two pairs of neurons had essentially the same functional relationship at the level of the first principal component, that is, the response strength, even though the cross-correlation was quite different.

visual patterns, it is clear that more than a few bits of information must be represented in the visual system. Because of the devastating impact of even a small amount of redundancy, we propose that the amount of redundancy must decrease as the distance between the neurons increases. Our analysis of the effect of even small amounts of mutually redundant information shows that, for preserving information, it is important that neurons be as independent as possible.

Noise

To the extent that two neurons are performing weighted sums of common afferents, the noise carried in their responses should tend to be correlated to the same degree that the signal is. Our results show that there is far less correlation of the noise in the responses of adjacent neurons than of the signal. This could arise from a combination of at least three mechanisms.

First, most of the noise in the output of a neuron could be intrinsic to the neuron, that is, not present in its inputs, but generated by internal processes such as random fluctuations in neurotransmitter release or membrane potential. Although intuitively it might seem that noise is detrimental, there is a line of thought suggesting that adding noise to a neuron's responses could serve to linearize its properties (Krölller et al., 1988). In that case large amounts of intrinsic noise might even be a useful property. In spite of this theoretical possibility, it seems unlikely to us that a neuron could have an intrinsic noise significantly greater than the noise in any one of its inputs because these inputs come from other neurons, which would presumably have similar levels of intrinsic noise.




Figure 13. Graphical illustration of how information adds with increasing numbers of mutually redundant neurons. The amount of transmitted information is represented by the area of the circles. With neuron *a* alone, the total amount of information is just the area of *a* (*A*). When we add neuron *b* to neuron *a*, the total information is increased by less than the area of *b* alone, represented by the unshaded area of *B*. The total amount of information is the union of *a* and *b*, which is

$$a \cup b = a + (b - (a \cap b)).$$

When a third neuron is added, the amount of new information added to the total is smaller than when *b* was added to *a*, again represented by the unshaded area in *C*. The total amount is now

$$a \cup b \cup c = a + (b - (a \cap b)) + (c - (c \cap a) - (c \cap b) + (a \cap b \cap c)).$$

We assume a population of neurons that each carry one bit of information, and that each pair has the same fraction of information that is redundant between them, y , where y varies from 0 (no overlap) to 1 (complete overlap). Then the total amount of information when a second neuron's response is considered along with that of the first will be $T_{\text{total}} = 1 + (1 - y)$, where y corresponds to the $a \cap b$ term. When a third neuron is added, carrying information that is mutually redundant with each of the other two, then the total information will be $T_{\text{total}} = 1 + (1 - y) + (1 - y - y + y^2)$. The y^2 term corresponds to the $a \cap b \cap c$ term above. It can be inferred that, for a pool of n neurons that satisfy these constraints, the total amount of information is $T_{\text{total}} = \sum_{i=0}^{n-1} (1 - y)^i$. This converges to the value $1/y$, so even for arbitrarily large numbers of neurons, if they are all mutually redundant to the same degree, then no more than this amount of information can be carried. If neurons carry other than 1 bit, this result should be multiplied by that amount. For the neurons in this study, the joint information was an average of 0.9 of the sum of the individual information (see Fig. 8), indicating that the value of y for these neurons is 0.2. While one might at first assume that this value should be 0.1, consider that the total amount of information is twice the 0.8 of the information for each neuron, plus the redundant 0.2 amount of information counted only once. The expression $n = \log(1 - f)/\log(1 - y)$ gives the number of neurons n needed to carry f percent of the maximum possible information in a pool with y amount of mutual redundancy between all pairs. For a value of y of 0.2, 95% of the maximum amount of information would be carried with just under 14 neurons.

Second, the inputs could be segregated, and adjacent neurons could be receiving afferent fibers from completely different neurons, yet still be constructing responses that have some overlap. This is possible, but seems unlikely, because the significant cross-correlograms indicate that a majority of the pairs of neurons in this study share some common inputs or are closely connected in some way.

Third, the independence of the noise may be a result of different temporal filtering. Suppose that two adjacent neurons have different time courses in their sensitivities to their inputs, that is, different temporal filtering properties. Different temporal filters can have overlapping sensitivities but uncorrelated noise, because they look at different parts of the same input signal. If this turns out to be the case, then not only is temporal modulation of neuronal signal present, but it is essential for the normal functioning of neurons in IT cortex, and perhaps for the whole visual system as well.

Relationship between spike trains and signals

Over 60% (17 of 28) of the pairs of neurons in this study had significant cross-correlograms. The cross-correlogram defines an effective connectivity between pairs of neurons, from which constraints on the pattern of local connections are inferred (Melsens and Epping, 1987). The cross-correlogram is highly stimulus dependent, and a nonsignificant cross-correlogram may be due to the specific stimulus set and/or experimental conditions, and does not rule out the possibility of shared connections (Melsens and Epping, 1987). Therefore, the 11 pairs that did not have significant cross-correlograms might well have had them under different conditions. We can only conclude that at least 17 pairs shared common inputs and/or were relatively closely connected in some way. However, the responses of these pairs were all nearly independent, suggesting that even when neurons share the same connections, both the signal they carry and the noise on this signal show only a little correlation. If these neurons were redundant, that is, processed their inputs identically, then the commonality of the connections shown by the cross-correlograms would have made it difficult if not impossible to obtain the low values of noise correlation actually found. Thus, the action potentials may be considered the carrier, and the response (as quantified by the mean firing rate and temporal modulation) the signal, which is transmitted by modulating the carrier.

As mentioned earlier, we have defined the signal as consisting of the first few principal components, because they account for most of the variance in the responses of single neurons, and for most of the stimulus-related information as well (Optican and Richmond, 1987). However, it is possible that there is stimulus-related information in the later principal components when the responses of two or more neurons are considered jointly. If this were true, then the cross-correlogram could be related to the actual signal, and the individual spikes could not be considered to be solely part of the carrier. Testing this possibility would require a different experimental design with many more trials per stimulus, to compensate for the increased dimensionality of the data, and was beyond the scope of this study.

Sampling bias

Any study using currently available techniques will yield a biased sample of neurons. It is well known that the characteristics of the microelectrodes used will select for neurons of a particular size (Abeles, 1982). For example, electrodes like the ones we used with tip impedances on the order of 1 M Ω will tend to select from the electric fields of cell somas, and are more likely to isolate signals from larger than from smaller cells. But the experimental paradigm and stimulus set also introduce bias. As noted earlier, we encountered many neurons that were apparently not affected by our stimulus set, and we recorded only when at least one neuron was clearly driven by it. Presumably, there were many neurons in IT cortex that did not respond at all, even though they would almost certainly be active with either different stimuli or different behavioral paradigms. The heterogeneity of a local area of cortex is probably far greater than the results of this study show, and by recording only from cells that were active under our conditions, we have biased our results to show more similarity between neurons than actually exists. In short, the degree of correlation between adjacent neurons that we found is an upper bound.

Our stimulus set was restricted in that it consisted of only 32

black-and-white patterns, all of the same scale and position, and the same behavioral relevance. Yet in other ways it was quite varied, in that it consisted of an orthogonal set of patterns that completely spans the space of black-and-white, 4×4 element patterns. The bias of an experiment is an inverse function of the richness of the stimulus set: if we had used two stimuli that differed from each other in only one feature, then by definition we would have found that the responses to them were highly correlated. The more restricted the stimulus set, the lower the chance of correctly determining the functions and relationships of two neurons. Furthermore, selecting stimuli based only on neuronal responses that seem favorable to the experimenter can easily lead to biased results that cannot be regarded as generic.

Independence versus similarity

In one respect two neurons that respond to the same stimuli are similar, for the obvious reason that they are both responding to the same stimuli. We found that the responses of both neurons of an adjacent pair were usually strongly modulated by our stimulus set, even though we only required this behavior of one neuron. Considering that we encountered many regions where we were unable to find any neurons that responded to our stimuli, this suggests that neurons in IT cortex are arranged in patches. However, if the responses of one neuron can only be poorly predicted from the responses of the other neuron, as we found in this study, then the responses must be nearly independent. Just because two neurons tend to respond under the same stimulus conditions does not mean that they are performing identical functions. It is true that most of the neurons in this study were similar in some ways, in that they responded strongly to the same restricted stimulus set. But, as shown by the results of the linear regression and information theoretic analyses, how they responded to the different elements of this set was very different, that is, poorly correlated, from one neuron to the next.

Local field potential

It is possible that there is massive redundancy in a local area of cortex, but that the technique of recording pairs of neurons from a single microelectrode is not capable of detecting it. For example, consider that there are three kinds of cells, *A*, *B*, and *C*, arranged in an alternating sequence *ABCABCABCABC*. In this example, adjacent neurons would always be different, but there would still be massive redundancy.

Evidence against the presence of neurons with similar responses existing in a local population of IT neurons comes from other studies. Spike-triggered averaging of the local field potential (LFP; 10 Hz to 10 kHz) of IT neurons in the same monkeys studied here showed no relationship between the LFP and the action potentials of a visually responsive cell (Gawne et al., 1991), suggesting that neurons in a local area of cortex respond to a visual stimulus in ways that are only very weakly correlated.

Multiunit separation

As mentioned before, our technique of multiunit separation cannot separate the signals from different units that occur simultaneously (see Materials and Methods). Because the time during which overlapping spikes cannot be detected is on the order of half a millisecond, the combined rate of firing of a pair of neurons would have to become a significant fraction of 2000 spikes per second for the number of overlaps to affect our results significantly. The neurons recorded in this study had combined

firing rates well below this level: even the peak rates rarely exceeded 100 spikes per second (see Fig. 5).

Although unlikely, pairs of spikes could occur simultaneously in a much greater proportion than would be predicted by the firing rates alone. In an extreme case, two spikes, each from a different cell, might only occur simultaneously, and be recorded as a single unit. In this extreme case the two neurons would be identical in their signal processing, and could be different only in their efferent targets. All existing technologies for isolating extracellular potentials—including those that only attempt to isolate a single unit—can be fooled by specific, systematic relationships in the firing pattern of the neurons. Only simultaneous intracellular recordings can completely rule out all possible artifacts stemming from the need to isolate the signals from multiple units.

Other studies

In another recent study that investigated the relationships between adjacent neurons in the IT cortex of awake behaving macaques (Gochin et al., 1991), only 80 out of 155 neurons showed a change in activity with presentation of a stimulus, and only 55 out of 155 neurons showed significant differential response modulation by the stimulus set. We isolated 28 responsive pairs, compared with 4 isolated pairs that were determined to be not responsive and 23 areas with differentiable units where we did not even attempt to isolate them because they were not clearly driven by our visual stimuli. Therefore, the proportion of responsive neurons we isolated was roughly the same as in Gochin et al. (1991). Also, although the measure of similarity across neuronal responses was different (they did not separate out the responses of a neuron into signal and noise), the basic result was that, whereas adjacent neurons are somewhat more correlated than neurons that are farther apart, they are still only weakly correlated, just as in the present study.

Comparisons of psychophysical and single-unit data show that massive redundancy across neurons is not needed for at least some functions of cortex. Tolhurst et al. (1983) estimated that psychophysical detection may be explained by the activity of only a few (two to eight) neurons. Newsome et al. (1989) found similar results for the detection of motion by both psychophysical measures and single cells in area MT. We suggest that if there is no advantage to be gained from having multiple identical neurons, then perhaps these neurons do not exist.

Tolhurst and Thompson (1982) found great variability of spatial frequency tuning in neurons that were relatively close (about 100 μm) in visual cortex. They hypothesized that, whereas a cortical column may be specialized for one orientation, there are many different subgroups for different spatial frequencies. The situation may be analogous in IT cortex: adjacent neurons often respond to the same stimulus set, but in different ways. This conclusion is supported by data from IT cortex reported recently by Fujita et al. (1992).

Barlow (1989) has argued that neurons should tend toward independence, because this would maximize the total amount of information about a sensory stimulus in a population, and because multiple independent elements have many advantages in signal processing. Linsker (1988) proposed, on theoretical grounds, that there is a balance between independence and redundancy in the responses of neurons that maximizes the information-carrying capacity of a population. This study explicitly assumed that noise is independent, even with redundant neurons. To the extent that two neurons are redundant and share




Figure 14. Proposed model for how adjacent neurons are organized. Two neurons each get sets containing identical afferents. Each afferent carries a combination of signal (S_1 and S_2) and noise (N_1 and N_2). The two neurons have different functions of the afferents, represented as lines of opposite slope, such that their outputs are independent. The intrinsic noise terms represent sources of noise that are not a function of the noise on the afferents, such as random fluctuations in membrane potential. This type of transformation will preserve information, and prevent the noise from being correlated.

the same afferents, it isn't possible for the noise to be independent. We suggest that the need to avoid correlated noise may push the balance between redundancy and independence even further toward independence than indicated by Linsker's analyses.

Occasionally mixtures of action potentials are recorded with the assumption that nearby neurons have similar properties. To reach conclusions about the information processing of neurons, our results suggest that this assumption is unwarranted, and care must be exercised to isolate reliably the waveforms generated from different neurons.

Techniques that average signals across many neurons, for example, recording, metabolic recording, and so on, are used frequently. By definition, averaging techniques remove the influence of individual variations. If the individual neurons in a local area all carried similar signals, then averaging techniques would reflect the messages actually used by the nervous system. However, our data show that the messages carried by individual neurons are quite different, and therefore averaging techniques will not be able to measure the messages carried by the neurons. These techniques can be quite valuable, in that they may give an indication of the relative change of the amount of processing going on in an area of cortex, but they are inappropriate for determining how that processing is carried out.

Conclusions

This study suggests that adjacent IT cortical neurons are not organized in massively redundant populations, but instead are heterogeneous, and are more independent than redundant. The noise between adjacent neurons was even more uncorrelated than the signal, and we suggest that ensuring that this is the case is a major requirement for the organization of the nervous system. If neurons have identical characteristics, and share the same inputs, then the noise must be highly correlated, and there will be no advantage to having many neurons, save of course

as replacements for lost ones. We suggest further that having different temporal filtering properties is one way in which the noise between two neurons can be made even more uncorrelated than the signal. Avoiding the consequences of correlated noise and redundant information would provide strong pressure favoring selection of connection patterns and weights that keep the responses of adjacent neurons as nearly independent as is biologically possible both for signal and noise (Fig. 14).

References

- Abeles M (1982) Studies of brain function local cortical circuits. Berlin: Springer.
- Abeles M, Goldstein MII (1977) Multispikes train analysis. *Proc IEEE* 65:762-773.
- Abramson N (1963) Information theory and coding. New York: McGraw-Hill.
- Ahmed N, Rao KR (1975) Orthogonal transforms for digital signal processing. Berlin: Springer.
- Barlow HB (1989) Unsupervised learning. *Neural Comput* 1:295-311.
- Bruesch SR, Arey LB (1942) The number of myelinated and unmyelinated fibers in the optic nerve of vertebrates. *J Comp Neurol* 77:631-665.
- Drapere N, Smith H (1980) Applied regression analysis, 2d ed. New York: Wiley.
- Eggermont JJ, Epping WJM, Aertsen AMHJ (1983) Stimulus dependent neural correlations in the auditory midbrain of the grassfrog (*Rana temporaria* L.). *Biol Cybern* 47:103-117.
- Eskandar EN, Richmond BJ, Optican LM (1992) Role of inferior temporal neurons in visual memory: I. Temporal encoding of information about visual images, recalled images, and behavioral context. *J Neurophysiol* 68:1277-1295.
- Fujita I, Tanaka K, Ito M, Cheng K (1992) Columns for visual features of objects in monkey inferotemporal cortex. *Nature* 360:343-346.
- Gawne TJ, Richmond BJ (1990) Determining the independence of messages carried by multiple simultaneous neurons. *Soc Neurosci Abstr* 16:230.
- Gawne TJ, Eskandar EN, Richmond BJ, Optican LM (1991) Oscillations in the responses of neurons in inferior temporal cortex are not driven by stationary visual stimuli. *Soc Neurosci Abstr* 17:443.
- Gochin PM, Miller EK, Gross CG, Gerstein GL (1991) Functional interactions among neurons in inferior temporal cortex of the awake macaque. *Exp Brain Res* 84:505-516.
- Harmuth HF (1977) Sequence theory: foundations and applications. New York: Academic.
- Judge SJ, Richmond BJ, Chu FC (1980) Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res* 20:535-538.
- Kröllner J, Grüsser OJ, Weiss LR (1988) Superimposing noise linearizes the responses of primary muscle spindle afferents to sinusoidal muscle stretch. *Biol Cybern* 60:131-137.
- Krüger J, Aiple F (1988) Multimicroelectrode investigation of monkey striate cortex: spike train correlations in the infragranular layers. *J Neurophysiol* 60:798-828.
- Linsker R (1988) Self-organization in a perceptual network. *Computer* 21:105-117.
- McClurkin JW, Gawne TJ, Optican LM, Richmond BJ (1991a) Lateral geniculate neurons in behaving primates: II. Encoding of visual information in the temporal shape of the response. *J Neurophysiol* 66:794-808.
- McClurkin JW, Optican LM, Richmond BJ, Gawne TJ (1991b) Concurrent processing and complexity of temporally encoded neuronal messages in visual perception. *Science* 253:675-677.
- Melssen WJ, Epping WJM (1987) Detection and estimation of neural connectivity based on crosscorrelation analysis. *Biol Cybern* 57:403-414.
- Morrison DF (1967) Multivariate statistical methods. New York: McGraw-Hill.
- Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. *Nature* 341:52-54.
- Optican LM, Richmond BJ (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex: III. Information theoretic analysis. *J Neurophysiol* 57:162-178.
- Optican LM, Gawne TJ, Richmond BJ, Joseph PJ (1991) Unbiased

- measures of transmitted information and channel capacity from multivariate neuronal data. *Biol Cybern* 65:305–310.
- Richmond BJ, Optican LM (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex: II. Quantification of response waveform. *J Neurophysiol* 57:147–161.
- Richmond BJ, Wurtz RH, Sato T (1983) Visual responses of inferior temporal neurons in the awake rhesus monkey. *J Neurophysiol* 50:1415–1432.
- Richmond BJ, Optican LM, Podell M, Spitzer H (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex: I. Response characteristics. *J Neurophysiol* 57:132–146.
- Richmond BJ, Optican LM, Spitzer H (1990) Temporal encoding of two-dimensional patterns by single units in primate primary visual cortex. I. Stimulus–response relations. *J Neurophysiol* 64:351–369.
- Robinson DA (1963) A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE Trans Biomed Eng* 10:137–145.
- Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27:379–423.
- Silverman BW (1986) *Density estimation for statistics and data analysis*. London: Chapman and Hall.
- Tolhurst DJ, Thompson ID (1982) Organization of neurones preferring similar spatial frequencies in cat striate cortex. *Exp Brain Res* 48:217–227.
- Tolhurst DJ, Movshon JA, Dean AF (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23:775–785.
- Wurtz RH (1969) Visual receptive fields of striate cortex neurons in awake monkeys. *J Neurophysiol* 32:727–742.