

Responses of neurons in primary and inferior temporal visual cortices to natural scenes

Roland Baddeley(1,2), L.F. Abbott(3), Michael C.A. Booth(1),
Frank Sengpiel(2), Toby Freeman(2),
Edward A. Wakeman (1), and Edmund T. Rolls(1)

(1) Department of Experimental Psychology, University of Oxford

(2) Department of Physiology, University of Oxford

(3) Volen Center and Department of Biology
Brandeis University, Waltham, MA 02254 USA

September 16, 1997

The primary visual cortex (V1) is the first cortical area to receive visual input, and inferior temporal (IT) areas are among the last along the ventral visual pathway. We recorded, in area V1 of anaesthetised cat and area IT of awake macaque monkey, responses of neurons to videos of natural scenes. Responses were analysed to test various hypotheses concerning the nature of neural coding in these two regions. A variety of spike-train statistics were measured including spike-count distributions, interspike interval distributions, coefficients of variation, power spectra, Fano factors and different sparseness measures. All statistics showed non-Poisson characteristics and several revealed self-similarity of the spike trains. Spike-count distributions were approximately exponential in both visual areas for eight different videos and for counting windows ranging from 50 ms to 5 seconds. The results suggest that the neurons maximise their information carrying capacity while maintaining a fixed long-term-average firing rate (Baddeley, 1996; Levy & Baxter, 1996), or equivalently, minimise their average firing rate for a fixed information carrying capacity.

1 Introduction

It has been suggested that visual representations are optimised to transmit the maximum information about the images encountered in everyday life (Uttley, 1973; Linsker, 1988; Barlow, 1989). This simple assumption has proven sufficient to account for the characteristics of large monopolar cells in the fly (Srinivasan *et al.*, 1982; Hateren, 1992; Laughlin, 1981), the temporal characteristics of retinal ganglion cells (Dong & Atick, 1995), human spatial frequency thresholds (Atick & Redlich, 1992; Van Hateren, 1993), and the psychophysics of orientation perception for short presentation times (Baddeley & Hancock, 1991).

Maximisation of information is a powerful theoretical principle that leads to testable predictions about the firing patterns of neurons. However, to generate specific predictions we must make some assumptions about the nature of the neural code and the type of constraint that limits its information carrying capacity. To apply information maximisation to neuronal spike trains, we must identify which of their characteristics carry information. In our analysis, we will consider two possibilities: that firing rates, or more precisely, spike counts over discrete intervals of time, are the information carrying elements; or that interspike intervals play this role. Without any constraints on the rate or precision of neuronal spiking, the information carrying capacity of a spike train is infinite. Thus, constraints play a crucial role in any information maximisation procedure. We will consider three possibilities, constraints on the maximum firing rate, the average firing rate, or a quantity known as the sparseness of the firing-rate distribution. Identifying the nature of the constraint that limits information carrying capacity has important implications for the biophysical mechanisms that underlie neural coding.

Assuming the firing rates carry information, Laughlin (1981) proposed a constraint on the maximum firing rate for neurons in the insect eye, and noted that the optimal firing-rate distribution in this case is flat (Figure 1A). A different proposal, that the distribution of firing rates should be sparse (Field, 1994; Olshausen & Field, 1996), has led to predictions for the receptive field properties of neurons in V1 similar to those seen in the data (Olshausen &

Field, 1996). An example of a sparse distribution is shown in Figure 1B. Here the number of ‘active’ neurons, neurons with significantly elevated firing rates, is small. A third proposed constraint is that the average firing rate of the neuron is held fixed while the information carrying capacity is maximised (Levy & Baxter, 1996; Baddeley, 1996). If information is carried by firing rates, this predicts that the distribution of firing rates should be exponential as in Figure 1C (Shannon & Weaver, 1949; Levy & Baxter, 1996; Rieke *et al.*, 1997)). If, instead, information is encoded by interspike intervals, this constraints predicts an exponential distribution of interspike intervals, that is, a Poisson distribution (DeWeese, 1996; Rieke *et al.*, 1997). Simulation results suggest that maximising information transmission for a fixed average firing rate also leads to receptive fields qualitatively similar to those found in V1 (Harpur & Prager, 1996).




Figure 1: Three proposals for optimal firing-rate distributions. All distributions are plotted on log-linear axes so that an exponential distribution appears as a straight line. **A)** If the maximum firing rate is constrained, the information carrying capacity of firing rates is maximum for a flat distribution. **B)** An example of a sparse distribution. **C)** If the constraint is on the average firing rate, the optimal distribution is exponential.

To investigate the nature of the constraints relevant for neural coding, we measured firing rate distributions and other spike-train statistics of neurons in two visual areas. We chose visual areas at opposite ends of the cortical visual processing stream, V1 and IT, to compare coding strategies. Since the distribution of the images used to generate the spike trains affects the statistical properties we study, we used videos that resemble, as nearly as possible, the normal visual environments of the animals from which the neurons were being recorded.

2 Methods and Stimuli

To stimulate the natural visual environment of a cat as closely as possible, we recorded three videos by ‘walking’ a video camera 30 cm above the ground in locations where cats are commonly found. The *mixed* video contained both natural and man-made features (length of video, 10 minutes; number of cells, 22). The *natural* video consisted almost entirely of natural vegetation (10 min, 6 cells), and the *indoor* video was shot in the laboratory and contained only man-made features (5 minutes, 4 cells). For comparison, a subset of the neurons was also tested with high contrast white-noise stimuli generated by removing the video input from the video player (10 minutes, 16 cells). The stimuli were recorded using a high quality video camera using VHS format and displayed on a black and white monitor (14 inch for the V1 cells, 20 inch for the IT cells).

While recording from two macaques, we presented three different videos: a *laboratory* video (5 min, 26 cells), a *colony* video taken of a wild monkey colony (5 min, 4 cells), and an *artificial* video of a popular Australian soap opera (5 min, 2 cells). Again for comparison we recorded cells while the monkeys were watching a blank screen (5 min, 19 cells).

Details of the recording and anaesthesia (for the cats) can be found elsewhere for cat V1 (Sengpiel *et al.*, 1994), and for macaque IT (Rolls & Tovee, 1995). After characterising the neurons using standard techniques (optimal orientation, spatial frequency, and direction tuning for V1; face and object selectivity for IT), the monitor was positioned 57 cm away from the cats and 1 m away from the monkeys, and spikes were recorded while a given video sequence representative of natural scenes was played. Data were collected from 3 cats, and 2 macaques. The cats were anaesthetised and the macaques were awake and free viewing. Only visually responsive cells were used.

2.1 Calculating spike-count distributions

We computed spike-count distributions by sliding a window of size T along the recorded spike trains, in steps of $T/6$ (to minimise the effects of window boundaries), and counting how many spikes occurred within the temporal window for each window position. The normalised spike-count frequency distribution provides an estimate of the probability distribution of spike counts for a given window size. This is equivalent to a firing-rate distribution up to a factor of T . To facilitate comparison of cells with different average firing rates and to assure that all cells were analysed with the same resolution, many of the results shown used a window size adjusted so that the average number of spikes per window was fixed. To do this we used a window size of N/f where N is the desired average number of spikes per window, and f is the average firing rate of the cell.

3 Analysis and Results

3.1 Firing rates

We first measured the average firing rates of the cells. For V1 of the anaesthetised cat, the firing rates for the video-stimulated neurons were low (mean = 3.96 Hz, sd = 3.61 Hz). This was lower than has been previously reported (Legéndy & Salcman, 1985) for the unanaesthetised cat (mean = 8.9 Hz, sd = 7.0 Hz), but was significantly higher than when the cells were stimulated with high contrast white noise (mean = 2.45 Hz, sd = 2.18 Hz). It is proposed that the low average rates were partly due to the effect of the anaesthetic (which could be tested by systematically varying the level). For the macaque IT cells, generally in the upper bank of the superior temporal sulcus at sites similar to those in (Rolls & Tovee, 1995), the average rate was higher for both video stimulation (mean = 18 Hz, sd = 10.3 Hz), and blank screen viewing (mean = 14 Hz, sd = 8.3 Hz).

These rates are low compared to the firing rates observed when these cells are stimulated with optimal stimuli. For instance, when the V1 neurons were stimulated with optimally oriented high contrast sinusoidal gratings the firing rates were on average 37.0 Hz (sd = 33.9 Hz). The low firing rates observed are probably attributable to the fact that natural video sequences rarely contain visual structures that are optimal for activating the cells, and also

due to the inhibitory effects that arise when cells respond to complex natural scenes (Gallant, 1996).

3.2 Spike-count distributions

Examples of the spike-count distributions computed with a window size of $T = 250$ ms, are shown in Figure 2. Like the neurons shown, the majority of neurons measured had spike-count distributions that were approximately exponential. The graphs in Figure 2 use log-linear axes so an exponential appears as a straight line.

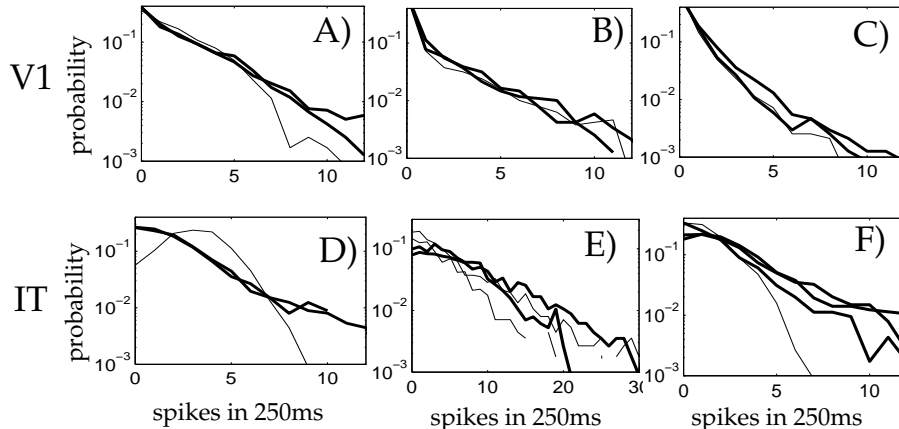


Figure 2: A representative collection of firing rate distributions. The spike-count probability distributions for 6 different cells computed using a sliding window of 250 ms. **A**, **B** and **C** shown three cat V1 cells, and **D**, **E** and **F** three macaque IT cells. The thick lines correspond to different replays of the same natural video sequence. The thin lines are for responses to either high contrast white noise (V1) or a blank screen (IT). The vertical axis is logarithmic. Other than the noise condition for **D** and **F**, the results are well approximated by exponential distributions.

An exponential distribution was found for the majority of cells when rate was computed using a 250 ms time window. To test the generality of this result, we calculated spike-count distributions while varying the window size over one order of magnitude. Figure 3 shows the spike-count distributions, averaged over all neurons, for window sizes chosen for each cell so that there were an average of 1, 2, or 10 spikes per window. Choosing the window size corresponding to a fixed average number of spikes rather than a fixed time allowed us to average results from cells with different firing rates. Except for the low spike-count regions of the distributions corresponding to 10 spikes per window, all the distributions are well approximated by an exponential. The fact that the same distribution applies for a wide range of window sizes, except for a change of scale, indicates ‘self similarity’ (Teich, 1989; Teich *et al.*, 1997). This is not a property of most distributions, for instance the shape of the Poisson spike-count distribution is highly dependent on the time scale.

Next, we investigated whether the distribution depended on which video sequence was used to elicit the responses. Figure 4 shows that the form of the distribution was essentially

independent of the particular video used to stimulate the neuron. Indeed, exponential distributions even appeared for the white noise and blank screen viewing conditions, even though they resulted in lower average firing rates. The similarity of the distributions for video and non-video viewing conditions might suggest that the neurons were not responding to specific features of the videos. However, a study of the neural responses in relation to the video images showed that this was not the case. For example, IT neurons known to respond to face images, fired vigorously when a face appeared in the video.

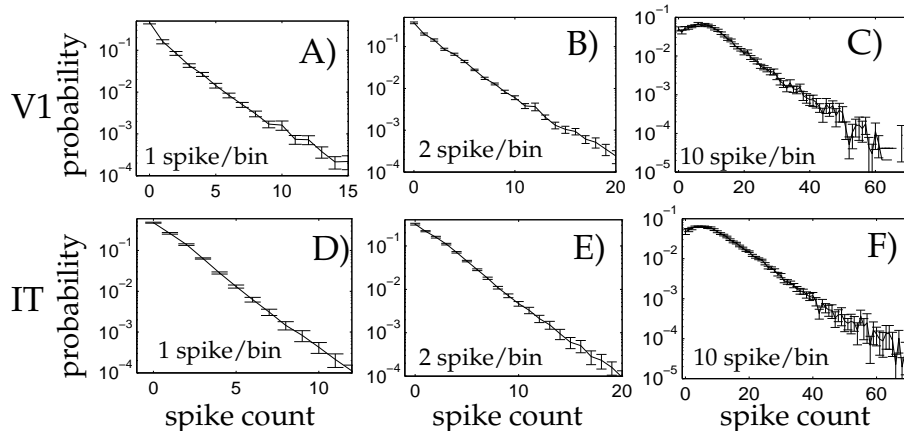


Figure 3: The average spike-count distributions for the V1 and IT neurons is relatively insensitive to the window size used to calculate the rate. To display the average distributions for all the video stimulated cells (each with a different average firing rate), we used a window size adjusted so that the average number of spikes per window was the same for all the cells, either 1, 2 or 10 spikes per window. Error bars are standard errors.

3.3 Interspike intervals

In addition to spike-count distributions, a number of other statistics provide information about the nature of neural coding. The first is the distribution of interspike intervals (ISIs) which can be used to identify a Poisson distribution, for example. We computed the ISI distribution (shown in Figure 5A and B after normalising to the mean interspike interval for each cell and averaging across cells). The ISI distributions do not appear to be Poisson and are described better by a power-law relation. This is indicated in Figure 5A and B by the approximately straight lines appearing on a log-log plot.

The coefficient of variation, CV , measures the normalised variability of interspike intervals, and has been the subject of much theoretical interest (see for instance (Softky & Koch, 1993; Holt *et al.*, 1996)). This measure can be used to identify a stationary Poisson process for which $CV = 1$. CV for N interspike intervals of duration T_j is given by

$$CV = \frac{\sigma}{\langle T \rangle} \quad (1)$$

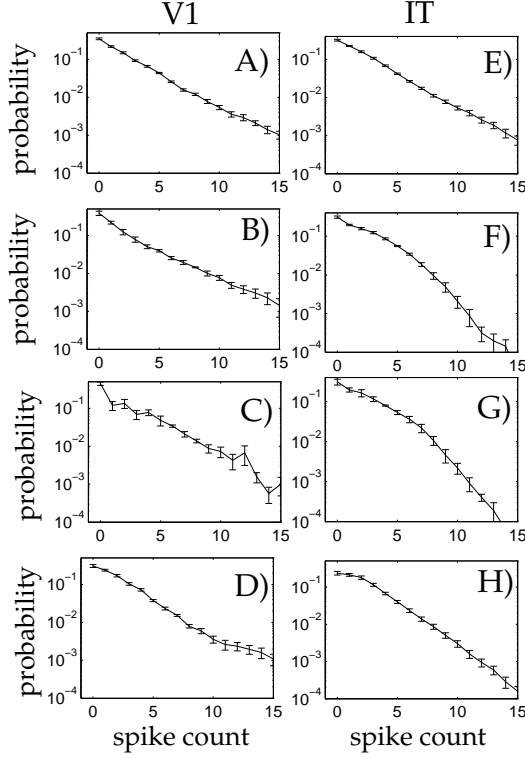


Figure 4: Spike-count distributions for six different video conditions and two noise conditions. In each case the window size was chosen so that the average number of spikes per window was one. Near exponential spike count distributions were found for all three video conditions: *mixed* **A**, *natural* **B**, *indoor* **C**, *laboratory* **E**, *colony* **F**, and *artificial* **G**. More surprisingly, the white noise **D**, and blank screen **H** conditions also produced approximately exponential distributions. Error bars correspond to standard errors.

where

$$\langle T \rangle = \frac{1}{N} \sum_{j=1}^N T_j \quad (2)$$

and

$$\sigma = \sqrt{\frac{1}{N-1} \sum_{j=1}^N (T_j - \langle T \rangle)^2}. \quad (3)$$

Using this measure, neither the V1 nor the IT spike trains are well characterised by a Poisson process. The coefficients of variation were, on average, 1.91 for V1 cells (sd = 0.42, $p(CV \leq 1) < 0.005$) in the video condition. For IT neurons the mean CV was 1.84 (sd = 0.49, $p(CV \leq 1) < 0.005$).

CV is most useful if the spike train is stationary, but this is not the case for most *in vivo* recordings. A local version of CV , CV_n , has been proposed in this case, to determine whether groups of n spikes can be described as arising from a Poisson process (Holt *et al.*, 1996). To compute CV_n , we define the mean ISI between interspike interval i and interspike interval




Figure 5: Normalised ISI distributions, and CV_n as a function of n . For each cell, the distribution of interspike intervals was found, the intervals were divided by the average ISI for each cell and the resulting normalised ISI distributions were averaged over all cells. Results are shown for video stimulated of cat V1 cells (**A**) and macaque IT cells (**B**), with the error bars corresponding to standard errors. **C**) CV_n as a function of n . The thick line indicates the result for cat V1 and the thin line is for monkey IT.

$i + n$ as

$$\langle T \rangle_i^n = \frac{1}{n} \sum_{j=i}^{i+n-1} T_j \quad (4)$$

and the local standard deviation as:

$$\sigma_i^n = \sqrt{\frac{1}{n-1} \sum_{j=i}^{i+n-1} (T_j - \langle T \rangle_i^n)^2}. \quad (5)$$

The local coefficient of variation is then

$$CV_n = K_n \sum_{i=1}^{N-n+1} \frac{\sigma_i^n}{\langle T \rangle_i^n} \quad (6)$$

where K_n is a value chosen so that CV_n is equal to one for a Poisson process. The computed values of CV_n as a function of n are shown in Figure 5C. The values are always great than one. CV_n values for the IT spike trains with small n are closer to the Poisson value which may reflect the persistence time of the images IT is sensitive to (faces, objects). Images that make V1 neurons respond vigorously (oriented lines at particular locations) are likely to have shorter persistence times in the videos. The relationship between CV_n and n appears approximately power law implying that it does not define any particular spike number or time scale.

An additional non-Poisson feature we found was positive correlation between successive interspike intervals. Pearson's product moment correlations, r , between adjacent interspike intervals took average values for the V1 cells under video conditions of 0.11, (sd = 0.1, $p(r \leq 0) < 0.005$) and, under noise conditions, 0.08 (sd = 0.12, $p(r \leq 0) < 0.025$). For IT the correlation for the video conditions was 0.194 (sd = 0.050, $p(r \leq 0) < 0.001$).

3.4 Power spectra

ISI distributions only characterise pairwise relations in firing patterns. A more powerful measure is the Fourier power spectrum of the spike train. Power spectra were computed

using a sampling frequency of 100 Hz (with a resulting Nyquist frequency of 50 Hz) based on overlapping samples of 2.56 seconds, after windowing with a Bartlett window (Press *et al.*, 1992). The results for 3 V1 cells and 3 IT cells, shown in Figure 6, are reasonably typical of the entire collection of neurons. For a Poisson process, the power spectrum is flat, but this was found for only 4 of the 99 cells. The vast majority of cells (92/99) had the greatest power at low frequencies (0 to 5 Hz) with the power decreasing monotonically for higher frequencies. For 3 cells the peak power was not at the low frequency end. For 52 out of the 99 cells, the plot of power for frequencies between 1 and 10 Hz was straight on log-log scales, another indication of self-similarity. This low frequency bias contrasts with LGN cells stimulated with natural scenes (Dan *et al.*, 1996), and may reflect a difference in the correlation times of the objects coded for (efficiently coded intensity versus the presence of faces and line segments).

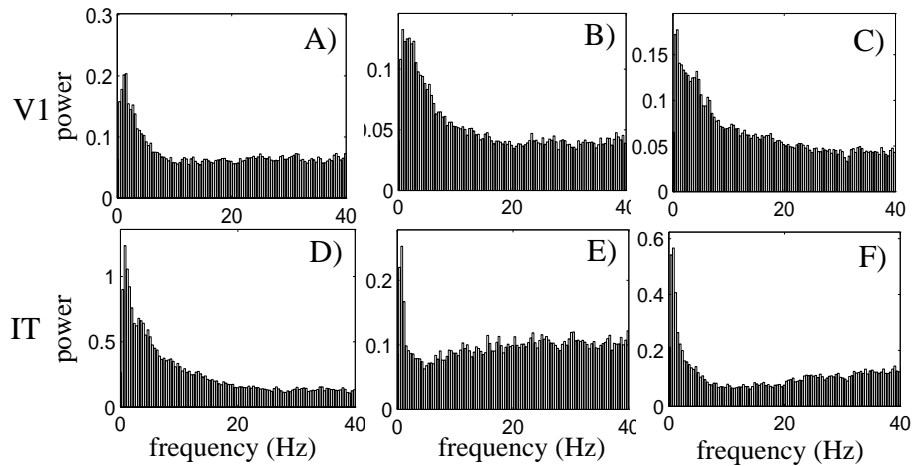


Figure 6: Representative power spectra for neurons from cat V1 (**A** complex, **B** simple, and **C** simple) and macaque IT (**D-F**). The vast majority of cells had excess power at low frequencies.

3.5 Fano factor analysis

Fano factors (Teich *et al.*, 1997) provide a convenient statistic for describing spike count variability over a range of different time scales. The Fano factor, $F(T)$ is the ratio of the variance to the mean of the spike-count distribution computed using a window of size T .

For a Poisson process, $F(T) = 1$ for all T , and if the system is fractal, the Fano factor is a power law function of T , $F(T) \propto T^\alpha$, where α is the fractal or scaling exponent (Teich *et al.*, 1997). Fano factors for our data are shown in Figure 7A. None of the cells appeared Poisson. Instead, a power law provides a very good approximation for the average Fano factors indicating the fractal nature of the firing patterns. The exponents for the IT cells (Figure 7C) are consistently higher than those for the V1 cells (Figure 7B) indicating more low frequency variation and confirming the analysis using CV_n . For V1 cells, the average α was 0.29 (sd = 0.165) and, for IT neurons, 0.53 (sd = 0.12).




Figure 7: **A)** The average Fano factors for the V1 neurons (thick line) and IT neurons (thin line). The error bars correspond to standard errors. The average Fano factor is well-approximated by a power law (straight line for log-log axes). We calculated the exponents α for all cells and histograms of the exponents are shown for the V1 cells (**B**) and IT cells (**C**). The arrows show the average exponents.

3.6 Sparseness

Both V1 and IT neurons had sparse firing-rate distributions as measured using Olshausen and Field’s definition (Olshausen & Field, 1996). A possible interpretation of the sparseness idea is that successively sparser representation should be generated at higher levels of the visual pathway. Indeed, the raw input is less sparse than the representation in LGN, which in turn is less sparse than in V1 (Field, 1994). We find no evidence for increasing sparseness as we move from the representations in V1 to IT. As seen in Figure 8, the degree of sparseness is not dramatically different for the two cortical areas and which area has the higher sparseness depends on the particular measure used.

4 Discussion

4.1 The empirical findings

The main empirical contribution of this study is to extend various statistical analyses of neural spike train data performed on artificially stimulated neurons, to neurons stimulated with naturalistic video sequences. Our main results are:

1. The average firing rates of the cells were low compared to those typically reported for neurons responding to optimal stimuli. Low firing rates have previously been reported for naturally stimulated cells in V1 (Legéndy & Salcman, 1985) and in other visual areas (Gallant, 1996).
2. The firing patterns for naturally stimulated cells are far from Poisson. This is indicated by ISI statistics, power spectra, Fano factors, and spike-count distributions. Non-Poisson characteristics of spontaneous and stimulus-induced firing have been reported by (Teich *et al.*, 1997) and (Bair *et al.*, 1994).
3. The spike trains show evidence of self similarity as demonstrated by the Fano factor analysis, by the spike-count distributions, and also by the approximately power law form of the ISI distribution and some of the power spectra. Self-similarity has been

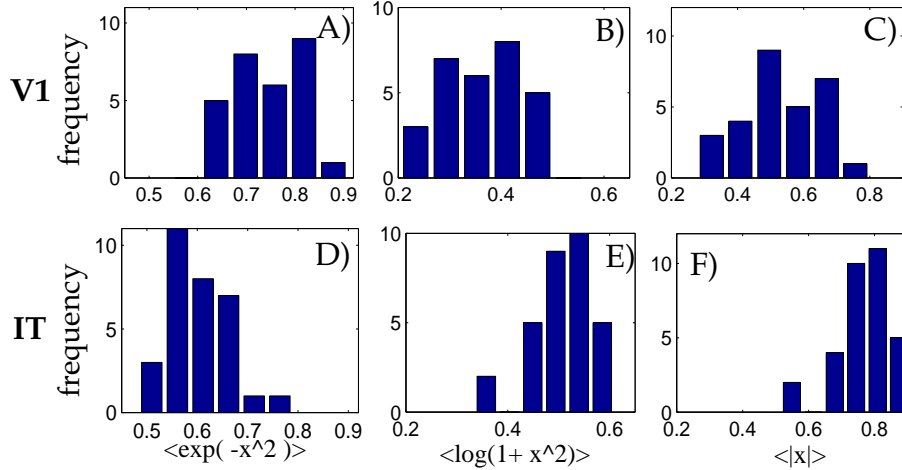


Figure 8: Sparseness distributions for the V1 and IT cells using the three measures proposed by Olshausen and Field (1996). **A** and **D** show $\langle \exp(-\mathbf{x}) \rangle$, **B** and **E** show $\langle \log(1 + \mathbf{x}) \rangle$, and **C** and **F** show $\langle |\mathbf{x}| \rangle$ where \mathbf{x} is the vector of neuronal firing rates divided by the root mean square firing rate for each cell, and $\langle \dots \rangle$ denotes the average over the entire spike train. The bin size used to estimate sparseness was 500 ms. There is one technical difficulty in directly using the original definitions since the theoretical neurons of Olshausen and Field could fire with both positive and negative rates. We therefore divided by the root mean squared output rather than the standard deviation in calculating sparseness.

reported previously for artificially stimulated V1 neurons (Teich *et al.*, 1996), but this is the first report for naturally stimulated neurons.

4. The spike-count distributions are approximately exponential over a large range of window sizes. This has been suggested for theoretical reasons (Levy & Baxter, 1996; Baddeley, 1996) but has not been reported previously as an experimental finding.

4.2 Implication for coding

The statistics of the spike trains that we observe are consistent with an optimised rate code. The observed spike-count distributions are of nearly maximum entropy if the average firing rate is constrained (see Figures 9A and B). Over a two and a half order of magnitude variation in the window size used for spike counting, the entropy is approximately constant and near its maximum possible value. An alternative way to express this result is that for a given entropy, the code minimises the average firing rate. This is seen in Figures 9C and D which compare the actual average firing rate to the minimum rate consistent with the observed spike-count entropy.

This interpretation assumes that firing rate is the appropriate variable to consider when studying coding issues. The output of a neuron presents a potentially severe bottleneck to communication. Using each interspike interval to carry information, rather than spike counts over discrete time intervals, might appear to be a more efficient coding method. The optimal distribution for such a code is a Poisson distribution and this is not consistent with

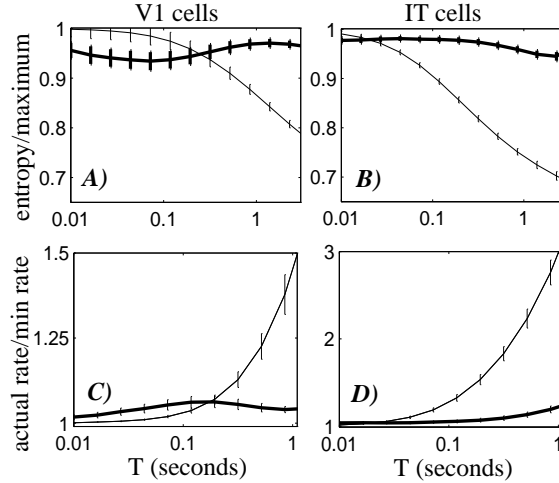


Figure 9: The observed firing rates are consistent with an efficient rate code if average firing rate is a constraint. For each cell, the entropy of the spike-count distribution was calculated and divided by the maximum possible spike-count entropy for a neuron firing at the same average rate. The value of this quantity averaged over all cells is indicated by the thick line for V1 neurons (**A**) and for IT neurons (**B**). For comparison, the thin lines show the same quantities for a Poisson process. **C** and **D** show ratios of the actual average firing rates of the neurons to the minimum possible rates for cells with output distributions of the same entropy (again the thin line corresponds to a Poisson process). The firing rate is near minimal for the given output entropy over two orders of magnitude of window size. Error bars represent standard errors. The ‘time’ indicated on the horizontal axis is the spike counting window size.

our results. Psychophysical experiments indicate that visual (Watt, 1987), and auditory (Viemeister, 1996) resolution for many tasks increases over a period of up to a second. For a representation based on interspike intervals, a complicated mechanism for integrating over successive interspike intervals would be required to display such behaviour. A rate code would naturally show an increase of resolution over time.

The proposal that average firing rate is being constrained while information is being maximised has many similarities to the ‘sparsity’ proposal (Field, 1994; Olshausen & Field, 1996). Since one of the definitions of sparsity equates with average firing rate ($\langle |\mathbf{x}| \rangle$), networks that maximise information transmission for fixed average firing rate results in V1-like receptive fields when trained to represent collections of natural images (Harpur & Prager, 1996; Olshausen & Field, 1996). However, rather than restricting the number of active elements, we suggest that the code is designed to maximise information transmission at fixed average rate with all rates conveying information, not just the most active neurons.

The visual cortex has among the highest oxygen consumptions of any part of the brain and accounts for 10% of brain volume in a typical sighted animal (Diamond, 1996). In children the brain can account for up to 50% of the resting oxygen consumption (Sokoloff, 1989). Therefore a code that minimised the average firing rate, and hence metabolic activity, could make a significant difference to the energy consumption of an animal. Assuming that the code in these visual areas is a rate code, the measurements presented here show that information

carrying capacity is near maximum values for a given average firing rate over a large range of time scales, and at two opposite poles of the cortical visual pathway.

Research supported by National Science Foundation grant DMS-9503261, the W.M. Keck Foundation (LA), the MRC centre for Cognitive Neuroscience, and the Wellcome Trust.

References

- Atick, J.J., & Redlich, A.N. 1992. What Does the Retina Know About Natural Scenes. *Neural Computation*, **4**, 196–210.
- Baddeley, R.J. 1996. An Efficient Code in V1? *Nature*, **381**, 560–561.
- Baddeley, R.J., & Hancock, P.J.B. 1991. A statistical analysis of natural images matches psychophysically derived orientation tuning curves. *Proc.R.Soc.London Ser.B*, **246**, 219–223.
- Bair, W., Koch, C., Newsome, W., & Britten, K. 1994. Power spectrum analysis of bursting cells in area MT in the behaving monkey. *Journal of Neuroscience*, **14**, 2870–2892.
- Barlow, H.B. 1989. Unsupervised Learning. *Neural Comp*, **1**, 295–311.
- Dan, Y., Atick, J.J., & Reid, R.C. 1996. Efficient Coding Of Natural Scenes In The Lateral Geniculate-Nucleus - Experimental Test Of A Computational Theory. *Journal Of Neuroscience*, **16**(10), 3351–3362.
- DeWeese, M. 1996. Optimization principles for the neural code. *Network: Computation in Neural Systems*, **2**(May), 325–331.
- Diamond, J.M. 1996. Competition for brain space. *Nature*, **382**(6594), 756–757.
- Dong, D.W., & Atick, J.J. 1995. Temporal Decorrelation - a Theory of Lagged and Nonlagged Responses in the Lateral Geniculate-Nucleus. *Network Computation in Neural Systems*, **6**(2), 159–178.
- Field, D.J. 1994. What is the Goal of Sensory Coding? *Neural Comp*, **6**, 559–601.
- Gallant, J.L. 1996. Cortical responses to natural scenes under controlled and free viewing conditions. *Invest. Ophthalmol. Vis. Sci.*, **37**(3), 674.
- Harpur, G., & Prager, W. 1996. Development of low entropy coding in a recurrent network. *Network: computation in neural systems*, **7**(2), 277–284.
- Hateren, J.H. van. 1992. Real and optimal neural images in early vision. *Nature*, **360**, 68–70.
- Holt, G.R., Softky, W.R., Koch, C., & Douglas, R.J. 1996. Comparison of discharge variability in-vitro and in-vivo in cat visual cortex neurons. *Journal of Neurophysiology*, **75**(5), 1806–1814.
- Laughlin, S. 1981. A Simple coding Procedure Enhances Information Capacity. *Zeitschrift fur Naturforschung, Section C-Biosciences*, **36**, 910–912.

- Legéndy, C.R., & Salcman, M. 1985. Bursts and Recurrences of Bursts in Spike Trains of Spontaneously Active Striate Cortex Neurons. *Journal of Neurophysiology*, **53**(4), 926–939.
- Levy, W., & Baxter, R.A. 1996. Energy-Efficient Neural Codes. *Neural Computation*, **8**(3), 531–543.
- Linsker, R. 1988. Towards an organising principle for a layered perceptual network. *Pages 485–494 of: Touretzky, D.S. (ed), Advances in Neural Information Processing Systems*, vol. 2. Denver 1989: Morgan Kaufmann, San Mateo.
- Olshausen, B., & Field, D. 1996. Emergence of simple-cell receptive-field properties by learning a sparse code for natural images. *Nature*, **381**(6583), 607–609.
- Press, W.H., Flannery, B.P., Teukolsky, S.A., & Vetterling, W.T. 1992. *Numerical Recipes in C (second edition)*. Cambridge: Cambridge University Press.
- Rieke, F., Warlan, D., de Ruyter van Steveninck, R.R., & Bialek, W. 1997. *Spike: exploring the neural code*. Bradford Books: MIT Press.
- Rolls, E.T., & Tovee, M.J. 1995. Sparseness of the Neuronal Representation of Stimuli in the Primate Temporal Visual Cortex. *Journal of Neurophysiology*, **73**(2), 713–726.
- Sengpiel, F., Blakemore, C., Kind, P.C., & Harrad, R. 1994. Interocular suppression in the visual cortex of strabismic cats. *Journal of Neuroscience*, **14**, 6855.
- Shannon, C.E., & Weaver, W. 1949. *The mathematical theory of communication*. Illinois Press.
- Softky, W.R., & Koch, C. 1993. The highly Irregular Firing of Cortical-Cells is Inconsistent with Temporal Integrations of Random EPSPs. *Journal of Neuroscience*, **13**(1), 334–350.
- Sokoloff, L. 1989. Circulation and energy metabolism of the brain. *Pages 565–590 of: Agranoff, B.W., Albers, R.W., & Molinoff, P.B. (eds), Basic neurochemistry: molecular, Cellular, and medical aspects*, 4'th edn. Raven Press, New York.
- Srinivasan, M.V., Laughlin, S.B., & Dubs, A. 1982. Predictive coding: A fresh view of inhibition in the retina. *Proc.R.Soc.London Ser.B*, **216**, 427–459.
- Teich, M.C. 1989. Fractal Character of the Auditory Neural Spike Train. *I.E.E.E. Transactions on Biomedical Engineering*, **36**(1), 150–159.
- Teich, M.C., Turcott, R.G., & Siegel, R.M. 1996. Temporal correlations in cat striate-cortex neural spike trains. *IEEE Eng Med Biol Mag*, **15**(5), 79–87.
- Teich, M.C., Heneghan, C., Lowen, S.B., Ozaki, T., & Kaplan, E. 1997. Fractal character of the neural spike train in the visual system of the cat. *Journal of the Optical Society of America A*, **14**(3), 529–546.
- Uttley, A.M. 1973. *Information transmission in the nervous system*. Academic Press (London).

- Van Hateren, J.H. 1993. Spatiotemporal Contrast Sensitivity of Early Vision. *Vision Research*, **33(2)**, 257–267.
- Viemeister, N. 1996. Auditory Temporal Integration: What is being Accumulated. *Current Directions in Psychological Science*, **5(1)**, 28–32.
- Watt, R.J. 1987. Scanning from coarse to fine scales in the human vision system after the onset of a stimulus. *Journal of the Optical Society of America*, *4A*, 2006–2021.