1992

Backpropagation Learns Marr's Operator

Anupam Joshi

Chia-Hoang Lee

Report Number:
92-059

Joshi, Anupam and Lee, Chia-Hoang, "Backpropagation Learns Marr's Operator" (1992). *Department of Computer Science Technical Reports*. Paper 980.
<https://docs.lib.purdue.edu/cstech/980>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries.
Please contact epubs@purdue.edu for additional information.

BACKPROPAGATION LEARNS MARR'S OPERATOR

**Anupam Joshi
Chia-Hoang Lee**

**CSD-TR-92-059
September 16, 1992
(Revised 3/93)**

Backpropagation learns Marr's operator

Anupam Joshi

Department of Computer Science

Purdue University

West Lafayette, IN 47907 , USA

Chia-Hoang Lee

Department of Computer and Information Science

National Chiao-Tung University

Hsinchu, Taiwan 30050, ROC

Correspondence should be addressed to Anupam Joshi at the above address.

Phone: (317)-494-0361

Fax: (317)-494-0739

Email: joshi@cs.purdue.edu

Abstract

This paper describes a neural network model of the retinal responses to stimuli, whose architecture is inspired by neurophysiological data. Suitable assumptions are identified which enable the development of a simple model for an individual X-type ganglion cell using backpropagation. This is then used to make a model of retinal processing. We present here our model of the individual ganglion cells and the underlying assumptions. We show that backpropagation leads to a model which is similar to the mathematical descriptions of retinal processing advanced by Marr. We present the results obtained when our model is used to simulate the effect of retinal processing on images. Empirical results about the speedups obtained when this model is implemented on parallel architectures are also reported.

1 Introduction

The “Early Vision” system, of which the retina is an important part, is the first stage in the visual pathway. This system acts on the raw data (intensity values) coming in from the environment and converts it into representational forms that are then used by the higher stages. An understanding of this system and an ability to model its processing can thus be of great value in trying to find out the kind of information that the latter stages of visual processing use. Works, amongst others, like (Marr 1982, Marr & Hildreth, 1979), (Enroth-Cugell & Robson, 1966), (Daugman 1983, 1988) and (Oğuztöreli 1988, 1989) postulate various mathematical models of the retinal processing.

In this effort our aim was to develop a neural network model of the retina, using a fairly simple learning paradigm, namely backpropagation. Backpropagation has been extensively used for various applications. Efforts like (Lehky et. al., 1990) amongst others have demonstrated the use of this paradigm in modelling the binocular depth perception. Simple learning methods like competition have also been used by (Barrow, 1987) to model the receptive fields in Area 17 of the cortex. The contention that evolution has “fine tuned” the neural architectures in biological systems to the tasks they perform, is self evident. It is thus reasonable to expect that an Artificial Neural System that bases itself on its biological counterpart will perform close to optimally. In keeping with this philosophy, our model attempts to mimic the retinal architecture.

We describe a model of the retinal cells (Joshi & Lee 1991) that was built and trained using backpropagation so that the cells in the output layer would have the same centre surround response and the ability to ignore the effect of the ambient light that the X-ganglion

cells in the retina have. We then analytically demonstrate that the connection strengths between neurons to which backpropagation converges causes the model to approximate the Laplacian of the Gaussian. This function (Marr 1982, Marr & Hildreth 1979) is commonly accepted (Watt 1988, Richter & Ullman 1986) as a good approximation describing the processing in the retina. Finally, we present the results of using this model to simulate the retinal processing on grayscale images as further support of the model.

Allow us to point out that we do not seek to portray our model as an exact description, in entirety, of the retinal processing. However, Marr's operator is a commonly accepted mathematical description of the retinal processing, and in as much as our model learns it, it captures the processing of its biological counterpart.

2 Modelling the Receptive Fields of the Retinal Ganglion Cells

The retina is a complicated organ. It consists of five main type of neural cells, namely the photoreceptors, the horizontals, the bipolars, the amacrine and the ganglions (Fig.1). The photoreceptors receive input from the external world and the output of the ganglion cells is regarded as the output of the retina. While a "front to back" information flow can thus be identified from the photoreceptors to the ganglions, there are feedback paths, and it is far from being a layered, feedforward system. There is a net fan in of connections going from the photoreceptors (about 1.26×10^8) to the ganglions (about 10^6). The photoreceptors are signal transducers, converting the light energy to electrical/chemical energy. Photoreceptors are known to hyperpolarise in response to light. They are connected

to bipolar cells, some directly and some through horizontal cells. The horizontal cells too hyperpolarise in response to light. We refer here to the L-type horizontal cells, not the C-type cells which show biphasic or triphasic response, and are involved in colour vision. The bipolar cells come in two varieties, those that hyperpolarise in response to light, and those that depolarise. Bipolars in general have an antagonistic surround region, which is mediated by horizontal cells. Bipolar cells are connected to ganglion cells, as are the amacrine cells. For a detailed anatomical and physiological description of the retina, the reader is referred to (Hubel 1988, Dowling 1987, Buser & Imbert 1992).

To model the retina using Neural Networks, we need to study its output, as represented by the activities of the ganglion cells in response to the input stimuli on the retinal photoreceptors. In other words, we need to study the receptive fields of the ganglion cells. Neurobiologists have shown that these receptive fields are of *Centre Surround* type, *On Centre*, *Off Surround* and *Off Centre*, *On Surround*. We will confine our discussion in this paper to those ganglion cells which show a sustained response to the stimuli, namely the X-type ganglion cells. We are only interested in this work to model the retinal responses to static stimuli. The ganglion cell fires only if there is a difference in the illumination in its centre and surround areas. Another important feature of the ganglion cell is that its response is largely independent of the ambient light. Thus when we talk of On or Off regions, we do so not in terms of the actual illumination but its relative value. An On centre, Off surround ganglion cell will fire if the centre region is 2% brighter than the surround region. This feature might contribute to what is known in psychophysics as “lightness constancy”, which refers to the fact that we are able to perceive objects to be the same regardless of the level of the ambient illumination (Wallach 1948).

In order to model the ganglion cell receptive fields, the obvious approach is to use neural network models which allow arbitrary connectivity between their component “neurons”. While such models are undoubtedly better able to capture the actual retinal architecture, they are computationally intensive and involve complicated learning strategies. It is thus worthwhile to see if suitable assumptions can be identified which would enable the construction of a model using layered, feedforward networks. Such networks are not as computationally intense, and have well defined, well tested and simple learning rules. We identified the following assumptions as being suitable to this end:

- We ignore all but the primary, feedforward, retinal circuitry. Secondary retinal paths, which involve feedback, are not considered.
- We ignore the function of the amacrine cells.

These assumptions reduce the retinal circuitry to a three layer, feed forward only network. Since such systems are amenable to training by gradient descent rules, we used backpropagation to train our system.

Both these assumptions can be supported by anatomical and physiological data. First, consider the assumption regarding feedback paths. The feedback paths essentially involve inhibitory connections from horizontal cells to cones, first demonstrated in turtles (Baylor et.al., 1971). This enables a rapid depolarisation of the membrane after the initial hyperpolarising response. It was demonstrated in (Kleinschmidt & Dowling, 1975) that blocking these feedback paths did not alter the sensitivity or response to dim flashes of the photoreceptor, but altered its response to bright flashes, making recovery from the hyperpolarisation slower. We can thus expect that ignoring these connections would degrade the response of

our model in presence of bright ambient light. The horizontal cells mediate the surround response of bipolar cells. The bipolar cells connect directly to photoreceptors that constitute their centre. The receptors constituting the antagonistic surround though connect to the horizontal cells, which in turn connect to the bipolar. The horizontal cell in effect can be seen as integrating the response of the surround receptors and giving it to the bipolar cell. Since we have chosen to ignore the possible feedback effects of the horizontal, we can subsume the horizontal cell and connect the bipolar directly to receptors not just constituting its centre but surround as well. Let us demonstrate the veracity of this assertion for the simple case where the output of each “neuron” in our model is linearly related to the weighted sum of the inputs. Note that this is a first order approximation to the sigmoidal function that actually relates the weighted input sum and the output. Consider Fig 2. We wish to show that given a weight vector $L = (l_1, l_2, l_3 \dots l_n)$ reflecting the weights of the connections from the surround receptors to the horizontal cell, and a weight m associated with the horizontal to bipolar connection, we can find a weight vector $W = (w_1, w_2, w_3, \dots, w_n)$ reflecting the weights of the direct connections from surround photoreceptors to the bipolar such that the input to the bipolar cell remains the same. In other words,

$$\forall L, \exists W \text{ s.t. } m(k \sum_i l_i x_i) = \sum_i w_i x_i$$

where k is a constant of proportionality. Note that this is trivially true by putting $w_i = mkl_i$. We may point out here that the connections between the photoreceptors and bipolar are now a combination of the receptor horizontal and horizontal bipolar connections.

Next, consider our other assumption, namely disregarding of the role of the amacrine cells. Many studies, including (West & Dowling 1972, West 1976), have shown that in

terms of synaptic input, ganglion cells can be broadly divided into two types, those that receive input primarily from amacrine cells, and those that receive input primarily from bipolar cells. It has been noted that the X-type ganglion cell, which produce sustained contrast sensitive response are of the latter type and receive on an average more than 70 percent bipolar input in the cat retina (Kolb 1979). Moreover, the intensity response characteristics of this type of cell seem in agreement with those of bipolar cells as evidenced by intracellular recordings in the mudpuppy retina (Thibos & Werblin, 1978). Further evidence in this connection comes from receptive field sizes, as well as from the experiments with neurotransmitters of the inner plexiform layer, the layer that contains amacrine cells. It was shown (Caldwell & Daw, 1978) that in rabbits, antagonists to inner plexiform layer neurotransmitters do not fundamentally change the receptive field organisation of the X ganglion cells.

Disregarding the role of amacrines should not thus, in theory, prevent a model from closely approximating the function of X ganglions. We point out here that the above evidences should not be taken to mean that X type ganglion cells do not receive any amacrine cell input, for there is clear evidence that they do. Rather, it demonstrates that the response characteristics of such cells are based on those of bipolars. Hubel (Hubel, 1988) aptly sums up the situation by saying that “ *Amacrine cells are not known to be involved in the centre surround organisation of ganglion cell receptive fields, although we cannot rule out the possibility.* ”.

We will present in later sections the results of our simulations, which are based on these simplifying assumptions, that will provide further support of their veracity.

2.1 Centre Surround Field Structure

The first step in building the model was to come up with a system that would develop centre-surround response for binary valued inputs of zero representing darkness and one representing light. Our model used five cells in the receptor layer, and one cell each in the bipolar and ganglion layers. The five cells in the first layer were all connected to the cell in the second layer, which was in turn connected to the cell in the third layer. The middle cell of the receptor layer served as the centre, and the two cells each on either side of it were taken to be the surround. This structure was trained on patterns of centre surround stimuli and the expected responses. The input stimuli were values of 0 and 1 which encoded the ON/OFF condition at the given photoreceptor. The expected output was taken to be 1 if the input was of the appropriate centre surround type, and zero otherwise. As is conventional with backpropagation, the weights were initialised to small random values before training.

It was observed that when the system had learned the input output pairs, the patterns of inhibitory and excitatory connections (Fig. 8) as reflected in the weights were such that the model was a close analogue of the biological system. So the model did more than just learn the Input-Output relation, it also captured some elements of the neuronal circuit in the retina. To understand this better, let us briefly examine the pattern of synaptic connections in the retina here. For greater details, the reader is referred to (Dowling 1979, Dowling 1987, Hubel 1988, Buser & Imbert, 1992). The connections between horizontal cells and receptors are excitatory, since horizontal cells, like receptors, hyperpolarise in presence of light. The bipolar cells are of two types. Some (the B_H cells) hyperpolarise in response to light, and the photoreceptors have excitatory synapses to such cells. Others (the B_D

cells) depolarise, and the receptors have inhibitory synapses on such cells. The connection between the horizontal and bipolar cells is such as to make the surround antagonistic to the centre. The connections between bipolars and ganglions are in general excitatory. The B_D cells feed to on centre ganglion cells, and the B_H cells feed the off centre ganglions.

In our simple model, however, there are no hyperpolarising cells. Notice, however, that the pattern of connections learnt is still such that the surround receptor-bipolar connections (which now present the cumulative effect of the receptor-horizontal-bipolar pathway) are antagonistic to the centre connection. Moreover, note that the pattern of excitatory-inhibitory connections learnt by our model for an on centre ganglion cell is the same as for an actual on-centre ganglion cell, except that the second to third layer connection is inhibitory instead of excitatory. This is because the connections from the first to second layer, which would have produced a B_H type of second layer cell if the receptors were hyperpolarising, now produces a B_D type cell. Accordingly, the second to third layer connection is the reverse of the actual biological connection. An analogous situation is observed for the model off centre cell as well. The initial weight set had to be heavily predisposed for the learning to result in any other feasible set of final weights.

2.2 Ignoring the Ambient

Another important facet of retinal processing is the ignoring of the ambient lighting. The on and off conditions of the centre and surround areas are defined not in terms of absolute amount of light, but in terms of contrast. To achieve this, the input patterns were made to take continuous values between zero and one rather than just binary values. A difference of 0.4 was deemed sufficient to make a distinction between on and off. The output continued

to be 0 or 1 depending on the input.

When our original 5 - 1 - 1 structure was trained on these new patterns, it correctly learned the desired behaviour. The training instances contained some data with differences of 0.4 between the centre and surround. Upon completion of training, the network was found to be capable of making distinctions when the on and off regions differed by more than 0.3. However, when the overall illumination was high, the performance of the network degraded and it required differences of at least 0.5 to be able to make an on off distinction. It may be noted that even humans, for instance, fail to notice minor differences in illumination if the overall illumination is too high. This probably occurs because at very high illumination levels, the photoreceptors saturate. Since neurons in the backpropagation paradigm have a similar saturation of the output level due to the sigmoidal squashing function, it is no surprise that this behaviour is observed in our model. However this degradation was observed earlier than expected. Also while the human retina can distinguish intensity differences as low as 2%, our model needed much higher differences (30%) as shown in (Joshi & Lee, 1991). This happens because in our model for the ganglion cells, feedback and intra layer connections have been ignored. As we have pointed out earlier, it has been experimentally shown that ignoring these connections leads to a change in the response of the receptors at high intensity levels. Such behaviour then is an expected outcome of our simplifying assumption. In other words, this assumption is likely to affect only the overall level of the contrast sensitivity of our model at high ambient illumination.

In Fig. 9, we plot the response of an on centre ganglion cell vs. the log of the intensity of the illumination at the centre. The surround region has a constant illumination. Notice that the responses of our model are very similar to the actual ganglion cell responses that

have been recorded intracellularly (Thibos & Werblin, 1978). Further, as the intensity of the annular surround region is increased, the curve still maintains its shape but as expected, shifts to the right.

Interestingly, the model correctly responded if the surround region was not uniformly illuminated, but was overall brighter than the centre region (or *vice versa*). This is an instance of generalisation that backpropagation learning often exhibits.

3 Analysis of the Model

In this section we shall present a mathematical analysis of our 5 - 1 - 1 model and show that backpropagation has learned the Marr operator. A few preliminaries, however, need to be noted. Firstly, this analysis disregards the “bias” term in computing the state of a neuron in backpropagation. We have run experiments in which the bias term was set to zero and still obtained the correct results, and so this is a reasonable assumption to make. Secondly, it needs to be recognised that the activity of a neuron is merely a squashed version of its state, *i.e.* the function $activity = \frac{1}{1+e^{-state}}$ simply forces the activity, that can potentially span the $-\infty$ to ∞ range, to lie between 0 and 1. At the peril of ignoring the nonlinearities introduced by the sigmoidal function, it can be argued along similar lines that the activity of the ganglion cell can be viewed as merely a squashed version of the state of the bipolar cell. In terms of actual “processing”, thus, we need only to study how the state of the bipolar cell is obtained from the inputs.

A commonly accepted mathematical description of the retinal processing is due to Marr (Marr 1982, Marr & Hildreth 1980) who hypothesised that the output of the retinal stage

can be expressed as the convolution of an Isotropic operator with the Image function.

Limiting ourselves to the one dimensional case, this means that the output is

$$F(x) * I(x) \quad (1)$$

where $I(x)$ is the image intensity function. $F(x)$ is an isotropic operator, which Marr argues has of the form $\nabla^2 G(x)$ where $G(x)$ is a gaussian function with zero mean, and ∇^2 denotes the Laplacian. The analysis that follows is also, *mutatis mutandis*, true for a 2 dimensional system.

We will now show that our system “learns” this function. Consider convolution . Using the *ab initio* definition of convolution, we can write :

$$F(x) * I(x) = \int_{-\infty}^{\infty} F(\delta) I(x - \delta) d\delta. \quad (2)$$

since $F(x)$, which is $\nabla^2 G(x)$, is symmetric around 0 , we can rewrite this as

$$F(x) * I(x) = \int_{-\infty}^{\infty} F(-\delta) I(x - \delta) d\delta. \quad (3)$$

This integral can be discretised into a sum, with a stepsize of 1, as

$$\sum_{\delta = -\infty}^{\infty} F(-\delta) I(x - \delta). \quad (4)$$

Now, since we consider the receptive field to be finite in size, we can change the limits of summation to the limits of the receptive field size, which in our case extends from -2 to 2, for $F(x)$ will be zero beyond this.

$$\sum_{\delta = -2}^2 F(-\delta) I(x - \delta). \quad (5)$$

This equation is identical to the state equation of the bipolar neuron,

$$\sum_{\delta=-2}^2 W_{-\delta} I_{x-\delta}. \quad (6)$$

where $W_{-\delta}$ would be the weight between the bipolar neuron and the receptor neuron δ away from the centre.

We have shown earlier that the output of the ganglion cell, by virtue of our simplifications, can be seen as merely a squashed version of the state of the bipolar. Hence, if the weights to which backpropagation converged should take the form of $\nabla^2 G$, the output of our model would be $\nabla^2 G * I$. In other words, the model would have learned the Marr operator by training on some samples of centre surround data. The experimental results show that the weights do, indeed approximate this function. In Figures 3 and 4, we plot the actual weight values obtained by the learning algorithm for different input data. In both the figures, we also plot an idealised $\nabla^2 G$ function, which best fits the data points. In other words, the variance (σ) of the function is chosen so as to minimise the least square error at the points for which we have the weight values, *i.e.*

$$\min_{\sigma} \sum_{\delta=-2}^2 (W_{\delta} - \nabla^2 G(\delta))^2 \quad (7)$$

4 Simulation of the Retinal Processing

In the preceding sections, we have postulated and analysed a system to model the actions of a ganglion cell. Since the output of the ganglion cells represent the retinal output, we can use this model to obtain the response of the retina when a grayscale image is presented.

In the photoreceptor layer, we have one neuron for each pixel. So for an $n \times n$ image, we need a $n \times n$ input layer. While the model we developed for the ganglion cell assumed

a one dimensional input, the images here are two dimensional. So we convert the model by essentially replicating the connections. A node $[i,j]$ in the second layer gets input from nodes $[i,j]$, $[i-1,j]$, $[i-2,j]$, $[i+1,j]$, $[i+2,j]$, $[i,j-1]$, $[i,j-2]$, $[i,j+1]$, $[i,j+2]$ in the first layer. The weight for the connection from node $[i+/- m,j]$ is used for the connection from node $[i,j+/- m]$ too (m can take the value 1 or 2). *Ipso facto*, nodes in the second layer whose i or j coordinates are either less than 3 or more than 254 will not have some of their inputs defined. reducing it to $(n-4) \times (n-4)$ in size. Since node $[i,j]$ in layer three gets input only from node $[i,j]$ in layer two, layer three is $(n-4) \times (n-4)$ also. Note that we are only approximating the circular symmetry of the receptive field by confining our connections to the NS and EW axes. Moreover, we are assuming a uniform density of photoreceptors and unchanging receptive field sizes across the whole visual field.

Several simulations were done with the model using 256×256 images with 256 gray levels as input. It was observed in general that the output contained only the outlines of the object(s) in the image, with everything else blacked out. The effect was noticeably clearer when the input image had good contrast and was less prominent for images with poorer contrast. Figs 5& 6 serve to illustrate this phenomenon. While Fig.5 clearly shows the outlines marked, in Fig.6, we can observe the faint outlines even though the background and the interior are not completely blacked out. This occurs because our simplifications lead to a lower contrast sensitivity as explained in section 2.2. As a comparison to the kind of “edge detecting” that our retinal model seems to be doing, Fig.7 shows the output of a sobel edge detector on the same images as Fig. 5& 6.

It is widely believed that the retina sends to higher visual processing stages only information regarding the changes in the input image. As early as 1906, Mach(Mach 1906) showed

the visual system is sensitive to d^2I/dx^2 and d^2I/dy^2 . This information seems very much in line with our simulation results, lending further credence to our neural model. Note that the outline marks the area in the input image where there was an intensity change. So the model appears to be filtering out all the constancies in the image, and producing as output the contours across which an intensity difference exists. The higher visual stages can then, presumably, reconstruct the image by “filling in” the constant information.(Grossberg & Wyse 1991).

5 Exploiting Parallelism

One of the “*ground truths*” about the human early visual system is its extremely fast processing, the times involved being of the order of tens of milliseconds(Oğuztöreli 1988). Any artificial system must then aim for similar, if not better processing speeds. Our model is able to match up to its biological counterpart in this aspect too.

Once the connection strengths between various neurons are determined, the model becomes particularly suited for parallel implementations. This is because each neuron does a local computation based on the activity of its neighbours in the preceding layer. The parallelism in this case is extremely fine grained, and this is a task particularly suited to SIMD machines. In order to exploit this parallelism, we ported our simulator to a MasPar¹. As a comparison, we also ported the programme to a CRAY YMP 4/446, which is a vector machine. Table 1 shows the processing times on these machines, as well as the speedups, compared to the time required to run the programme on a Sun Sparc IPC. Since the MasPar

¹For use of this facility, the authors wish to thank the PPL of the EE Dept at Purdue, which is supported by NSF award no. 9015696-CDA

has only 16384 processors, configured in an array of 128×128 , the timing measurements here refer to the processing of an 128×128 image. Images of greater size can be run on the MasPar using a simple divide and conquer technique, although the recombination steps involved would slow down the processing slightly. Also, the timing here does not include times for Input/Output. We can clearly see here that on machines like the MasPar that allow its fine grained parallelism to be exploited, the model does its computations in the order of milliseconds.

6 Conclusions

In this paper, the authors have introduced a model of retinal ganglion cells and used it to simulate the information processing in the retina. The model is an approximation of the human retina and is based, architecturally, on it. A mathematical analysis of the model has also been presented.

Our model is able to duplicate the behaviour of ganglion cells in terms of their centre surround field structure and their ability to ignore the ambient illumination conditions. The model was developed by examining the neural circuitry in the retina, and using it as a guide to establish the neural architecture of our model. Suitable assumptions which enabled us to reduce the retinal circuitry to a three layer, feedforward only network have also been identified. This is unlike some of the pioneering works in this area like (Linsker 1988), which start with arbitrary multilayer networks. While our simplifying assumptions do, to some extent, adversely affect the ability of the model to notice small contrast changes, they lead to a much simpler model which can be trained using conventional backpropagation

algorithms. The output of the ganglion layer of the retinal model helps us to visualise and understand the kind of information that the retina passes on to the higher cortical stages.

The connection strengths between neurons which are established by the learning process of backpropagation, we observe, are quite similar to those actually found in the retina, at least in terms of the pattern of excitatory and inhibitory connections. Also, we demonstrate that the operation that our model does on the input, as dictated by the learned weights, is an approximation to the Laplacian of the Gaussian. This function, according to the works of David Marr and others is a mathematical description of the functionality of the retina. It is interesting to observe that evolution and gradient descent, both of which are trying to find an optimal system to perform a certain function, lead in this case at least to very similar systems.

The results we have obtained by simulating the processing of the retina with grayscale images as input seem to be in conformity with the opinion that many other researchers hold regarding the functionality of the retina. Our model can thus serve as a “front end” for other systems that model higher visual stages like (Braham & Hamblen, 1990). Also, while useful in themselves in studying the processing done by the retina, the outputs of the model can serve as a segmentation of the input image. The authors are currently developing a method to use these segmentations to establish correspondence in image sequences.

We would like to add that our assumptions about ignoring the connections that involve feedback and the role of amacrine cells of the sustained type do detract from this model. Further research work in neural net models will undoubtedly enable us to model arbitrarily connected networks more efficiently in the future. Given such new learning paradigms, it would be extremely instructive to do away with the simplifying assumptions made here. We

would also like to stress once again that our modelling, while trying to maintain a retina like architecture, essentially tries to learn the input output behaviour of the retina. It is not an exact model in terms of actual intracellular connections. However it does capture the basic architecture in the retina responsible for the X type ganglion cells, and is a good compromise between simplicity and accuracy.

In conclusion then, we propose a biologically motivated neural network model of the retina that is able to closely duplicate the kind of information processing that the human retina does.

7 Acknowledgements

The authors would like to thank the anonymous referees for their comments, and Dr. Z. Pizlo for some very fruitful discussions regarding those comments. The comments and discussions have helped in improving the presentation of this paper.

References

- [1] Barrow HG(1987), Learning receptive fields, *Proceedings IEEE First International Conference on Neural Networks*, v4, pp 115-121.
- [2] Baylor DA, Fuortes MGF, O'Bryan PM(1971), Receptive fields of single cones in the retina of turtle, *J. Physiol*, **214**:256-294 .
- [3] Braham R, Hamblen JO(1990), The Design of a Neural Network with Biologically motivated Architecture, *IEEE Trans. Neural Networks*, **1**:251-262 .

- [4] Buser P, Imbert M(1992) (translated by Kay RH), *Vision*, The MIT Press, Cambridge, Massachussets.
- [5] Caldwell JH, Daw NW(1978), Effects of picrotoxin and strychnine on rabbit retinal ganglion cells: changes in centre surround receptive fields, *J. Physiol.* , **276**:299-310.
- [6] Daugman J(1983), Six formal properties of two dimensional anisotropic visual filters, *IEEE Trans. SMC*, **13**.
- [7] Daugman J(1988), Pattern and Motion Vision without Laplacian zero crossings, *Journal Opt. Soc. of America A*, **5**.
- [8] Dowling JE(1987), *Retina, the approachable part of the brain*, Harvard Press, Cambridge.
- [9] Enroth-Cugell C, Robson J(1966), The contrast sensitivity of retinal ganglion cells of the cat, *Journal Physiol. (London)*, **187**:517-522.
- [10] Grossberg S, Wyse L(1991), Neural network architecture for figure-ground separation of connected scenic figures, *Neural Networks*, **4**:723-742.
- [11] Hubel DO(1988), *Eye, Brain and Vision*, Scientific American Library, New York.
- [12] Joshi A, Lee CH(1991), On modelling the retina using Neural Networks, in *Proceedings, Intl. Joint Conference on Neural Networks '91*, Singapore, pp 2343-2348.
- [13] Kleinschmidt J, Dowling JE(1975), Intracellular recordings from gecko photoreceptors during light and dark adaptation, *J. Gen. Physiol*, **66**:617-648a.

- [14] Kolb H(1979), The inner plexiform layer in the retina of the cat: electron microscopic observations, *J. Neurocytol.*, **8**:295-329.
- [15] Leaky SR, Puget A, Sejnowski TJ(1990), Neural models of binocular depth perception, *Cold Spring Harbor symposia on Quantitative Biology*, **55**:765-777.
- [16] Linsker R(1988), Self Organisation in a Perceptual Network , *Computer*, **21**:105-117.
- [17] Mach E(1906), Über den Einfluss räumlich und zeitlich variieender Lichtreize auf die Gesichtswahrnehmung, *Sitzungsber. Math. Naturwiss. Kl. Kaiser. Akad. Wiss.*, **115**:633-648.
- [18] Marr D(1982), *Vision*, W.H. Freeman and Company.
- [19] Marr D, Hilderth E(1980) The theory of edge detection,*Proc. Royal Soc. London B*, **207**:187-217.
- [20] Oğuztörel M(1988), Underlying Neural Computations for some visual phenomenon, *Biol. Cybernetics*, **60**:89-106.
- [21] Oğuztörel M(1989), Neural Computations in some visual processes, in *Proceedings, 1989 IEEE International Conference on Systems, Man and Cybernetics* ,pp 664-670.
- [22] Richter J, Ullman S(1986), Non linearities in cortical simple cells and the possible detection of zero crossings, *Biol. Cybernetics*, **53** :195-202.
- [23] Thibos LN, Werblin FS(1978), The response properties of the steady antagonistic surround in the mudpuppy retina, *J. Physiol.* , **278**:79-99.

- [24] Wallach H(1948), Brightness constancy and the nature of achromatic colours, *Journal of Experimental Psychology*, **38**:310-324.
- [25] Watt R(1988), *Visual Processing: Computational Psychophysics and Cognitive Research*, Lawrence Erlbaum Associates.
- [26] West RW(1976), Light and electron microscopy of the ground squirrel retina: Functional considerations, *J. Comp. Neurol.* , **168**:355-377.
- [27] West RW, Dowling JE(1972), Synapses onto different morphological types of retinal ganglion cells, *Science*, **178**:510-512.

Figure Titles

Fig 1: Neural organisation in the Retina

Fig 2: Neural Connections in retina and the model.

Fig 3: Actual and Idealised Weight values

The X axis plots the distance from centre, and Y the Weight.

Fig 4: Actual and Idealised Weight values

The X axis plots the distance from centre, and Y the Weight.

Fig 5: Image with good contrast. Input to and Output of the model

Fig 6: Image with poor contrast. Input to and Output of the model

Fig 7: Output of sobel edge detector applied to the images.

Fig 8: Excitatory and Inhibitory connections learnt

for an On Centre ganglion.

Fig 9: Intensity response curve of the ganglion cell

The X axis plots the intensity at the Centre on log scale, and Y axis the response.

The response is shown for three different surround intensities.

Sparc	Cray		MasPar	
Time	Time	SpeedUp	Time	SpeedUp
0.98s	0.038	25.78	0.00237s	412

Table 1: Times and SpeedUps

Fig. 1

*Neural Connections
in Retina*

*Connections in
the model*

Figure 2:

Y

X

Figure 3

Y

X

Figure 4

Figure 5a

Figure 5b

Figure 6a

Figure 6b

Figure 7a

Figure 7b

Figure 7c

Figure 7d

Connection Patterns

Figure 8:

Response

Figure 9: