Incorporating Developmental Factors into Models of Accommodation and Vergence

Patricia M. Riddell & John A. Bullinaria

Department of Psychology, University of Reading
3, Earley Gate, Reading, RG6 6AL, UK.

15 July 1999

Grant Support: EPSRC grant GR/L82274, MRC grant G9608874N.

Keywords: *accommodation, vergence, visual development, depth perception, eye movements, neural network models*

Abstract

Purpose: To investigate how known developmental factors can be incorporated into neural network models of the accommodation and vergence systems. To demonstrate that models of this type can learn adult-like linkages between accommodation and vergence by simple associative mechanisms.

Methods: Existing linear systems models were used as a starting point for the design of a series of multi-layer neural network models of accommodation and vergence. Unlike the systems models that are generally specified in detail by hand, the neural models learn for themselves how to perform appropriately. A general connectionist framework was set up and trained, by updating initially small random connection weights using gradient descent, so that the network outputs for realistic patterns of input cues (blur, disparity, etc.) corresponded to appropriate patterns of accommodation and vergence responses. Maturational factors were incorporated into the models by using empirical infant data to define starting points, by varying parameters such as depth of field during training and by bringing in cues at appropriate developmental stages in the model. The outputs and internal connection weights were recorded throughout the training to determine the developmental progression of the model.

Results: Training the models in this way resulted in developmental progressions which matched empirical infant data and automatically generated linkages between the accommodation and vergence subsystems with strengths within the normal adult range.

Conclusions: A range of known maturational factors together with empirical infant data can be used to constrain developmental models of accommodation and vergence. Starting with small random weights, these neural network models can learn to associate changes in accommodation with related changes in vergence, and develop adult-like linkages between accommodation and vergence. This provides a demonstration of how such linkages might develop through simple associative mechanisms.

Introduction

The literature contains many linear systems models of accommodation and vergence that provide a good account of accommodation and vergence responses to unpredictable stimuli in normal adults. However, there has been little attempt to determine how the accommodation and vergence systems develop over time. There are two historical perspectives on how these control systems might develop over time. Hering[1] believed that much of the development of vergence depended on maturational factors and that, for instance, the yoking between the eyes was present at birth. Helmholtz[2], on the other hand, thought that the development of accommodation and vergence was learnt through experience. He believed that symmetric vergence responses were a learnt property of the vergence system. If Helmholtz was correct then it should be possible to model the development of the interactions between accommodation and vergence using neural network models that learn by simple associative mechanisms.

We have attempted to model the development of accommodation and vergence from birth to adulthood by designing neural network models which can learn appropriate, adult responses for accommodation, vergence and the linkages between these systems. Since the time of Hering and Helmholtz more has been learnt about the state of accommodation and vergence in the infant. It is known, for instance, that stereopsis develops suddenly and rapidly at 14-16 weeks of age, and it is thought that this development is dependant on maturational changes in the primary visual cortex – the development of the ocular dominance columns. In this study, therefore, it is appropriate to base our approach on a combination of learnt and maturational elements. Maturational elements have been incorporated into the model where these are known to occur, but, otherwise, the model is free to learn the associations between the accommodative and vergence systems. The assumptions we have made are outlined in this paper along with a description of the initial set of models which have resulted from this approach.

The models were evaluated by two basic criteria:

1. Do the models learn in a similar developmental progression to human infants? This was assessed by comparing the accommodation and vergence responses of the model at different stages with the responses of infants at matched developmental stages.
2. Are the models able to learn appropriate adult linkages between vergence and accommodation? This was tested by comparing the gain of the accommodative convergence/convergence response (AC/A ratio) and the vergence accommodation/accommodation response (CA/C ratio) learnt by the model with empirical adult values.

The results, so far, suggest that many of the parameters of the accommodation and vergence systems, and their reciprocal linkages, can be learnt by simple associative mechanisms.

Methods

Our starting point was, what is now, a standard linear systems model of accommodation and vergence[3-5] (Figure 1). This was used as a basis for the design of a series of multi-layered, neural network models. The neural network models were designed by taking each of the components of the systems models e.g. controllers, tonic leaky integrators, etc, and mapping

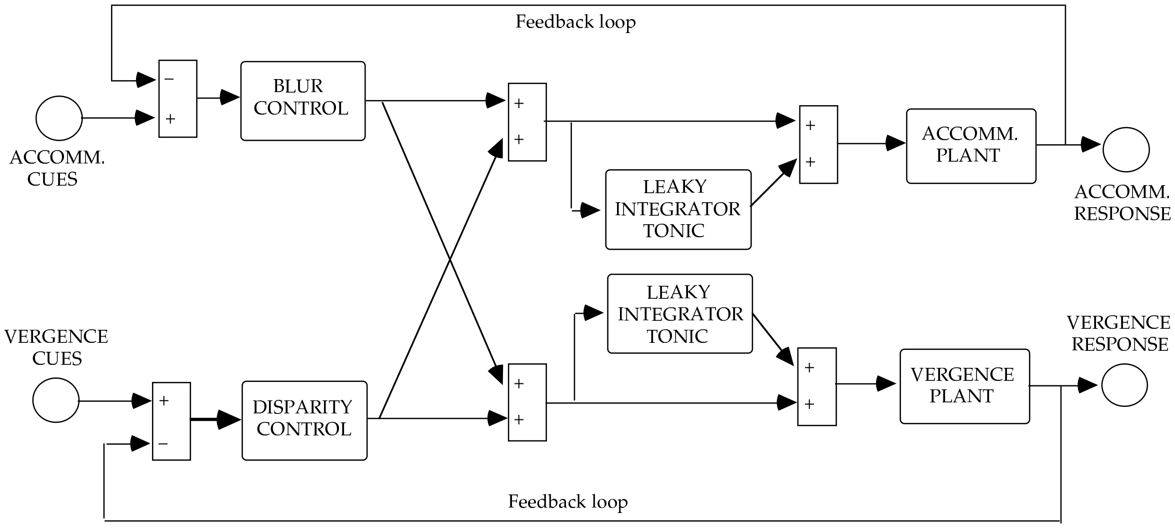

Figure 1 – Schematic diagram of a standard linear systems model of accommodation, vergence and the linkages between these systems.

them to leaky integrator neurons in the neural network models, with each gain in the linear systems model becoming a learnable connection weight. Then, unlike the systems models, which generally are specified in detail by hand, the neural network models are expected to *learn* appropriate behaviour.

A general connectionist framework was set up as shown in Figure 2 and trained using a gradient descent algorithm that iteratively updates initially small, random, connection weights to produce appropriate vergence and accommodation responses for realistic input cues (blur, disparity etc.). In this system, components such as the cross-links are learnt if they are useful for producing good performance rather than being specified *a priori* in an attempt to account for actual human performance. Similarly, both fast and slow neurons (defined by their different time constants) are used in the model. This allows the network to learn useful structures such as the two component phasic controllers proposed by Krishnan and Stark[6] and Schor and Kotuluk[7], and to produce leaky integrator tonic components of the appropriate strengths, without these being imposed by the modeller. The aim of the modelling is twofold: 1) The pattern of learning in our model should correspond to the developmental changes found in infants and children, and 2) The performance of the final, fully trained, network should match empirical patterns of competent adult behaviour.

Although many of the details and parameters of the neural network model follow directly from the linear systems models, in particular that of Schor et al[3], several complications due to developmental factors need to be taken into account:

1. Panum’s fusional area.

The accuracy of vergence response will be at least partially related to the known developmental change in the size of Panum’s fusional area. According to Aslin: “Panum's fusion area in adults is approximately 10 arc minutes... If the size of Panum's fusion area in infants is proportional to peak contrast sensitivity (as it is in adults), the extent of Panum's fusion area in 9- to 10- week olds should be 200 arc minutes,” (Aslin[8], p35). We used this to

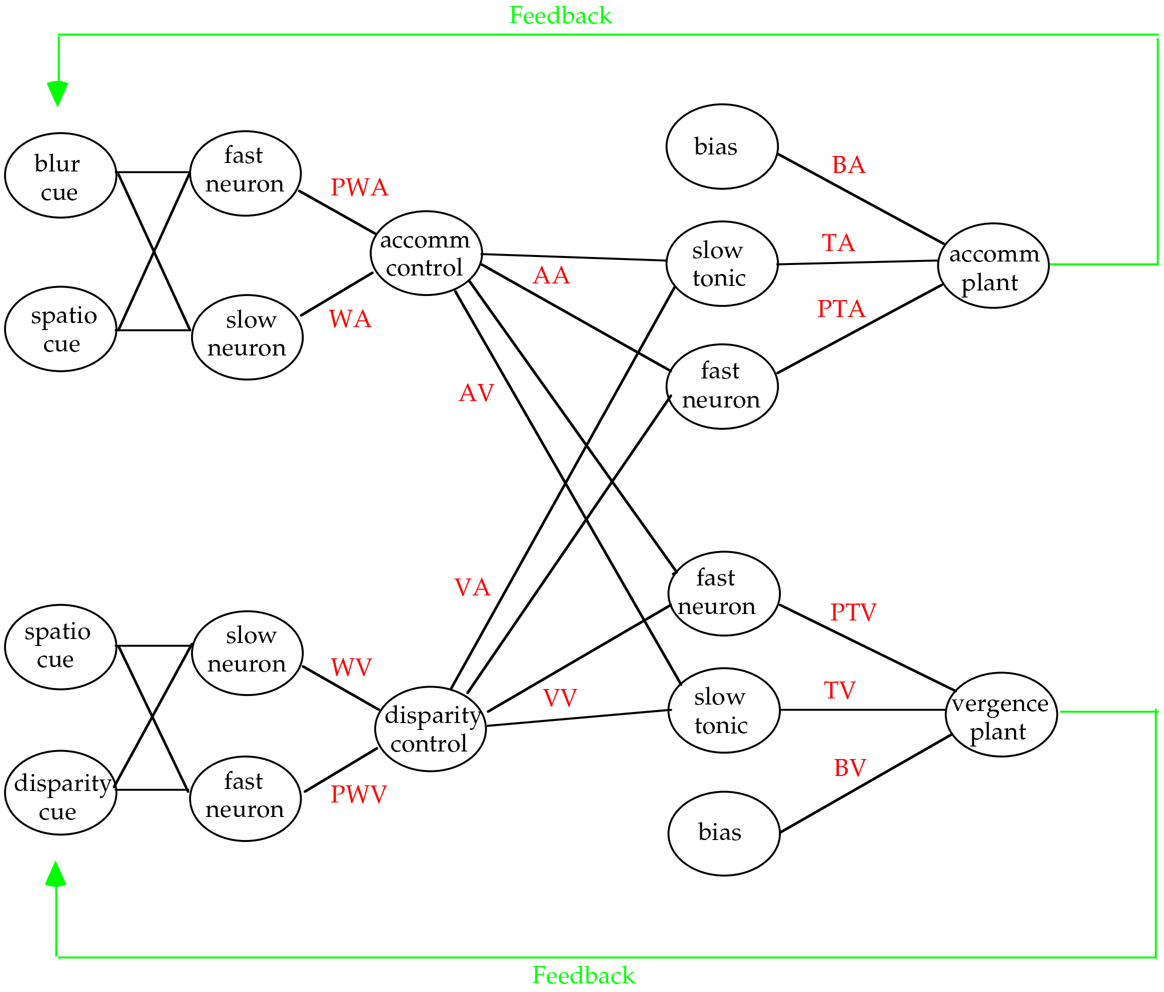

Figure 2 – Schematic diagram of the neural network model of accommodation, vergence and the linkages between these systems. Connection weights are designated by the letters above connections.

define the starting point of Panum's area, and therefore the resulting accuracy of the vergence response, for the model, and then decreased the size of Panum's area as a function of the known improvement in contrast sensitivity with age[9]. This resulted in vergence responses that changed in the appropriate direction for target distance, but were of a smaller gain than found in adults. This was considered to be a suitable match to the data on early infant vergence responses[10].

2. Depth of field

There are various data that describe the early accommodative abilities of the infant. The earliest work[11] suggested that infants have relatively fixed accommodation at about 4 dioptres. Later work[12-13] suggested that some infants are able to change their accommodation to targets from 1 to 6 dioptres. One explanation of these results has been that infants with flat accommodative responses are failing to attend to the stimuli[8]. However, all of these studies used small populations of infants. In a much larger population, Hainline, Riddell, Grose-Fifer and Abramov[14] described a range of behaviours, with around 50% of infants younger than 60 days demonstrating appropriate accommodative changes with target distance. In a further 36%, however, there was no change in accommodation despite changes in simultaneously measured vergence responses. This rules out the idea that the lack of change in accommodation is the result of inattention in all infants. It is also clear from their data that, in at least 40% of infants, the linkages between accommodation and vergence are not developed at birth. Since we are interested in whether these linkages can be learnt, it is important to test the model with the most stringent initial conditions. We therefore chose to start the model with an accommodation response that did not vary with target distance, and that was fixed at about 4 dioptres. The vergence response, however, did vary appropriately with target distance. This is a pattern of response that was found in a significant proportion of the infants tested by Hainline and Riddell[10]. We used an initially high depth of field (DOF) to account for fixed accommodation in babies. There are many possible reasons for poor accommodation in newborn infants including sensory factors, and the inability of the infant to detect the sign of the blur[8]. What is important for this study is that by setting DOF large in infants, we could reproduce the flat accommodation functions reported in previous studies[11,14]. The adult DOF represents a typical value from Campbell[15].

3. *Initial network connection weights*

The choice of initial weights for a gradient descent learning process such as ours can have a big effect on the outcome of the training[16]. Currently we start the biases with estimated newborn baby values, and some of the links (namely those which in the corresponding systems models do not have a gain parameter) are set to start with unit weights. All other weights start with small random values. Initially, the vergence bias is set based on the empirical data of Aslin[17] who measured dark vergence responses in 1-4 month old infants. The accommodative bias is set to give the initial myopic behaviour seen in previous studies[11, 14]. How this affects the models' performance is currently under investigation.

4. *Maturational elements*

According to Birch: "After the abrupt onset of stereopsis at 3-5 months of age, the average infant attains a stereoacuity of 60 secs or better by 6 months of age. The averaging of ages obscures what is a remarkably rapid improvement in stereoacuity; that is, the developmental curve for stereoacuity is steeper for individual infants than for the sample as a whole.... Individual infants' stereoacuities improved from 60 minutes of arc to 60 secs of arc over a time span of 5-6 weeks." (Birch[18], p228)

It is therefore clear that vergence changes cannot be based on disparity calculations before 14-16 weeks of age and that we must build this into our models. Prior to the onset of the disparity system we assume that an effective vergence response is driven by independent foveation.

5. *Spatiotopic and retinotopic cues*

The use of spatiotopic and retinotopic cues for determining the distance to a target has been addressed in some detail[19] and it has been suggested that the visual system performs a weighted sum of cues to depth. We follow a much simpler approach as in the systems models of Schor, et al[3]. In particular, we assume that the spatiotopic cues (e.g. occlusion, texture, shading etc) drives the system into the range where the more accurate retinotopic cues (i.e.disparity and blur) can take over. Currently, no further age dependencies are added here, beyond those discussed in Points 1-4 above.

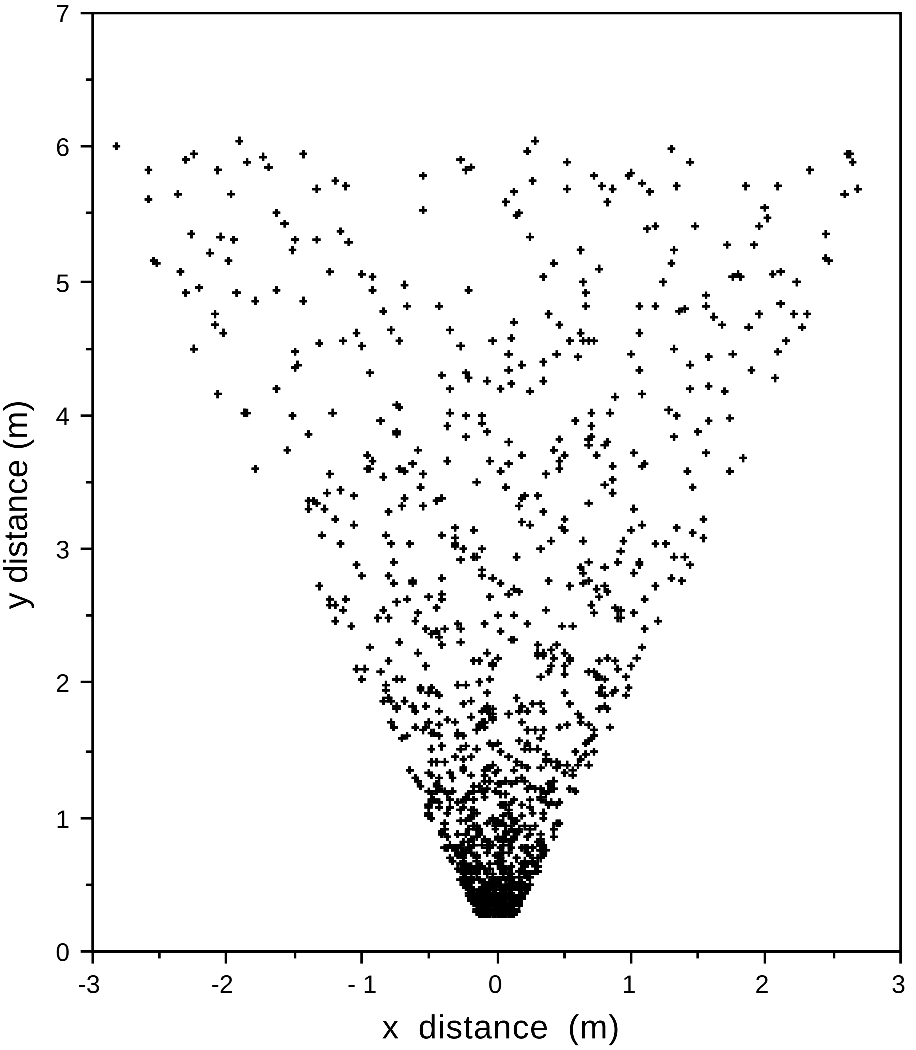

Figure 3 – Distribution of training points: The x-axis gives the horizontal spread of points while the y-axis shows the distribution of points in depth.

6. Monocular cues

Currently we take the disparity cue to be absent about 10% of the time. This is a very rough approximation to what is really happening, but it might provide a reason for the more accurate vergence system to make use of the information provided by the cross-links from the accommodation system. (Better accuracy is enough to induce the development of the vergence to accommodation cross-links.)

7. Cost function for gradient descent learning

Gradient descent learning works by performing an iterative sequence of small adjustments to the set of free parameters (in our case, connection weights) to minimise some appropriate cost function. Clearly we need to minimise the mean absolute values of both blur and disparity. We also need to prevent large oscillations about the required responses (i.e. under damping). This can be easily implemented by minimising the mean absolute velocities. Our gradient descent cost function is a trade off between these factors, and some experimentation is required to determine the appropriate relative strengths. Currently, we do not incorporate age dependencies here.

8. Training data

The training data is generated from random points in the horizontal plane in front of the eyes as in Figure 3. The retinotopic cues consist of appropriate values for either smooth or sudden transitions (50% of each) between these points with random fixation durations of between 2 and 20 seconds. The spatiotopic cues are currently taken to be noisy versions of the appropriate retinotopic cues. Each epoch of training comprises a small step in weight space in the direction that best reduces the cost function discussed above for a new 80 second section of this training sequence. We may need to re-assess these details later, in the process of refining our models' performance. Currently we do not vary these distributions with age, but we aim in the near future to investigate various plausible age and subject dependent variations.

9. Within- and cross-system weights

It is possible that it is easier to learn weights within a system than between systems. However, we currently assume our learning rates (i.e. gradient descent step sizes) to be the same for both within and between system weights. We may eventually need to re-assess this in the light of our models' performance.

10. Time-scales

In order to match the developmental changes seen in infants with the learning in our model, a time scale correspondence was set up based on the time of onset of stereopsis. Since we know that this occurs between 14-16 weeks of life, we can use the number of epochs of training before the disparity cue is introduced to calculate the number of epochs that corresponds to one week of life. In the model we present here, each week of a baby's life corresponds to about thirty-five epochs of training in our networks. We then use this to fix appropriate time scales for the other maturational factors discussed above (e.g. the developmental changes in Panum's fusional area and depth of field). The network learning rates (i.e. gradient descent step sizes) are then set so that the pattern of network development is synchronised with these maturational changes.

It is clear that existing empirical results leave many details of our model under- constrained, and, even given the assumptions outlined above and the conventional strategy of keeping the models as simple as possible to start with, we still have considerable room to maneuver. Here we shall present the results for one typical instantiation of our model that demonstrates how models of this type can be successful in learning appropriate accommodation and vergence behaviors. Often, related models with only slightly different design choices or parameter settings can be considerably less successful. Further details of our modelling framework, together with an analysis of all the consequences of the specific design choices and the sensitivity to all the remaining free parameters, will be published elsewhere (Bullinaria & Riddell, in preparation).

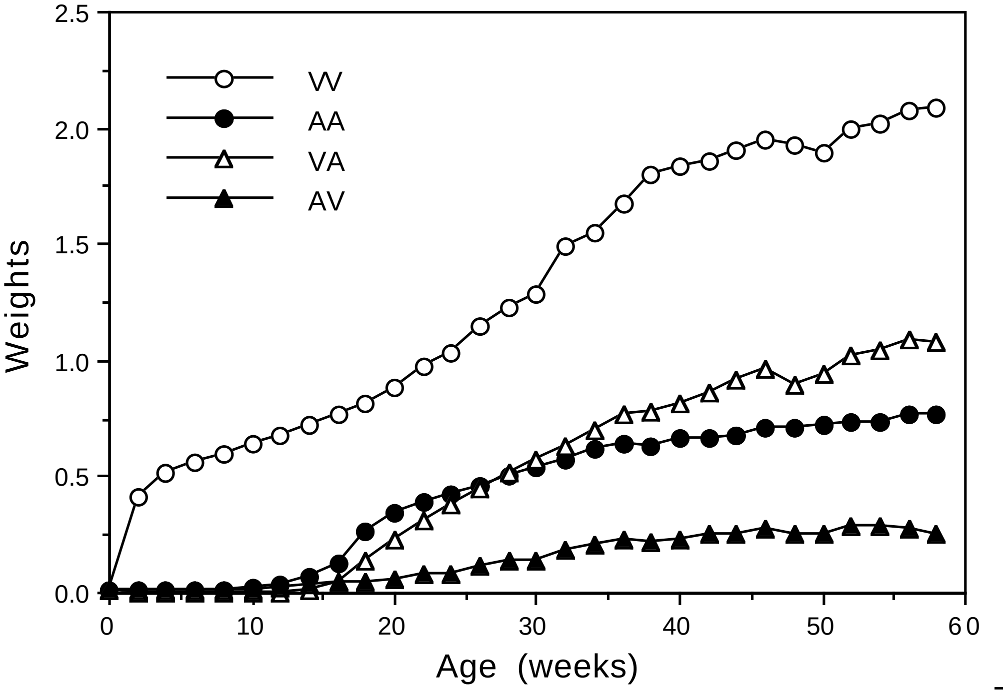

Figure 4 – Change in connection weights with development in weeks. It can be seen that, with time, the model learns to associate vergence and accommodation (VA and AV weights increase over time).

Results

The development of the networks' internal connection weights throughout the training process specify the developmental progress of the model. Given that our network is essentially a linear system, we can renormalise all of the weights in Figure 2 so that: WV = WA = PTV = PTA = 1. This does not affect the network's performance but does allow us to simplify the presentation of our results. Figure 4 shows how the renormalised weights change during learning in this model. From this figure, it can be seen that the vergence system weights develop earlier and rise to relatively higher values than the accommodation weights.

In these networks, the weights tend to level off to their adult values between weeks 50 and 60. The first thing to check is that the fully trained system responds appropriately. We find that it does: for example, Figure 5A shows its response to a typical series of vergence steps and transitions. For comparison, Figure 5B shows how easily the model can behave inappropriately, in this case, by choosing different relative values for the parameters in the gradient descent cost function.

Next, we look at the accommodation and vergence responses of the model at different time slices. These have been compared to infant data collected using photorefraction[20]. Infant accommodation and vergence are measured for a series of five targets that vary in distance from 25 cm to 200 cm. We then compare these with the responses of the model to the same set of targets. The results are converted to plots showing asymptotic response versus target demand for accommodation and vergence. Figures 6 A-E show the responses of the model

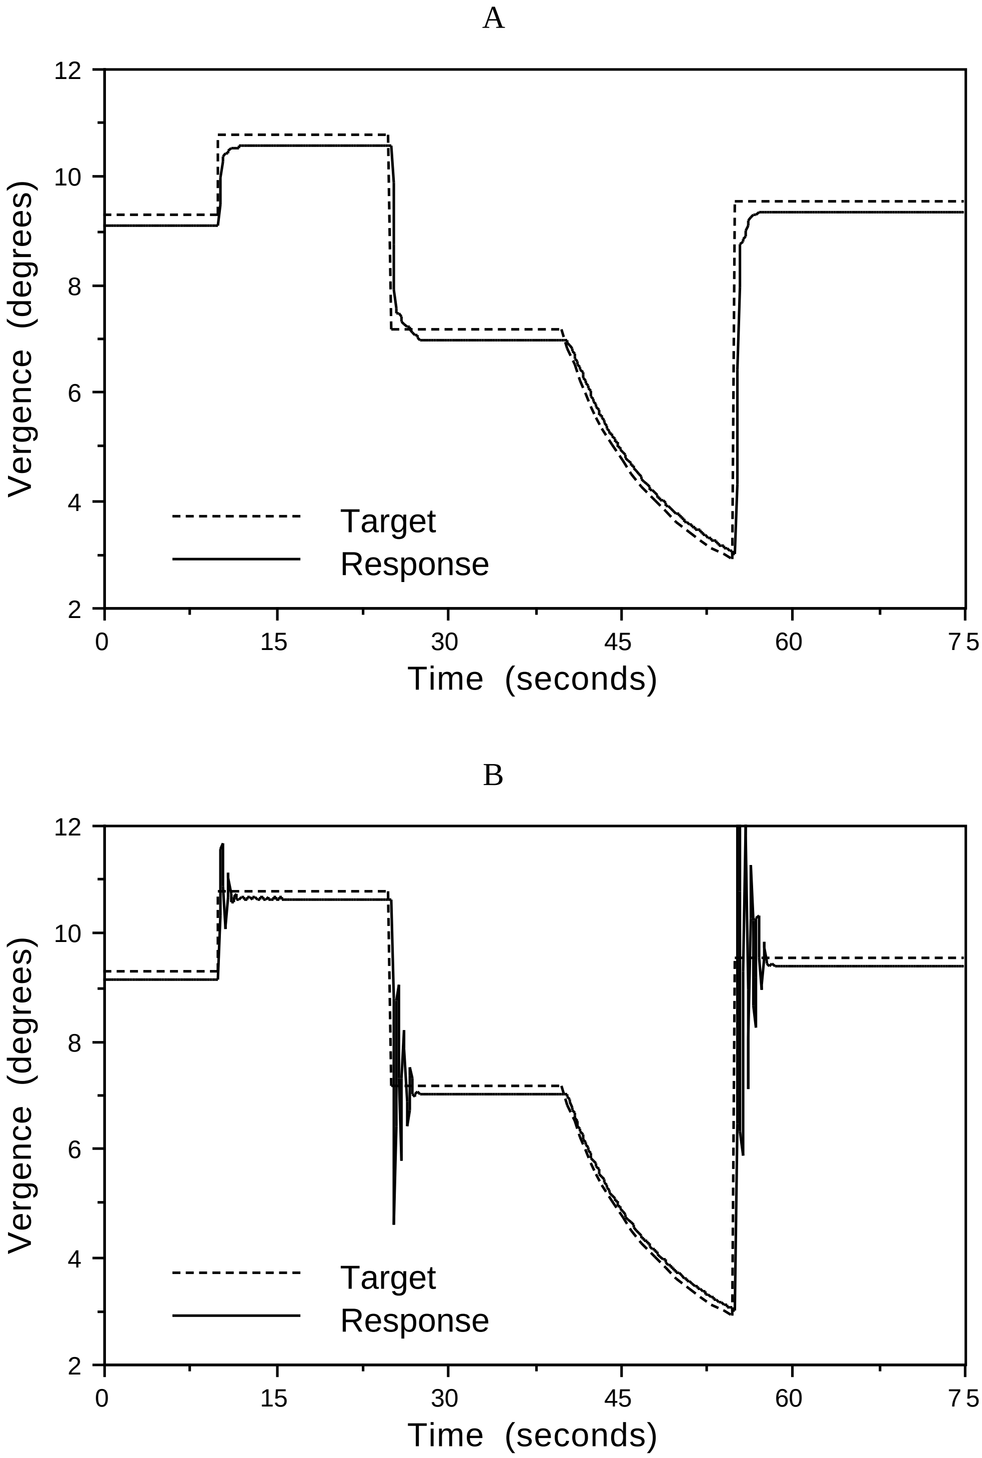

Figure 5 – The vergence responses of the model to step and continuous target movements in depth. A: the response of the model when the values for connection weights have been learnt appropriately. B: an example of a similar model where inappropriate connection weights were learnt, leading to overshoots and oscillations in vergence response when the target moved from one position to another.

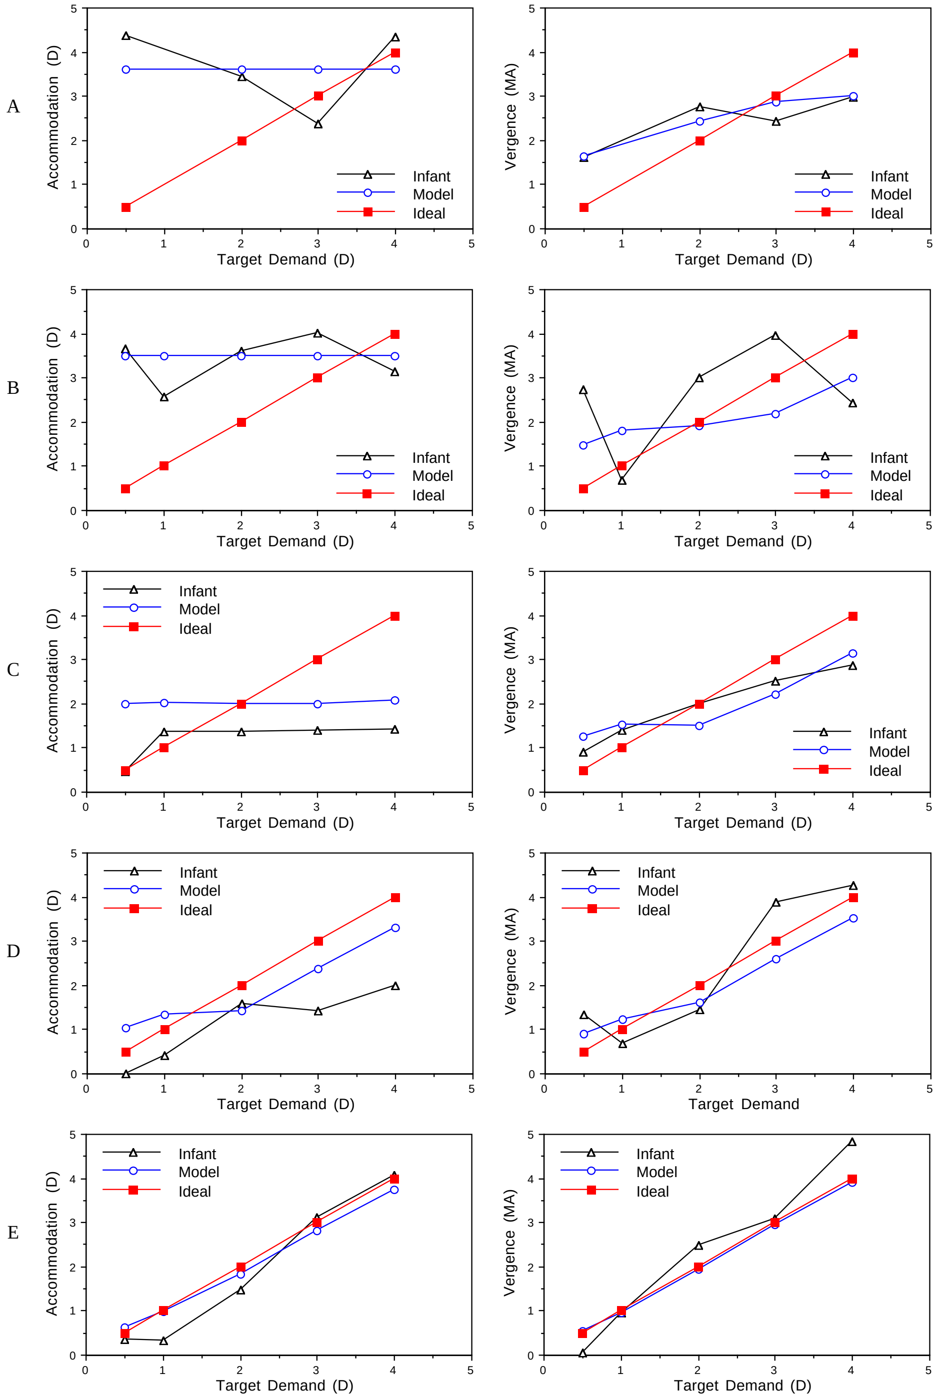

Figure 6 – Comparison of the response of the model to a pseudorandom sequence of targets at 3D, 0.5D, 4D, 1D and 2D with the response of an infant of an equivalent developmental stage to the same target sequence. The ideal response to this sequence is also given for comparison. In each case, the left-hand panel shows the accommodative response to the targets and the right-hand panel shows the vergence response. A) response at 1 week of age; B) response at 6 weeks of age; C) response at 10 weeks of age; D) response at 17 weeks of age; E) response at 35 weeks of age. Note that at each stage, the model performance is similar to that of the infant demonstrating the similarity between the developmental time-course and behaviour of the model and human infants.

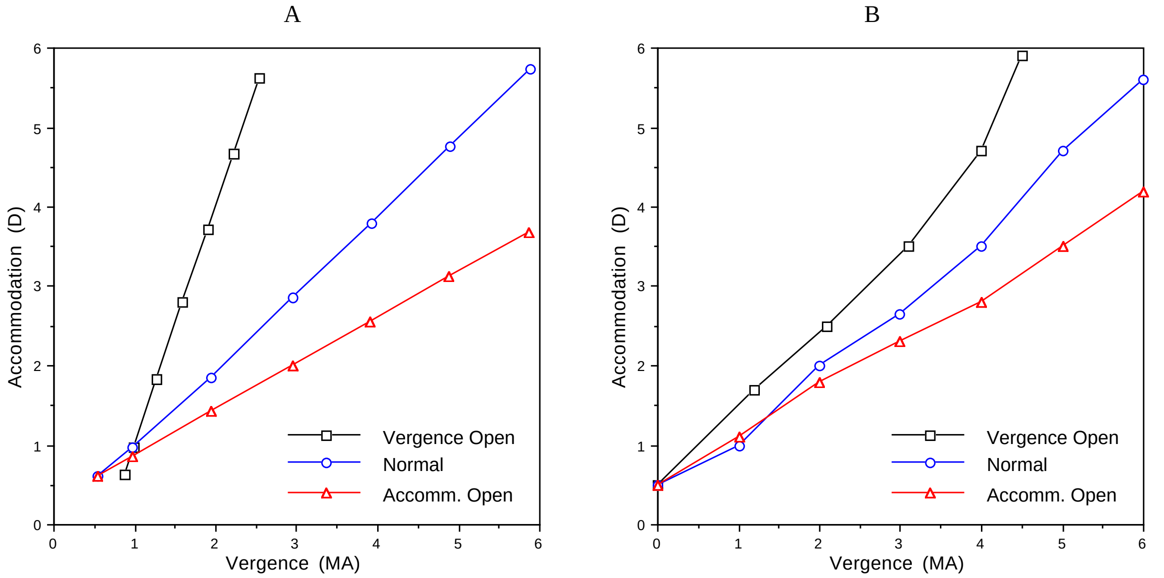

Figure 7 - A: Vergence (x-axis) and accommodative (y-axis) responses of the model to targets at different depths under three conditions: normal viewing conditions; vergence open loop (used to calculate the AC/A ratio); and accommodation open loop (used to calculate the CA/C ratio). B: Normal adult vergence and accommodation responses to targets at different depths under the same three conditions (data from Fincham & Walton[21]*).*

alongside those of representative infants of similar ages, and the ideal responses. It is clear from these plots that the model is learning accommodation and vergence responses in a manner which is similar to that of at least some infants

Finally, we ask if the model has learnt appropriate cross-links. Figure 7 a shows the adult response of the model under normal conditions and with vergence or accommodation open loop for target demands of 0.5, 1, 2,3, 4, 5, and 6 Dioptres. The corresponding empirical data from Fincham and Walton[21] is shown in Figure 7b. In this study, accommodation and vergence were measured simultaneously while subjects looked at targets at different distances with both eyes open (normal condition), with one eye shut (AC/A ratio) and with an artificial pupil of 0.5mm (CA/A ratio). Data are given for one characteristic adult. Comparing these figures reveals that there is a close match between the models' learnt CA/C ratio and the empirical findings for this subject. When the results for the AC/A ratio are compared, however, it can be seen that the model produces an AC/A ratio that is lower than the empirical adult data. There are, however, known age trends in the values of both the AC/A and CA/C ratios[22]. The AC/A ratio increases while the CA/C ratio decreases with age. This is due to changes in the elastic properties of the lens – a feature not yet incorporated into the model. The available empirical data is from a 32 year old adult, who will have a larger AC/A ratio but lower CA/C ratio than an infant and so might not be the appropriate comparison for the final developmental stage of the model. Unfortunately, data for children is currently unavailable.

Discussion

Our results show that it is possible to formulate neural network models that learn appropriate accommodation and vergence responses with linkages between the systems that are in line with empirical findings. Several factors seem to be crucial in producing exactly the right developmental sequence of events (e.g the onset of the disparity system relative to the decrease in depth of field with development). However, so far we have found that potentially realistic changes made to the model tend to result in behaviours that are within the normal range. Our modelling results confirm us in the opinion that producing appropriate accommodation and vergence responses relies crucially on a complex series of interactions between the various sources of information. There are several limitations to the models as they stand, but also many avenues for improvement. These will be addressed in future models. Three limitations of particular importance are described below:

1. We found lower AC/A ratios than found in normal adults, however, it would be more appropriate to compare the model AC/A ratio with those of younger subjects. The lower AC/A ratio in the model appears primarily to be due to the difference in accuracy between the blur and disparity cues. Since the disparity cue is more accurate than the blur cue, there is a strong vergence to accommodation link, but the model does not learn to drive vergence responses with accommodation to the same extent as is found in normal adults. This difference between the models' performance and normal adult performance requires further investigation.
2. We do not yet know what causes the poorer accommodation in newborns. One possibility is that this is limited by poor sensory abilities. It might be possible to test this hypothesis in adults by producing stimuli that mimic the quality of the images viewed by newborn infants. If adults produced reduced accommodation to such stimuli, this would suggest that the poor accommodation of the newborn is the result of sensory limitations. This hypothesis is currently being tested.
3. We currently use the bias in the accommodation system to produce the myopic accommodative performance seen in newborn infants. It is known, however, that many newborn infants have a hyperopic refractive error. We would like to use these neural network models to investigate the process of emmetropization in infants and the relation of this process to accommodation.

Clearly, further studies will be needed to fully investigate how neural network models of the type described here can be used to help us understand the processes involved in the development of accommodation and vergence in humans. In the long term, we plan to use our models to investigate which changes to their properties result in developmental trajectories that are similar to those found in oculomotor disorders, and hence identify precursors of potential abnormal function in children, with a view to designing plausible remedial actions.

Acknowledgements

The authors would like to thank Simon Rushton for valuable discussions and Judy Turner, Anna Horwood and Sheila Houston in the Baby Vision Lab for providing examples of infant accommodation and vergence.

References

1. Hering, E. *Die Lehre vom binocularen Sehen*. 1896 Leipzig:Engelmann.
2. Helmholtz, H. von *Helmholtz's Treatise on Physiological Optics*, Vol 3. 1925 J.P.C Southall (Trans) New York: Optical Society of America.
3. Schor CM, Alexander J, Cormack L, Stevenson S Negative feedback control model of proximal convergence and accommodation. *Ophthalmol Physiol Opt.* 1992;12: 307-318.
4. Hung GK, Semmlow JL Static behavior of accommodation and vergence: Computer simulation of an interactive dual-feedback system. *IEEE Trans Biomed Eng.* 1980; 27:439-447.
5. Judge SJ How is binocularity maintained during convergence and divergence? *Eye,* 1996; 10:172-176.
6. Krishnan VV, Stark L A heuristic model for the human vergende eye movement system. *IEEE Trans Biomed Eng.* 1977; 24:44-49.
7. Schor CM, Kotulak JC Dynamic interactions between accommodation and vergence are velocity sensitive. *Vision Res.* 1986; 26:927-942.
8. Aslin RN Infant accommodation and vergence. In: Simons K ed. *Early Visual Development, Normal and Abnormal*. Oxford:Oxford University Press; 1993:30-38.
9. Norcia AM, Tyler CW, Hamer RD Development of contrast sensitivity in the human infant. *Vision Res*. 1990; 30:1475-1486.
10. Hainline, L, Riddell PM Binocular alignment and vergence in early infancy. *Vision Res.* 1995*;* 35:3229-3236.
11. Haynes H, White BL, Held R. Visual accommodation in human infants. *Science.* 1965; 141:528-530.
12. Banks MS Infant refraction and accommodation. *Int Ophthalmol Clin.* 1980; 51:646-666.
13. Brookman KE Ocular accommodation in human infants. *Am J Optom Physiol Opt.* 1983; 60:91-99.
14. Hainline L, Riddell PM, Grose-Fifer J, Abramov I Development of accommodation and convergence in infancy. *Behav Brain Res.* 1992; 49:33-50.
15. Campbell FW The depth of field of the human eye. *Optica Acta*. 1957; 4:157-164.
16. Kolen JF, Pollack JB Back Propagation is Sensitive to Initial Conditions. In Lippman RP, Moody JE, Touretzky DS eds. *Advances in Neural Information Processing Systems.* San Mateo, CA: Morgan Kauffmann. 1991;*3*:860-867.
17. Aslin RN Dark vergence in human infants: implications for the development of binocular vision. *Acta Psychol (Amst).* 1986; 63:309-322.
18. Birch EE Stereopsis in infants and its developmental relation to visual acuity. In: Simons K ed. *Early Visual Development, Normal and Abnormal*. Oxford:Oxford University Press; 1993:224-236.
19. Landy MS, Maloney LT, Johnston EB, Young M Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision Res.* 1995; 35:389-412.
20. Abramov I, Hainline, L, Duckman, RH Screening infant vision with paraxial photorefraction. *Optom Vis Sci.* 1990; 67:538-545.
21. Fincham EF, Walton J. The reciprocal actions of accommodation and convergence. *J Physiol (Lond).* 1957; 137:488-508.
22. Bruce AS, Atchison DA, Bhoola H Accommodation-vergence relationships and age. *Invest Ophthalmol Vis Sci.* 1995; 36:406-413.