

Regularization in Oculomotor Adaptation

John A. Bullinaria, Patricia M. Riddell, Simon K. Rushton

Department of Psychology, The University of Reading
Reading, RG6 6AL, UK

Abstract. The oculomotor system remains plastic so that it can maintain clear single binocular vision during development and also in novel visual conditions (such as wearing new spectacles). It is important to understand this adaptation process so that we can predict in advance potential problems that might arise with new optical devices such as virtual reality head mounted displays. In this paper we present neural network models of adaptation to vertical disparities at different points in the visual field and argue that regularization (weight decay) provides a more realistic account of the empirical data than other approaches.

1. Introduction

To see objects clearly at different distances one needs to be able to make appropriate changes to the focus (accommodation) and cause appropriate rotations (vergence) of our eyes. The range of cues that drive the accommodation and vergence systems, the inevitable interactions between these systems, and the complicated dynamical responses found experimentally, render the modelling of this part of the visuo-motor system a non-trivial problem. A number of control systems models have already been formulated [2], but no one model has yet been universally accepted and there still remains some dispute over the strength and positioning of the links between the various sub-systems. These systems models are generally specified in detail by hand specifically to simulate adult performance. One advantage of connectionist models is that we are able to set up a general framework and let the network *learn for itself* how the available cues and cross-links are best used. This not only solves the problem of determining the correct adult architecture, but also provides an additional check of the model against developmental data, and allows us to simulate and understand the kinds of problems that might arise during the development process.

In this paper we focus on a particular sub-problem, namely the adaptation process that maintains clear vision as the visual system grows and when using new visual devices such as spectacles. Since there is good experimental data available showing how subjects adapt when exposed to targets with vertical disparities that vary with position in the visual field (e.g. [3, 5, 8, 9]) and a partially successful model already exists [6], we shall concentrate on formulating more realistic models of this task. Our starting point is data from Schor *et al.* [8] describing the spatial aspects of adaptation and Maxwell & Schor [5] on the temporal aspects of adaptation. Both studies had subjects presented with two targets, of equal but opposite disparities, at midline locations, until adaptation reached a steady state. The degree of adaptation was assessed for a range of disparities and target separations by measuring changes in vergence position when there was no fusion target present. The response gains were found to depend on both the spatial separation and the size of the disparity stimulus. Figure 1 shows the detailed pattern of data to be modelled, taken from Schor *et al.* [8].




Figure 1: The experimental pattern of response gains (from Schor *et al.* [8]).

2. The Neural Network Models

Following McCandless & Schor [6], we begin with the natural assumption that adaptation arises by learning associations between existing physiological outputs from brainstem neurons sensitive to eye position [4] and the required vergence response. They set up neural networks that map directly from a set of random linear input units representing eye position sensitive neurons to a single linear output unit representing the motor output of the vertical vergence system. The network training data was the same as the experimental stimuli, namely opposite local disparities placed midline in the upper and lower visual fields. A gradient descent learning algorithm was used to update iteratively the connection weights to reduce the error on the outputs for each eye position. The models were tested by comparing the vertical vergence responses at all eye positions, with no fusion target present, against the experimentally observed responses. The response as a fraction of the stimulus is the 'response gain'.

Such models are clearly capable of learning the training data exactly and hence always have a response gain of one. Since human subjects do not fully adapt to some stimuli (i.e. have gains less than one as in Figure 1), McCandless & Schor [6] stopped the training of their models when their adaptation matched the observed adaptation. Whilst their simulations do demonstrate how the output from eye position sensitive brainstem neurons could be used to drive vertical vergence changes, we feel that they accounted for the data of Figure 1 in a somewhat *ad hoc* manner.

We extended the McCandless & Schor model [6] in a number of ways. First we investigated plausible alternatives to the Linear input codings they used to represent the eye position signals from midbrain cells. Of particular interest were Sigmoidal input codings (which are simply more biologically plausible versions of their linear signals), and Difference of Gaussian codings (such as might be provided by retinal receptive fields). Each of these functions is defined by a set of parameters (e.g. height, width, centre) and a random set of these (within reasonable ranges) were assigned to each input unit to give its activation for each eye position. Typically, between 80 and 800 input units were sufficient to avoid randomization artefacts. Each network was trained by fixed step-size gradient descent on a sum squared error

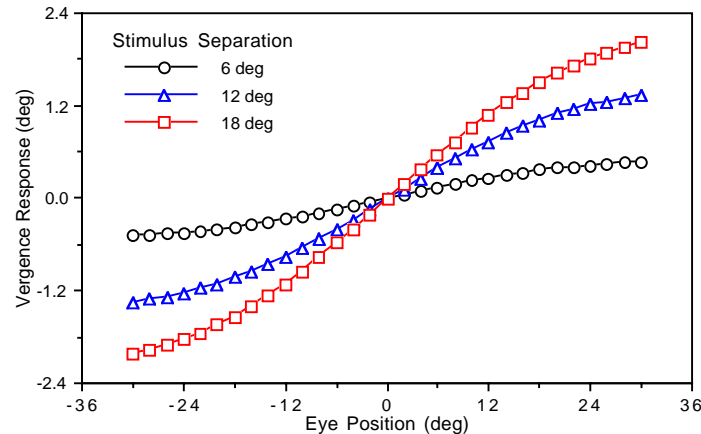


Figure 2: The output responses of three trained networks based on sigmoids.

measure; first on zero disparity to provide a steady pre-adaptive state, and then on two disparity stimuli. All the results we present will be for networks with 400 Sigmoidal input units connected directly to a single linear output unit. Naturally, these networks could achieve perfect adaptation, i.e. a response gain of one, in every case. We then considered what might be preventing this perfect gain in human systems.

We began by following McCandless & Schor [6] and restricting the amount of training. Figure 2 shows the outputs of three typical networks after 7000 epochs of training with a learning rate of 0.00001 on a stimulus disparity gradient (i.e. stimulus disparity difference over separation) of 0.1. To relate this to the experimental data in Figure 1, the response gradient is defined using the change in output between eye positions of ± 6 degrees and the response gain is this divided by the stimulus gradient. We find that the response gain is now less than one, as required, and moreover, it now depends on the stimulus separation in the right direction. The effect of the stimulus disparity gradient, however, cancels out in the gain calculation. Using the alternative input codings (e.g. Linear or Difference of Gaussian) changes the form of the network's extrapolation to larger eye positions, but has little effect on the pattern of response gains. Clearly, by picking an appropriate amount of training for each training set, we can get any pattern of gains we want, but before following that route, it makes sense to consider what other processes might account for the data.

3. Regularization

A particularly natural process that prevents networks from learning their training data exactly, and smoothes their response to potentially noisy data, is regularization in the form of weight decay [1, p338]. Poggio *et al.* [7] have reviewed the use of regularization principles in various areas of computational vision, and outlined how they could be implemented in single biological neurons or small networks of neurons. Normally, gradient descent learning operates by iteratively adjusting the network weights w to reduce some output activation error function E (e.g. the sum of squared output activation errors). If one adds a weight cost function to this error function, it

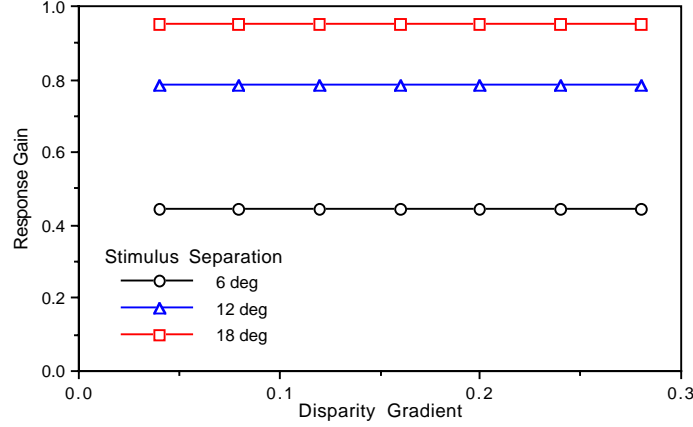


Figure 3: The pattern of response gains for our model with quadratic weight cost.

results in an extra contribution to the gradient descent weight changes $\Delta w \sim -\partial E/\partial w$ proportional to the derivative of the weight cost. If the cost function is quadratic (i.e. the sum of squared weights) we thus get a linear weight decay contribution to the weight changes. This keeps the weights low and smoothes out the network response at the expense of the output activation errors. Using such a standard quadratic weight cost in our model results in the asymptotic pattern of response gains shown (for 21 trained networks) in Figure 3. The weight decay during training has a similar effect on the network output pattern as the reduction in weights due to stopping the training early. The required effect of disparity gradient, seen in Figure 1, is still missing.

We can see why by considering a simplified model with just a single linear input. We then have just one weight w trained on a target output t for an input which is a scale factor α times half the stimulus separation s . For a sum-squared error term and n th order weight cost with a parameter λ , we train by performing gradient descent on

$$E_n = \frac{1}{2} (t - w\alpha \frac{1}{2}s)^2 + \frac{1}{n} \lambda w^n .$$

For the quadratic ($n = 2$) case we can easily solve for the learnt weight w_2

$$\frac{\partial E_2}{\partial w} = 0 \quad \Rightarrow \quad w_2 = \frac{2\alpha s t}{\alpha^2 s^2 + 4\lambda} .$$

The gain g_n is simply the ratio of the response gradient to the disparity gradient

$$g_n = \frac{\text{output/input}}{t/\frac{1}{2}s} = \frac{w_n \alpha i / i}{2t/s} = \frac{\alpha s}{2t} w_n$$

so, for the quadratic case, we obtain the exact expression for the gain

$$g_2 = \frac{1}{1 + x_2} \quad \text{where} \quad x_2 = \left(\frac{4\lambda}{\alpha^2} \right) \frac{1}{s^2}$$

which is indeed independent of the disparity gradient $d = 2t/s$.

It is clear from the above equations that, for a given stimulus separation, large disparity gradients will require larger weights than small disparity gradients, but the




Figure 4: The pattern of response gains for our model with quartic weight cost.

trade-off between the quadratic error term and the quadratic weight cost is leaving no net effect. A higher order weight cost, however, should have the desired result. We can also solve the above equations for the cubic ($n = 3$) weight cost and find

$$g_3 = \frac{2}{x_3} (\sqrt{1+x_3} - 1) \quad \text{where} \quad x_3 = \left(\frac{16\lambda}{\alpha^3} \right) \frac{d}{s^2}$$

where the desired disparity gradient dependence finally appears. Solving for the higher order weight cost cases becomes messy, but Newton's method starting at $g_n = 1$ allows us to iteratively approximate the gains using

$$\Delta g_n = - \frac{\left(\frac{\alpha s}{2} \right)^2 (g_n - 1) + \lambda \left(\frac{d}{\alpha} \right)^{n-2} g_n^{n-1}}{\left(\frac{\alpha s}{2} \right)^2 + (n-1) \lambda \left(\frac{d}{\alpha} \right)^{n-2} g_n^{n-2}},$$

and similarly for mixed order costs. A quartic ($n = 4$) weight cost, for example, will result in even stronger cubic weight decay which will prevent large disparity gradients from being learnt in a reasonably realistic manner. Implementing such a weight cost in our full model results in the pattern of response gains shown in Figure 4, which is now much closer to the empirical pattern of Figure 1. Unfortunately, the experimental data is too noisy to constrain our models much further at this stage.

4. Conclusions and Discussion

We have presented a preliminary investigation into possible improvements of the physiologically plausible neural network model of oculomotor adaptation of McCandless & Schor [6]. First, we have shown that the details of the biologically inspired eye position coding are not crucial to the models' pattern of response gains. Moreover, variations in them can be used to account for details of the saturation of the adaptation response as the eye position moves beyond the targets. Second, we have argued that weight cost (which, in the context of gradient descent learning, translates into weight decay) constitutes a natural process whereby local adaptations become smoothed out, and have shown how it restricts the adaptation response in line with the

empirical gain data over a wide range of stimuli. We consider this more appropriate than imposing individual limits on the training times for each stimulus set. However, McCandless & Schor [6] have noted subject reports that whilst small disparities were easily fused, the larger disparities often remained diplopic. This means that the subjects were not able to produce the required eye rotations to foveate the targets in each eye and hence had less training time with the eyes in the correct position. We intend to study the effect of incorporating such empirically measured variations in training times to see if this really can provide an alternative account of the data.

We have found that the known hemifield differences in temporal response [5] can be accommodated in our models either by varying the training data frequencies in line with attention, or by manipulating the parameters describing the input position coding. This is currently being investigated further. We also intend to examine the models' efficacy in simulating the experimental data for more complex target configurations. Finally, we are exploring how the use of different learning algorithms might affect the results. In particular, we propose to further improve the realism of the models by implementing them with more biologically plausible learning algorithms.

The main motivation behind our modelling work is to understand the fundamental biological mechanisms involved in local adaptation and to gain insight into the plasticity of the oculomotor system both in adults and during development. In this way, we can hope to study the development of oculomotor disorders and to resolve, ethically, issues of safety when altering visual environments (e.g. by the use of virtual reality headsets) both in adults and in children with their more plastic visual systems.

References

1. Bishop, C.M. (1995) *Neural Networks for Pattern Recognition*. Oxford, UK: Oxford University Press.
2. Eadie, A.S. & Cardin, P.J. (1995) Evolution of control system models of ocular accommodation, vergence and their interaction. *Medical & Biological Engineering & Computing*, **33**, 517-524.
3. Henson, D.B. & Dharamshi, B.G. (1982) Oculomotor adaptation to induced heterophoria and anisometropia. *Investigative Ophthalmology and Visual Science*, **22**, 234-240.
4. King, M.W., Fuchs, A.F. & Magnin, M. (1981) Vertical eye movement-related responses of neurons in midbrain near interstitial nucleus of Cajal. *Journal of Neurophysiology*, **46**, 549-562.
5. Maxwell, J.S. & Schor, C.M. (1994) Mechanisms of vertical phoria adaptation revealed by time-course and two dimensional spatiotopic maps. *Vision Research*, **34**, 241-251.
6. McCandless, J.W. & Schor, C.M. (1997) A neural net model of the adaptation of binocular vertical eye alignment. *Network*, **8**, 55-70.
7. Poggio, T., Torre, V. & Koch, C. (1985) Computational Vision and Regularization Theory. *Nature*, **317**, 314-319.
8. Schor, C., Gleason, G., Maxwell, J. & Lunn, R. (1993) Spatial aspects of vertical phoria adaptation. *Vision Research*, **33**, 73-84.
9. Schor, C.M. & McCandless, J.W. (1997) Context-specific adaptation of vertical vergence to correlates of eye position. *Vision Research*, **37**, 1929-1937.