Modelling The Prenatal Development of The Lateral Geniculate Nucleus

Stephen Eglen
School of Cognitive and Computing Sciences, University of Sussex,
Brighton, BN1 9QH, UK. Email stephene@cogs.susx.ac.uk

Abstract

The lateral geniculate nucleus is an important structure within the mammalian visual pathway. It organises its retinal inputs into multiple topographic maps according to their eye of origin. During prenatal development these retinal inputs are spontaneously active due to waves of activity that propagate across the retinae. A popular hypothesis is that these waves drive activity dependent mechanisms to generate the mature state of LGN connectivity. A model is presented here of the joint prenatal development of topography and ocular segregation in the LGN, based on earlier work. In contrast to previous models of ocular dominance, it is shown that the model can develop normally without any form of postsynaptic normalisation. Furthermore, presynaptic normalisation must be divisive to produce the projection columns found within the LGN. It is also shown that by changing the rate at which waves are generated, the standard covariance rule must be adapted to prevent it learning during the long periods of silence between waves.

1 Introduction

In the mammalian visual system, the lateral geniculate nucleus (LGN) is the major target for retinal axons from both eyes. In the mature LGN, the retinal axons are ordered according to two principles. Firstly, there is a topographic mapping of retinal axons into geniculate cells, such that neighbouring retinal cells project to neighbouring geniculate cells. The mapping of retina to geniculate is particularly interesting since neighbouring retinal cells project to a column of cells in the LGN, rather than the normal point to point mapping found in other areas. Secondly, the retinal axons segregate according to their eye of origin, forming layers within the LGN – cells from the contralateral eye innervate layer A of the LGN and cells from the ipsilateral eye innervate layer A1 of the LGN. A major problem in neuroscience is to understand the developmental processes that create these mappings from one group of cells to another.

It is commonly thought that there are two kinds of developmental processes – those requiring activity and those that do not (see [2] for a review). The development of some maps, such as cortical ocular dominance, relies heavily on activity dependent processes driven by postnatal patterned vision. In the case of cat LGN development however, by birth the LGN has already developed from an unordered state into a near adult form, and so we are only concerned with prenatal development. Furthermore, the retinal photoreceptors are not yet functioning, so processes relying on patterns of visually driven activity are ruled out. However, although there is no visually induced activity, neighbouring retinal ganglion cells are spontaneously active in the prenatal retina, producing "waves of activity" across the retina. These waves have been hypothesised to provide the correlations in activity of neighbouring retinal ganglion cells that drive hebbian learning rules to develop the patterns of ocular segregation and topography in the LGN [4].

To test this hypothesis, an initial computer model of the development of topography and ocular segregation was developed [3]. In this paper, the model is further developed, investigating in particular the normalisation and learning rule.

1.1 Related Work

Several models of topographic development in the retinotectal pathways have been created (for example [6]). These are not directly relevant since they only consider the development of point to point topographic mappings, rather than the development of projection columns as found in the LGN. They also do not consider the development of ocular segregation. Models have been presented showing the development of ocular dominance patches in the cortex, (eg [5]), but these assume a certain correlation between input cells which can be quite restrictive. For example, the model by Miller breaks down if positive between eye correlations are introduced. Furthermore, it did not consider the development of topography. A model of the joint development of ocular dominance and topography was developed by [1]. This model used a competitive learning rule, and produced ocular stripes even when the eyes were positively correlated, unlike Miller's model. However, the nature of the topography and ocular segregation is sufficiently different in the LGN than in the cortex, and so cannot be explained by these models alone.

2 The Model

Figure 1 describes the model used in the simulations. At each time step, retinal activity is propagated to the geniculate cells, and the weights adjusted with the covariance learning rule (with $\alpha = 0.1, \beta = 0.0125, \varepsilon = 0.01$):

$$\Delta w_{ij} = \varepsilon(r_i - \alpha)(o_j - \beta) \quad \text{subject to each weight } w_{ij} \geq 0$$

To prevent weights growing without bounds, the original model employed two forms of normalisation. Firstly, all nl (80) weights from each retinal cell are normalised divisively to sum to tR. Secondly, all nr (100) weights to each geniculate

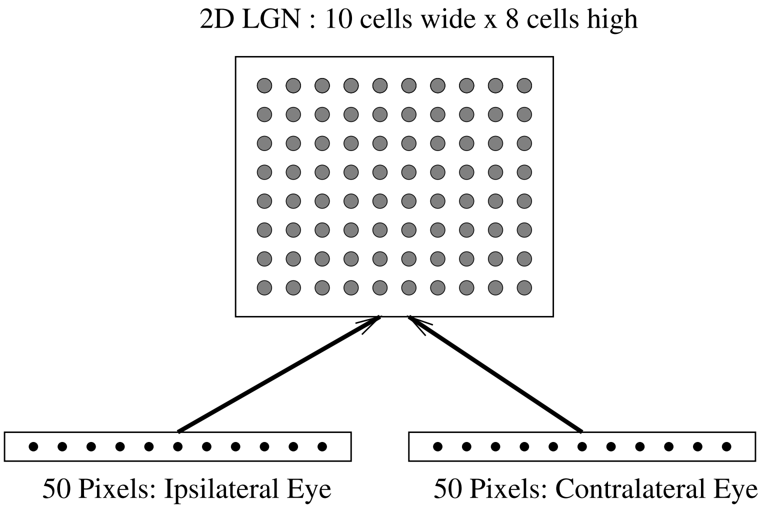

Figure 1: Model architecture. Each retina is one dimensional and fully connects to the two dimensional coronal slice of the LGN via weighted connections. Activity of the retinal and geniculate cells is denoted by r_i and o_j respectively. The strength of the connection between retinal cell i and output cell j is given by w_{ij}. The activation level of output cell j is given by $o_j = \sum_{i=1}^{nr} w_{ij} r_i$.

cell are normalised subtractively to sum to tL:

$$w'_{ij} = \frac{w_{ij} \cdot tr}{\sum_{k=1}^{nl} w_{ik}} \qquad \text{Retinal / Presynaptic Normalisation}$$

$$w''_{ij} = w'_{ij} - \frac{1}{nr}(tl - \sum_{k=1}^{nl} w'_{kj}) \qquad \text{Geniculate / Postsynaptic Normalisation}$$

Finally, the model also uses a growth rule, to encourage a retinal cell to increase the synaptic strength to neighbouring LGN cells. Specifically, if a retinal cell i synapses onto several neighbouring cells of a LGN cell j, the weight w_{ij} should be increased by some fraction:

$$\Delta w_{ij} = \gamma \sum_{k=\text{neigh}(j)} w_{ik}$$

The neighbours of a cell are those within a certain distance of the cell. To avoid border effects, the cells at the left and right edges of the LGN are considered to be neighbours.

Initial Weights. The weights are initially given random values, subject to two forms of bias which represent the initial state of the retinogeniculate pathway. Firstly, the contralateral eye innervates the cat LGN a few days earlier than the ipsilateral eye. This is modelled by allowing the contralateral eye to innervate all eight layers of the LGN, whereas the ipsilateral eye innervates only the bottom 4 layers. Secondly, we impose a topographic bias onto the top layer of weights from each eye, so that the left end of the retina maps onto the left end of each layer. This form of bias is common in topographic models, otherwise there is no global mechanism to specify the overall orientation of the mapping [6].

Nature of inputs. Here we assume that the two eyes are independently generating waves of activity that travel across the retina. At each time step, if there is no wave present, a wave can start at either end of the retina with probability p. Once a wave is started, it travels across to the other end of the retina, with the centre of the wave moving on one pixel at every time step. The wave is modelled here as a Gaussian ($\sigma = 0.85$). Once a wave finishes, the eye goes into a refractory period (normally one time step), before being able to generate another wave. This means that there can be at most one wave in each retina at all times.

Learning Paradigm. At each time step, the activity of each retina is updated either by moving the wave to the next position or seeing if a new wave should be generated. The activity of each geniculate cell is then calculated, and the learning rule applied to each weight in the network. Additionally, every 100 time steps or so, the growth learning and normalisation are applied. Learning typically continues for around 50,000 time steps, during which time the neighbourhood of growth learning slowly decreases as the map refines.

3 Results

Figure 2 shows the normal development of the system. These pictures verify the results given in the original work [3]. Here we see that the geniculate orders itself such that the top four layers respond to the contralateral eye and the bottom four layers respond to the ipsilateral eye. We also see from the topographic pictures that each layer has a smooth topographic map of the retina. In particular, we notice that each retinal cell projects to a column of four cells in the geniculate.

3.1 Effect of normalisation

Previous work on ocular dominance and topography has shown that divisive and subtractive normalisation have different properties [5, 1]. Both models required postsynaptic normalisation for normal development. Furthermore the type of normalisation was dependent on the nature of the correlations between the eyes: Miller showed that divisive postsynaptic normalisation could only be used if the eyes are anti correlated. If there are positive between-eye correlations, then the results of [1] show that postsynaptic normalisation should be subtractive and that a competitive rather than correlational rule should be used. To test whether the form of postsynaptic normalisation was crucial for development, nine different experiments were performed, varying the type of normalisation at each site: both presynaptic and postsynaptic normalisation could be either divisive, subtractive, or none. The results of these nine experiments are shown in Figure 3.

From this figure, two main points are raised. Firstly, we notice from the left hand column of weight diagrams that the form of postsynaptic normalisation does not matter – in all 3 cases, the weights evolve to almost the same normal pattern.

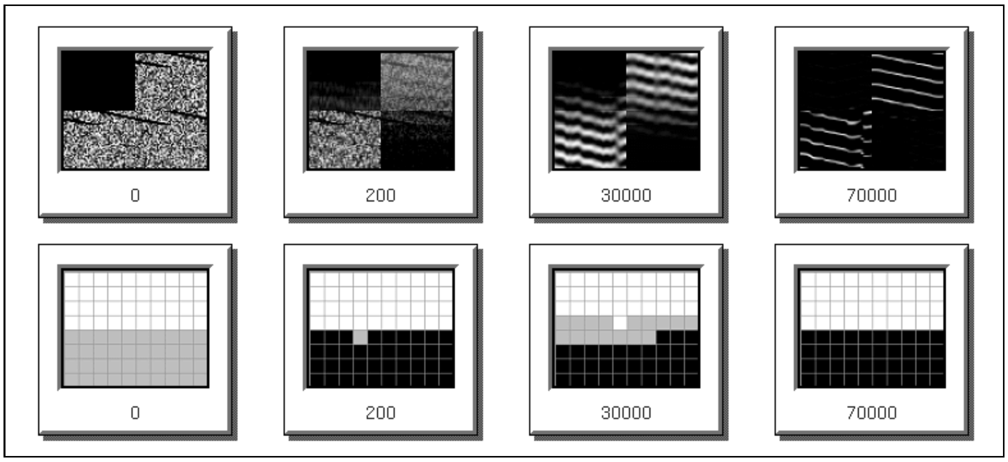

Figure 2: Normal development of the geniculate. Top row: Development of topography over time. Each geniculate cell's weight vector is displayed as one row of the raster image, where intensity of the pixel is proportional to the weight strength. Row n of the image corresponds to the weight vector for geniculate cell n at grid location (x, y), where $n = (10 \times y) + x$. The number in the frame of each picture corresponds to the time step of the simulation. The left hand picture shows the initial state of the weights before learning begins, showing the ocular bias and the topographic bias for rows four and eight of the geniculate. A smooth retinotopic mapping can be seen to develop for each of the four layers for each eye: neighbouring geniculate cells receive input from neighbouring retinal cells. Bottom row: Development of ocular dominance over time. Each cell in the 10×8 geniculate is coloured according to its inputs: Cells receiving at least 80% input from the contralateral (ipsilateral) eye are coloured white (black). Other cells are binocular and are coloured grey. Cells very quickly become monocular, and then for a long while there is a binocular region in the middle of the LGN. Eventually the system stabilises so that each eye innervates four rows of the LGN.

Secondly, we notice that the form of presynaptic normalisation is crucial for development: only if the presynaptic normalisation is divisive will the normal topography and ocular segregation develop. Hence, as long as the presynaptic normalisation is divisive, the model develops normally – postsynaptic normalisation of any form is redundant. When there is no presynaptic normalisation (pictures c,f,i of Figure 3), we find that most of the retinal cells disconnect from the geniculate, as there is no constraint to keep them connected.

3.2 Temporal Nature of Retinal Inputs

Earlier results from Miller [5] indicated that postsynaptic subtractive normalisation is necessary when the two retinae are not anti-correlated. The correlations between eyes in our model is a function of the probability of the waves firing. We therefore explored a range of values for p to see if postsynaptic normalisation was ever needed once p exceeded a certain threshold.

The top diagram in Figure 4 shows the results of using four different values of p for the waves. (Previous experiments here used $p = 0.020$.) Firstly, when p is very low, most of the time the geniculate is receiving no input, and so most of

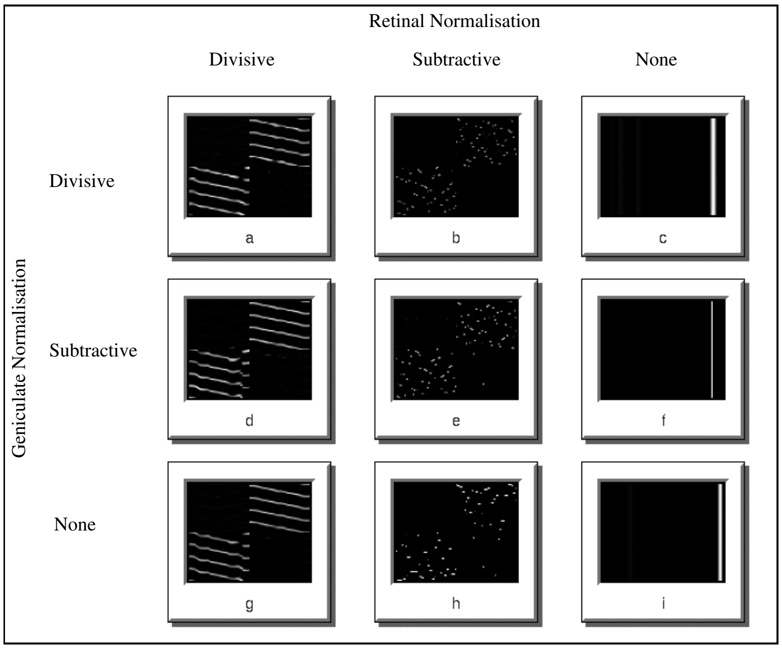

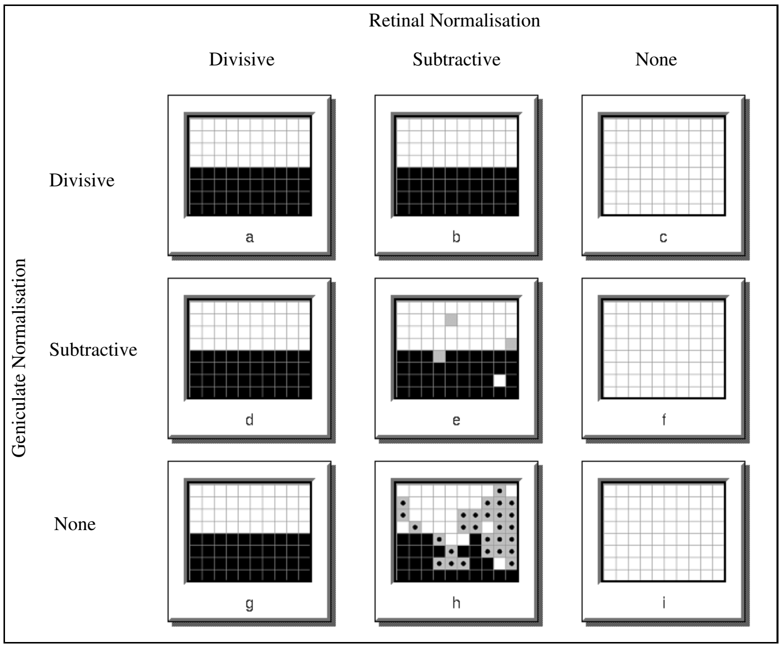

Figure 3: The effect of different normalisation schemes. See Figure 2 for description of visualisation methods. Grey cells with black circles in the ocular dominance plots indicate that the geniculate cell is dead and receives no input from either retina. Note that in the last case (i), there is no normalisation at either site, and so weights grow without bounds.

the time there is no output. When both input r_i and output o_j are low, the covariance rule causes the weight w_{ij} to increase. Hence, when p is low, most of the time the weights are being increased non selectively. The normalisation and growth mechanisms are then responsible for averaging out the total weight strength evenly amongst all weights. Secondly, we notice that as p increases, the topography and ocular segregation begin to break up, but still the majority of cells are monocular. So, normal development relies on a limited range of values of p, with topography and ocular segregation breaking up outside this range. To date, the exact nature of the correlations between the waves is unknown, but experimental data so far suggests that the waves are uncorrelated [7], and that p should be fairly low.

To see if normal topography and ocularity could develop when p is very small, we adapted the covariance learning rule so that no learning took place when both input and output were less than their thresholds. Results are shown in the bottom picture of Figure 4 and show that normal topography can develop when the eyes are not firing very often.

4 Conclusions

Here we have developed a model of ocular segregation and topography in the LGN originally created by [3]. Firstly we have shown that the presynaptic normalisation is an important component of the model, and that the postsynaptic normalisation scheme, at least using the retinal inputs we have presented here, is redundant. This compares against two popular models of ocular dominance, which both relied on postsynaptic normalisation [5, 1]. It is likely however that this result is due to the nature of the retinal inputs and the type of ocular segregation required (patches versus layers), and that postsynaptic normalisation is still required in the other models.

Secondly, we have shown that when the waves are not present in the retina very often, the standard covariance rule needs modifying to prevent it incrementing weights when both the inputs and outputs are below threshold. This has been implemented simply by not doing any learning on a weight when its input and output are below their thresholds. An alternative scheme could be to adapt the input and output thresholds (α, β) according to some temporal average of their past values to automatically handle this situation. However, initial experiments adapting the thresholds have only been of limited success.

Work is currently underway using two dimensional retinae and a three dimensional geniculate to see if the ideas presented here scale up to the real dimensions of the visual pathway. We are also investigating the nature of correlations between ON and OFF cells that can cause sublayers within each eye-specific layer to develop.

References

[1] G. J. Goodhill. Topography and ocular dominance – a model exploring positive correlations. *Biological Cybernetics*, 69:109–118, 1993.

[2] C. S. Goodman and C. J. Shatz. Developmental mechanisms that generate precise patterns of neuronal connectivity. *Cell*, 72:77–98, 1993. Joint Supplement with Neuron, Vol 10.

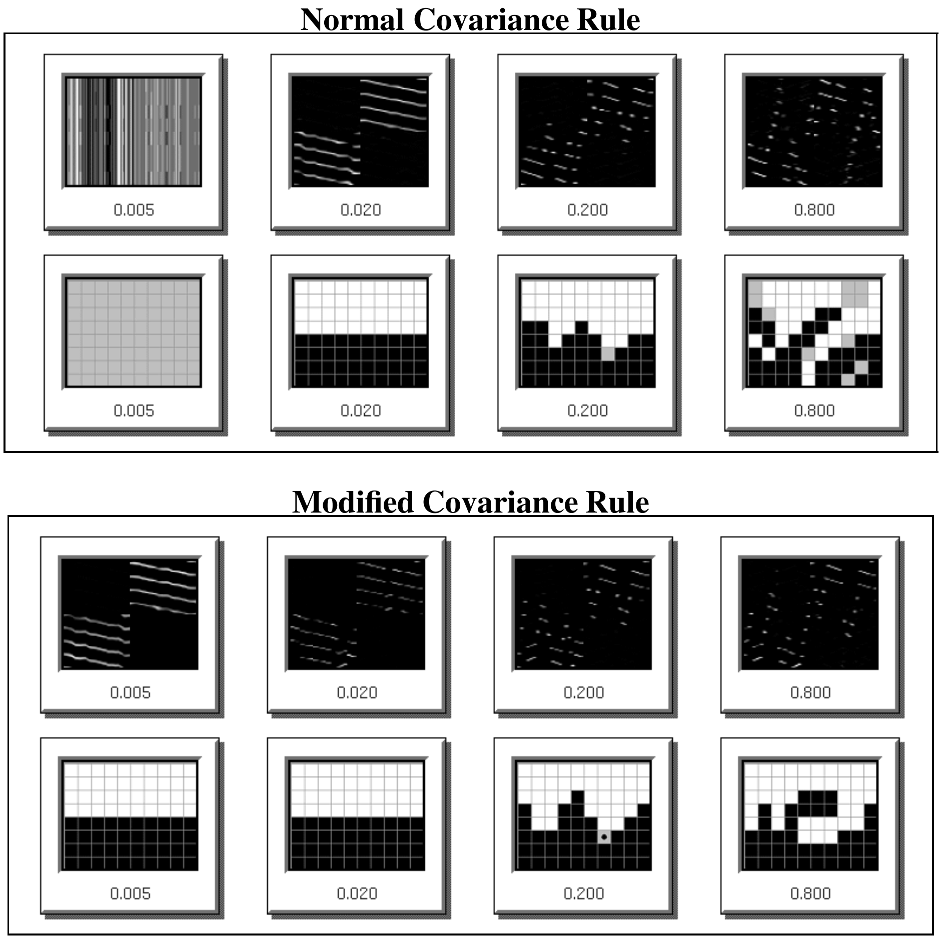

Figure 4: Effect of varying p on development. Top picture: Development with normal covariance rule. Bottom picture: Development with the modified covariance rule (no learning when input and output are less than their thresholds). For each value of p, (from left to right: 0.005, 0.020, 0.200, 0.800), the weight diagram and corresponding ocular dominance plot is shown.

[3] R. Keesing, D. G. Stork, and C. J. Shatz. Retinogeniculate development: the role of competition and correlated retinal activity. In J. E. Moody, S. J. Hanson, and R. P. Lippmann, editors, *Advances in Neural Information Processing Systems*, volume 4, pages 91–97. Morgan Kaufmann, 1992.

[4] M. Meister, R. O. L. Wong, D. A. Baylor, and C. J. Shatz. Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina. *Science*, 252:939–943, 1991.

[5] K. D. Miller, J. B. Keller, and M. P. Stryker. Ocular dominance column development — analysis and simulation. *Science*, 245(4918):605–615, 1989.

[6] D. J. Willshaw and C. Von der Malsburg. A marker induction mechanism for the establishment of ordered neural mappings: Its application to the retinotectal problem. *Philosophical Transansctions of the Royal Society, London B*, 287(1021):203–243, 1979.

[7] R. O. L. Wong, M. Meister, and C. J. Shatz. Transient period of correlated bursting activity during development of the mammalian retina. *Neuron*, 11(5):923–938, 1993.