



Synaptic energy efficiency in retinal processing

Benjamin T. Vincent ^{*}, Roland J. Baddeley

Laboratory of Experimental Psychology, School of Biological Sciences, University of Sussex, Falmer, Brighton BN1 9QG, UK

Received 11 April 2002; received in revised form 20 November 2002

Abstract

Recent work suggests that the visual system may represent early visual information in an energy efficient manner [Nature 381 (1996); Nature, 381 (1996) 607; Neural Comput. 3 (2001) 799; Curr. Opin. Neurobiol. 11 (2001) 475]. This paper applies the idea of energy efficient representations to understand retinal processing, and provides evidence that centre surround processing observed is efficient in terms of minimizing synaptic activity. In particular, it is shown that receptive fields at different retinal eccentricities and at different levels of noise, can be understood in terms of maximizing the transmission of visual information given a constraint on total synaptic strengths and hence energy consumption.

© 2003 Elsevier Science Ltd. All rights reserved.

Keywords: Retina; Centre surround; Energy efficiency; Synaptic activity; Retinal ganglion cells

1. Introduction

The human brain forms only 2% of body mass, but uses 20% of all blood oxygen (Rolfe & Brown, 1997). With such metabolically expensive tissue, it seems reasonable that minimising energy consumption has shaped the evolution of neural processing of the brain (Baddeley, 1996; Baddeley et al., 1997; Laughlin, 1999; Levy & Baxter, 1996). If equivalent computations could be carried out using less energy, savings from these energy efficient mechanisms could be used for other vital processes such as growth or reproduction.

The metabolic costs of neural signalling have been quantified on the organisms level (Rolfe & Brown, 1997), and on the neural level (Aiello & Bach-y-Rita, 2000; Ames, 2000; Attwell & Laughlin, 2001; Laughlin, van Steveninck, & Anderson, 1998). For this article it is useful to categorise energetic neuronal costs into three classes; costs associated with (i) firing rates (action potentials), (ii) synaptic transmission and, (iii) maintaining neurons (see Attwell & Laughlin, 2001).

Previous theoretical work can be interpreted as dealing with the first of these costs. The receptive fields of neurons in V1 can be understood as forming an energy efficient representation in terms of maximising in-

formation transmission, whilst minimising average firing rate (Baddeley, 1996). Olshausen and Field (1996), using natural images as input, maximised information transmitted whilst minimising a quantity they called sparsity and this resulted in receptive fields that are oriented bandpass, with quantitative similarities to V1 neuron receptive fields (van Hateren & van der Schaaf, 1998). One of their definitions of sparsity was the average absolute firing rate. Minimising the average firing rate whilst maximising information transmission, was found to be sufficient to account for V1 receptive fields (Baddeley, 1996). As well as accounting for the receptive fields, this ‘minimising firing rate’ proposal is consisted with observed firing rate in V1 and IT (inferiortemporal cortex) (Baddeley et al., 1997).

This previous work has concentrated on minimising energy costs associated with firing rates but as stated before, there are also costs associated with neural maintenance and synaptic activity. Here we explore the implications of representations that minimise costs associated with synaptic transmission, and compare the results to the receptive fields not in V1 but in the retina.

A common framework for modelling early visual representations is to model outputs of neurons as the result of linear filters operating on input images (Bell & Sejnowski, 1997; Harpur, 1997; Olshausen & Field, 1996; Srinivasan, Laughlin, & Dubs, 1982). In this framework, the magnitude of the filter weights is equated to the synaptic connection strengths, and the filter

^{*} Corresponding author.

E-mail address: b.t.vincent@sussex.ac.uk (B.T. Vincent).

outputs are equated to neuronal firing rates. It is within this framework that we attempt to understand the consequences of maximising information transmission whilst minimising the metabolic costs associated with synaptic transmission.

Following Bell and Sejnowski (1997), Harpur (1997) and Olshausen and Field (1996), we quantify the information transmitted in terms of the sum squared reconstruction error. If the signal (the images) and the noise are Gaussian, then minimising mean squared reconstruction error maximises the information that the outputs provide about the inputs (Baldi & Hornik, 1995). It is known that natural images are not Gaussian distributed, but we would propose this as a reasonable first approximation. Note here information refers to all variations in image intensity and does not distinguish between potentially biological relevant variations such as those generated by reflectance changes, and those of less relevance such as those generated by illumination variation. Despite this, we believe this common approximation is sufficient for the purposes of this study.

This leaves the problem of quantifying the metabolic cost associated with synaptic transmission. The arrival of an action potential at a presynaptic terminal alters the probability of vesicle release. The 'stronger' the synaptic connection the higher the release probability and the greater the number of release sites. The neurotransmitter within the vesicles result in hyperpolarisation or depolarisation of the postsynaptic potential, changing firing probability.

In relation to metabolic costs, the major contributions are those associated with action potential transmission and synaptic activation (Attwell & Laughlin, 2001), with these costs being approximately equal. Of these costs, those associated with synaptic 'strength' will be associated with glutamate release. To a first approximation, below saturation, ignoring sublinear summation at high-firing rates observed, metabolic cost will be proportional to glutamate release and hence to synaptic strength. This leads us to quantify in our simplistic model the metabolic cost of synaptic activity in terms of the sum of the absolute strengths of the weights.

This assumption that metabolic cost is a simple linear function of synaptic strength is by necessity a simplification, but we tested the robustness of our results to this assumption, by exploring costs of the form

$$\text{cost} \propto |\text{strength}|^P$$

with values of P ranging from 0.5 to 1.5. With this range of values of P we obtain qualitatively similar results. In the previous literature, people have explored constraints on weights, usually not from a metabolic perspective but simply to prevent connections from growing infinitely. Almost without exception, constraints on the sum of the

squares of the weights (e.g. $P = 2$) have been studied. This constraint results in qualitatively very different receptive fields (see later) than with the metabolically inspired constraint on the absolute values.

Using this linear framework, optimal filters were calculated under a constrained optimisation criteria on total synaptic strength for a range of convergence ratios (number of filters for a given input), energy budgets and signal-to-noise ratios (SNRs). The resulting filters were quantitatively modelled and compared to biological measurements. It is shown that this simple synaptic energy restriction is sufficient to account for many aspects of retinal processing.

2. Methods

The optimal filters were calculated with a modified symmetric error correction network (Baldi & Hornik, 1995). In its unmodified form this calculates the optimal (least squared reconstruction error) filters, for a set of M inputs which converge on N outputs. An input sample, denoted by the column vector \mathbf{x} , is forward propagated through filters \mathbf{W} to give outputs $\mathbf{y} = \mathbf{W}\mathbf{x}$. These are then back propagated to compute an error signal $\mathbf{e} = \mathbf{x} - \mathbf{W}'\mathbf{y}$. The filters are then updated by $\Delta\mathbf{w} = \lambda\mathbf{W}\mathbf{x}(\mathbf{x} - \mathbf{W}'\mathbf{W}\mathbf{x})'$, where the learning rate λ decreases over time, in order to minimise $E = \langle (1/2)\mathbf{e}^2 \rangle$.

The form described above simply minimises the sum squared reconstruction error. We modified this basic model to additionally impose a constraint on the metabolic cost of each filter. Given our cost scheme, defined earlier as $\text{cost} \propto |\text{strength}|^P$, the energetic constraint is defined as $\sum_{m=1}^M |\mathbf{W}_{n,m}|^P \leq \beta \forall n$. This constraint is enforced by a mechanism that can be thought of as a form of weight decay where all weights are moved toward zero by a small constant amount (see Goodhill & Barrow, 1996; Miller & MacKay, 1994, contrast this to more typical weight decay where weights decrease toward zero in proportion to their magnitude which implements a constraint on the sum of squared weights). Formally, after every weight update of the learning rule, the constraint $\Delta\mathbf{w}_n = -k \text{sgn}(\mathbf{w}_n) |\mathbf{w}_n|^{P-1}$ is applied iteratively to each filter \mathbf{w}_n whose cost is greater than the constraint, where k is a very small constant.

In this model, the weights used for each receptive field are also used in the reverse direction to reconstruct the image. Therefore the forward weights (\mathbf{W}) are implicitly constrained to be the transpose of the weights used for reconstruction (\mathbf{W}'). This model will only accurately reconstruct the variance of the input if the length of each weight vector is one. One consequence of length one weight vectors is that if the noise on the input is independent identically distributed and Gaussian of variance σ_{noise}^2 then every output neuron will have noise equal to σ_{noise}^2 . Since this noise will be constant, maximising the

signal (by minimising the sum squared reconstruction error) will also maximise SNR. For Gaussian signal and noise, maximising SNR maximises the transmitted information with $\text{information} = (1/2) \log_2[1 + \text{SNR}]$.

The weight matrix W was initialised as random Gaussian distributed noise with zero mean, normalised to satisfy the energetic constraint. The input data composed of 50,000 samples of size 16×16 from van Haterens' natural image dataset (van Hateren & van der Schaaf, 1998). The image intensities of each image sample were set to have zero mean but in contrast to Olshausen and Field (1996), inputs were not whitened.

The receptive fields of retinal ganglion cells have traditionally been modelled by a difference of Gaussian function. To enable quantitative comparison, the filters resulting from simulations were fitted with a difference of Gaussian model $\text{DoG}(d) = k_c \exp[-(d/r_c)^2] - k_s \times \exp[-(d/r_s)^2]$, where d is the distance from the receptive field centre (Enroth-Cugell & Robson, 1966). The six parameters correspond to the radius of the centre and the surround, r_c and r_s respectively and the sensitivity of the centre and surround k_c and k_s respectively. In addition x, y determined the location of the centre of the receptive field. This was done using a least squares criteria. The quality of fit was assessed by the correlation coefficient of the fit to the data.

The goal of these filters was to transmit information under a limited synaptic budget. The performance of a set of filters in transmitting information was measured by the inverse of the mean squared error of reconstruction $p = \langle (1/2)e^2 \rangle^{-1}$, over 50,000 image samples. To take into account the effect of different synaptic budgets, we defined an efficiency measure. Here efficiency was defined as performance p divided by total cost which is the budget per filter multiplied by the number of filters βN . This gives us efficiency $= p/\beta N$.

To calculate filters in the presence of noise, the filters were optimised to reconstruct images given noise-corrupted inputs. The only change in the calculation of the filters was the calculation of the outputs with $y = W(x + \eta)$, where η is a vector of zero-mean Gaussian distributed noise with variance set so $\sigma(x)/\sigma(\eta) = \text{SNR}$. The contrast sensitivity functions of Fig. 4 were calculated with Enroth-Cugell and Robson (1966, Eq. (9)). The parameters used in this equation were estimated from the best fitting difference of Gaussian filters fit to simulations where $M = 16^2$, $N = 64$, and $\beta = 4$.

3. Results

The benefits of synaptic efficiency would be lost if the filters were unable to encode visual information. To confirm that filters with lower energy budgets were capable of capturing visual information, performance was

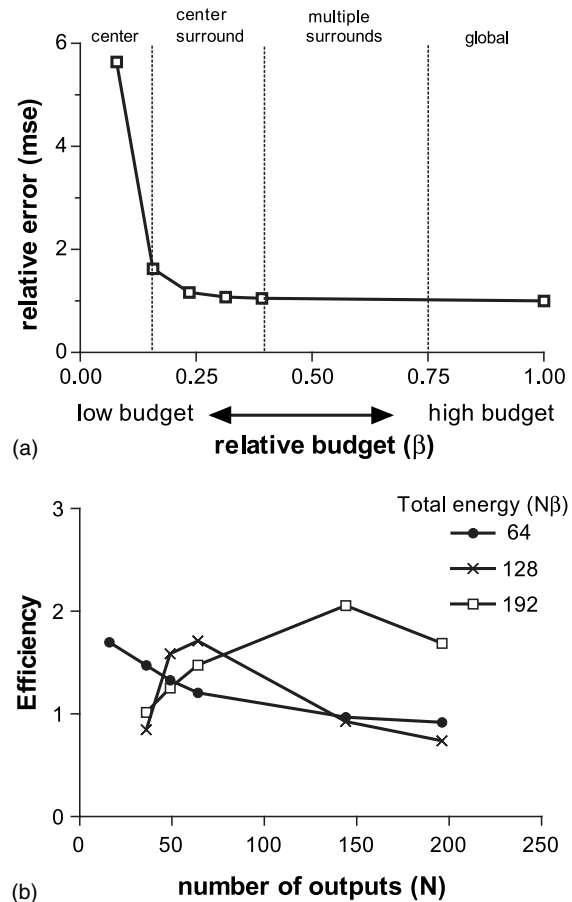


Fig. 1. Performance of filters. Performance of filters under different energy budgets is shown in (a). Both budgets and error are relative to an unconstrained set of filters. Very high performance can be maintained under low-synaptic budgets. Centre surround filters represent the best compromise between performance and synaptic budget. The model can produce filters with very low-performance decreases under very tight energy budgets. With 100 outputs, even at a 75% drop in energy budget, the error is only 17% greater than the lowest possible error. The regions labelled denote the general type of filter produced with a certain energy budget. Each curve in (b) shows the efficiencies for a given energy budget. For a low-energy budget (\bullet) it is more efficient to use a few output units. As the energy budget increases (\times) to (\square) it is more efficient to use more output units.

measured (see Section 2). Filters with an unlimited budget had the highest performance, but this high performance can be maintained under relatively lower energy budgets (see Fig. 1a). Below a certain point however, errors increase rapidly. The results suggest that the centre surround solution is the best balance between energy efficiency and performance.

A compromise between energy efficiency and reconstruction performance leads to a specific prediction of the optimal number of units to use for a given energy budget. Fig. 1b shows that when the energy budget is low, it is most efficient to use a low number of neurons. As the energy budget increases, greater efficiency will be obtained by using more output units.

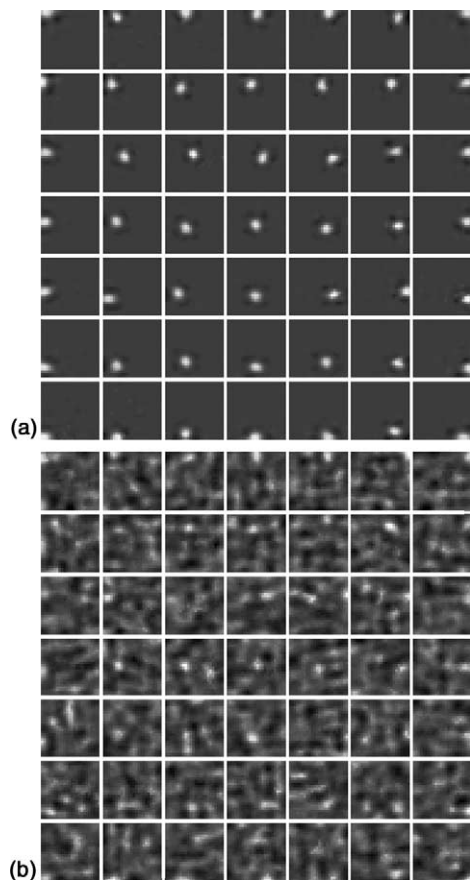


Fig. 2. Constrained and unconstrained filterbanks. A representative set 49 filters (with 16×16 input) is shown in (a) which have the energy budget $\beta = 4$. They demonstrate the localized centre surround organization. The filters that result with an unconstrained energy budget (b) are very different. They are non-localised and span the principal component subspace.

With very low-energy budgets, the filters consist of a Gaussian like centre component. With a wide range of energy budgets above this, the filters have a centre surround organisation (see Fig. 2a). At the top end of this range, filters had additional excitatory and inhibitory surrounds, an effect reported in the retina (Ikeda & Write, 1972). With unconstrained budgets filters were global (see Fig. 2b) and in this condition the filters span the same space as principal component (Harpur, 1997).

In the low-energy region which results in centre surround receptive fields the filters are well approximated by a difference of Gaussian model. This is shown both by the high correlation between the filters and their fits ($R = 0.92$), and can be seen visually in Fig. 3.

A comparison of receptive field coverage showed similarities to the mosaic organization of retinal ganglion cells. The receptive field centres were irregularly organized but were approximately evenly spaced (see Fig. 4). This can be quantified in terms of spacing ratios (Devries & Baylor, 1997), the distance between receptive field centres, divided by the sum of the centre standard

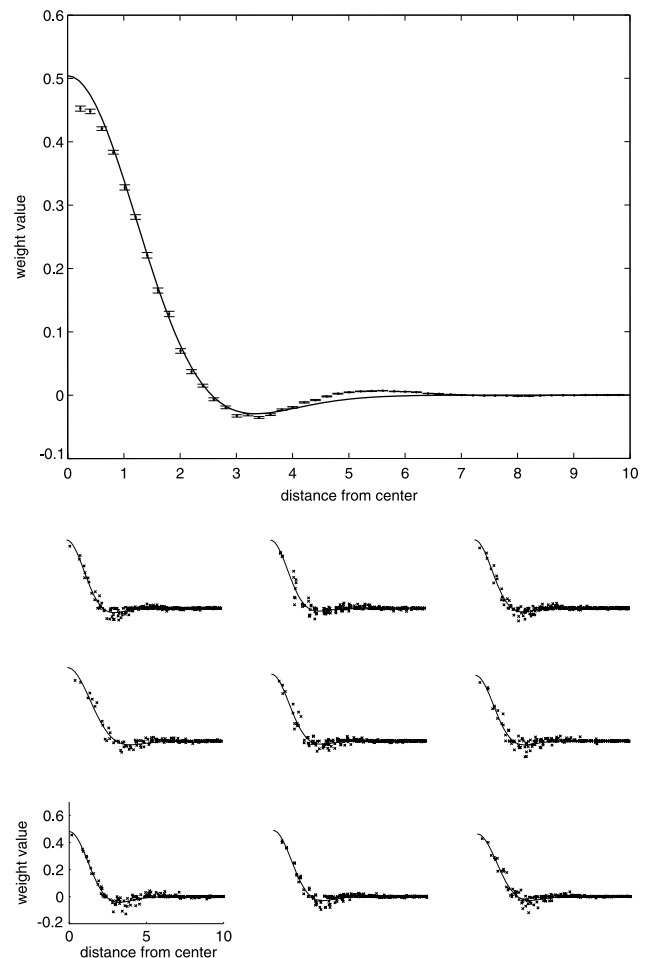


Fig. 3. Difference of Gaussian fits to the filters. The uppermost graph shows the average receptive field and the best fitting difference of Gaussian. The data points are the connection strengths (in 0.2 pixel windows) as a function of distance from the centre of the receptive field, error bars are standard error. The lower figures show fits to individual receptive fields. Data points are again connection strength as a function of distance (and hence the number of data points increase with distance from the centre), and the line is the difference of Gaussian fit. As can be seen a difference of Gaussians provides a good fit both to the average filter ($R = 0.92$) and the individual filters.

deviations ($s = d/(r_1 + r_2)$, where r_n is the standard deviation radius of filter n). For one class of ganglion cells the mean spacing ratio was 0.98 and mean 1.21 for a second. This compares favourably to a set of filters (mean = 1.15, std = 0.11) with $M = 16^2$, $N = 100$, $\beta = 5$. Similar results were found for other convergence ratios.

Filters were calculated for a variety of convergence ratios and were generally consistent with observations of the retina. The optimal filters when the number of inputs and outputs are equal ($N = M$) consisted of individual pixels which corresponds to the 1:1 connections between photoreceptors and ganglion cells (Yamada, Silveira, Perry, & Franco, 2001). These foveal filters had the highest performance and efficiency. Peripheral regions of the retina are approximated as the convergence ratio

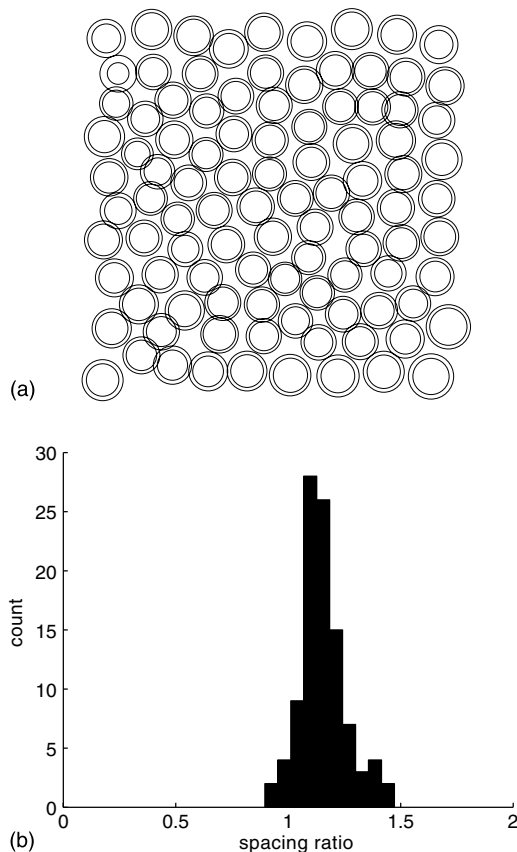


Fig. 4. Mosaic structure of filters with $M = 24$, $N = 100$, $\beta = 5$. (a) Shows a plot of receptive field locations. Each filter is represented by two circles, with the central circle corresponding to centre radius and the outer circle the surround radius. This is quantified in terms of spacing ratios (see text). (b) Shows a histogram of spacing ratios and the mean ratio is consistent with that observed for the retina.

increases. The resulting filters were coherent with biological observations of ganglion cells at different eccentricities. Filters become centre surround and increase in size in more peripheral regions (see Fig. 5).

A comparison between average X and Y retinal ganglion cells (Enroth-Cugell & Robson, 1966) and difference of Gaussian fits to the models receptive fields was made. Fig. 6 shows this comparison, fits to average Y cells were very close and fits to average X cells were reasonable. However, our model failed to produce filters

which adequately captured the high-relative surround radius of the X cell ($k_s/k_c = 4$); the fit shown has a relative surround radius of 1.95.

An important feature of retinal ganglion cells is their DC response. This describes responses to the level of illumination and is a key feature of ganglion cells (Enroth-Cugell & Robson, 1966; Linsenmeier et al., 1982). All filters calculated with a variety of energy budgets and convergence ratios had positive DC. This is of relevance to image coding because it implies that retinal ganglion cells do not do full redundancy reduction, nor respond only to edges.

An important aspect of retinal processing is how receptive fields alter over different luminance levels. As light levels decrease, the signal decreases and thus SNR decreases. One way of characterising receptive fields is by its contrast sensitivity function and for real neurons, this changes as a function of SNR. The optimal filters and corresponding contrast sensitivity function exhibit the same dependence on SNR as biological neurons (see Fig. 7; Barlow, Fitzhugh, & Kuffler, 1957; Enroth-Cugell & Robson, 1966). When there is no noise, the filters have the band pass characteristic caused by the inhibitory surround. As the level of noise increases, this disappears and gives way to a low-pass characteristic. In addition, filters become insensitive to the higher spatial frequencies due to a slightly increased centre radius. This dependence of contrast sensitivity function on SNR is not a unique prediction of this model (see Table 1), but despite the simplicity of the model, observed changes with SNR are the same as in the retina.

4. Discussion

Optimal filters that encode natural images under a simple synaptic energy constraint demonstrate many features present in retinal ganglion cells. The quantitative properties of centre surround organization at a range of eccentricities and SNRs, in addition to the positive DC and mosaic organization of retinal ganglion cells have resulted from this energetically inspired model.

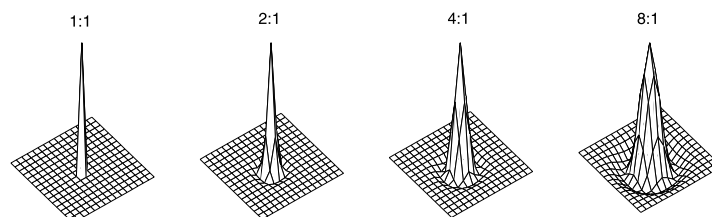


Fig. 5. This figure shows predicted receptive fields at different retinal eccentricities. Foveal regions are represented by a 1:1 convergence ratio (left) and receptive fields consist of a single excitatory connection with no surround. More peripheral regions are denoted by higher convergence ratios (right). Receptive fields consist of a surround region and become larger at more peripheral regions. Plots show the average of a difference of Gaussian fit to filters with $M = 16^2$ and $N\beta = M$.

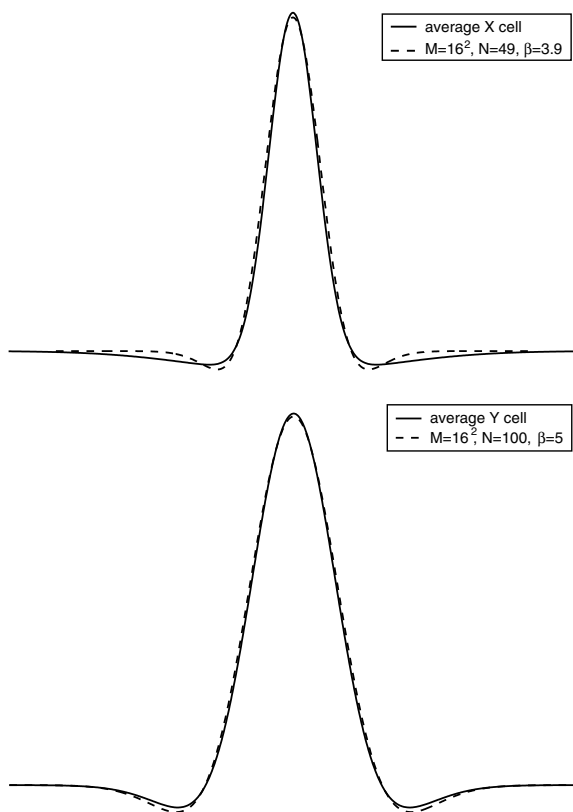


Fig. 6. Comparing model receptive fields to average X and Y retinal ganglion cells. These plots show the average X and Y cells (Linsienmeier, Frishman, Jakiela, & Enroth-Cugel, 1982) (solid lines) and difference of Gaussian fits to model filters (dotted lines) with simulation parameters shown in the legends. The top plot shows that our model achieves a reasonable fit to the average X cell, but fails to capture the full size of the surround. The bottom plot shows high similarity between the average Y cell and our model.

Why do the optimal filters have an antagonistic surround? The filters become local because the weight constraint acts to set many weights to zero and because of the autocorrelation function of the image dataset. Due

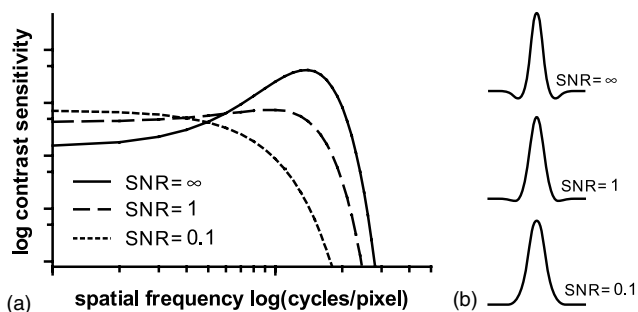


Fig. 7. The effect of noise on filters. With no noise filters consist of a centre and surround which gives rise to the band pass characteristic (—). As the SNR decreases, the surround diminished in strength and the centre increases in size slightly. This leads to decreased band pass characteristics and decreases sensitivity at high-spatial frequencies respectively (---). At very low SNRs, the surround is absent, leaving a low-pass filter (···). The units on the axes are arbitrary.

to the error minimisation, these overlapping weights tend to be as orthogonal as possible and the absolute values of all weights per input to have an equal sum. So the surround regions develop as a balance between the filters trying to be orthogonal and the weight constraint.

Two assumptions have been made with synaptic efficiency. Firstly that each connection works under the same cost scheme. It is possible that excitatory and inhibitory connections incur different costs, however, we propose that this homogeneous costing is a reasonable approximation to make for this initial investigation. The second assumption is that metabolic cost is a linear function of synaptic efficacy, which is a reasonable starting point. The parameter P was used to alter the non-linearity, similar results were achieved using values of P other than one so the centre surround solution to this problem is fairly robust as long as P is close to one.

Here we have interpreted our constraint within the metabolic framework, but it is also possible to understand it in terms of the more common “sparsity” framework (Field, 1994). By placing a constraint on the sum of absolute weight values (our energy constraint), and the second constraint on weights being length one (see Section 2), we are enforcing measure of sparsity which is more robust than other measure such as kurtosis. Therefore it is possible to understand the results as the implications of imposing a robust sparsity constraint not on the outputs of neurons, but on the weights.

Many models of retinal processing seek to explain which computational principal the retina is using. A comparison of previous models is made in Table 1 and are reviewed in Burton (1999). Predictive coding (Srinivasan et al., 1982) describes the inhibitory region (both spatial and temporal) very well over a range of SNRs. However, filters are zero DC and do not fully explain the central region. Collective coding (Tsukamoto, Smith, & Sterling, 1990) can account for a dome-shaped central region as the optimal solution to maximize SNR increase between photoreceptors and retinal ganglion cells but does not account for the surround. Atick and Redlich (1990, 1995), Haft and van Hemmen (1998) and van Hateren (1993a, 1993b) account for the centre and surround over SNRs, but explicitly impose locality so aspects such as the sampling mosaic are not addressable. No single alternative model fully explains all aspects of retinal processing (but this may well be because these aspects were not addressed rather than being fundamentally inexplicable).

The concept of efficient neural coding is not new, however, the formulation in terms of cost of neurons, action potentials, synaptic activity and wiring is relatively recent (Baddeley, 1996; Cherniak, 1995; Laughlin, 2001; Laughlin et al., 1998). The approach used in this paper is not specific to the spatial processing of the retina, and could be extended to account for spatio-temporal processing in the retina and to other sensory

Table 1

A comparison between models of spatial receptive fields of retinal ganglion cells

	Synaptic efficiency	Predictive coding	Collective coding	van Hateren	Infomax	Atick and Redlich
Centre and surround	y	Surround	Centre	y	y	y
Positive DC	y	n	n/a	y	y	y
Mosaic	Emerges	n/a	n/a	n	Imposed	n
Local filters	y	n/a	n/a	Imposed	Imposed	Imposed
SNR	y	y	y	y	y	y

modalities such as spectro-temporal processing in the auditory system. Lots of work has explored sparse coding in the visual cortex (Hoyer & Hyvarinen, 2000; Hyvarinen & Hoyer, 2001a, 2001b, 2001c; Olshausen & Field, 1996; van Hateren & Ruderman, 1998; van Hateren & van der Schaaf, 1998). These results should be interpretable in terms of energy efficient coding.

It is argued here that V1 receptive fields could be interpreted in terms of a constraint on mean firing rates (Baddeley, 1996; Levy & Baxter, 1996), and retinal codes in terms of a synaptic constraint. Why should this be? The answer may lie in the number of neurons in the retina and V1. Because there are only ~1 million retinal ganglion cells per eye, codes that constrain firing rates may yield few savings and hinder the accuracy of the ~100:1 compressed data. This leaves connection costs for potential savings. In the cortex however, sheer neuron numbers means massive savings could be made by having only a small proportion of neurons strongly active, in addition this massive divergence suggests that coding performance would not be hindered.

In conclusion this paper presents optimal filters designed to capture visual information with efficient use of synaptic resources. These filters have strong similarities to the receptive fields of X and Y ganglion cells. It was found that a constraint on synaptic efficiency was able to account for many properties of retinal ganglion cells, such as the centre surround mechanism, and how this alters over retinal eccentricity, and SNRs. This simple synaptic efficiency combined with cortical sparsification provides a common energy efficiency approach that can explain many aspects of retinal and cortical level processing in the visual system.

Acknowledgements

The authors wish to thank Ben Tatler, Rebecca Shortt and Daniel Osorio for useful suggestions during the preparation of this paper and the MRC grant G78-6812.

References

- Aiello, G., & Bach-y-Rita, P. (2000). The cost of an action potential. *Journal of Neuroscience Methods*, 103, 145–149.

- Ames, A. (2000). Cns energy metabolism as related to function. *Brain Research Reviews*, 34, 42–68.
- Atick, J. J., & Redlich, A. (1990). Towards a theory of early visual processing. *Neural Computation*, 2, 308–320.
- Atick, J. J., & Redlich, A. (1995). A convergent algorithm for sensory receptive field development. *Neural Computation*, 5, 45–60.
- Attwell, D., & Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow and Metabolism*, 21(10), 1775–1783.
- Baddeley, R. (1996). Visual perception, an efficient code in v1? *Nature*, 381.
- Baddeley, R., Abbott, L. F. A., Booth, M. C., Sengpiel, F., Freeman, T., Wakeman, E. A., & Rolls, E. T. (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society of London Series B—Biological Sciences*, 264.
- Baldi, P. F., & Hornik, K. (1995). Learning in linear neural networks: A survey. *IEEE Transactions on Neural Networks*, 6(4), 837–858.
- Barlow, H. B., Fitzhugh, R., & Kuffler, S. W. (1957). Change of organisation in the receptive fields of the cats retina during dark adaptation. *Journal of Physiology*, 137, 338–354.
- Bell, A. J., & Sejnowski, T. J. (1997). The ‘independent components’ of natural scenes are edge filters. *Vision Research*, 37(23), 3327–3338.
- Burton, B. G. (1999). *Problems and solutions in early visual processing. Information theory and the brain*. Cambridge: Cambridge University Press.
- Cherniak, C. (1995). Neural component placement. *TINS*, 18(12), 522–527.
- Devries, S. H., & Baylor, D. A. (1997). Mosaic arrangement of ganglion cell receptive fields in rabbit retina. *Journal of Neurophysiology*, 74(4), 2048–2068.
- Enroth-Cugell, C., & Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *Journal of Physiology*, 187, 517–552.
- Field, D. J. (1994). What is the goal of sensory coding? *Neural Computation*, 6, 559–601.
- Goodhill, G., & Barrow, H. (1996). The role of weight normalization in competitive learning. *Neural Computation*, 6, 225–269.
- Haft, M., & van Hemmen, J. L. (1998). Theory and implementation of infomax filters for the retina. *Network*, 9(1), 39–71.
- Harpur, G. F. (1997). *Low entropy coding with unsupervised neural networks*. Ph.D. thesis, Oxford University, Engineering Department.
- Hoyer, P. O., & Hyvarinen, A. (2000). Independent component analysis applied to feature extraction from colour and stereo images. *Network*, 11(3), 191–210.
- Hyvarinen, A., & Hoyer, P. O. (2001a). A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. *Vision Research*, 41(18), 2413.
- Hyvarinen, A., & Hoyer, P. O. (2001b). Topographic independent component analysis. *Neural Computation*, 13(7), 1527–1558.
- Hyvarinen, A., & Hoyer, P. O. (2001c). Emergence of phase- and shift-invariant features by decomposition of natural images into independent feature sub-spaces. *Neural Computation*, 12(7), 1705–1720.

- Ikeda, H., & Write, M. J. (1972). The outer disinhibitory surround of the retinal ganglion receptive field. *Journal of Physiology*, 226, 511–544.
- Laughlin, S. B. (1999). *Coding efficiency and the metabolic cost of sensory and neural information. Information theory and the brain*. Cambridge: Cambridge University Press.
- Laughlin, S. B. (2001). Energy as a constraint on the coding and processing of sensory information. *Current Opinion in Neurobiology*, 11(4), 475–480.
- Laughlin, S. B., van Steveninck, R., & Anderson, J. C. (1998). The metabolic cost of neural information. *Nature Neuroscience*, 1(1), 36–41.
- Levy, W. B., & Baxter, R. A. (1996). Energy efficient neural codes. *Neural Computation*, 8(3), 531–543.
- Linsenmeier, R. A., Frishman, L. J., Jakiela, H. G., & Enroth-Cugel, C. (1982). Receptive field properties of x and y cells in the cat retina derived from contrast sensitivity measurements. *Vision Research*, 22(9), 1173–1183.
- Miller, K., & MacKay, D. (1994). The role of constraints in hebbian learning. *Neural Computation*, 6, 100–126.
- Olshausen, B., & Field, D. (1996). Emergence of simple-cell receptive fields properties by learning a sparse code for natural images. *Nature*, 381(6583), 607–609.
- Rolfe, D. F., & Brown, G. C. (1997). Cellular energy utilization and molecular origin of standard metabolic rate in mammals. *Physiological Reviews*, 77(3), 731–758.
- Srinivasan, M. V., Laughlin, S. B., & Dubs, A. (1982). Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London Series B—Biological Sciences*, 216(1205), 427–459.
- Tsukamoto, Y., Smith, R. G., & Sterling, P. (1990). ‘Collective coding’ of correlated cone signals in the retinal ganglion cell. *Proceedings of the National Academy of Sciences of the United States of America*, 87(5), 1860–1864.
- van Hateren, J. H. (1993a). Spatiotemporal contrast sensitivity of early vision. *Vision Research*, 33(2), 257–267.
- van Hateren, J. H. (1993b). Three modes of spatiotemporal preprocessing by eyes. *Journal of Comparative Physiology A—Sensory Neural and Behavioral Physiology*, 172(5), 583–591.
- van Hateren, J. H., & Ruderman, D. L. (1998). Independent component analysis of natural image sequences yields spatiotemporal filters similar to simple cells in primary visual cortex. *Proceedings of the Royal Society of London Series B—Biological Sciences*, 265(1412), 2315–2320.
- van Hateren, J. H., & van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society of London Series B—Biological Sciences*, 265(11394), 359–366.
- Yamada, E. S., Silveira, L. C., Perry, V. H., & Franco, S. E. C. (2001). M and p retinal ganglion cells of the owl monkey: morphology, size and photoreceptor convergence. *Vision Research*, 41(2), 119–131.