

Modelling the development of the retinogeniculate pathway

Stephen Eglen

CSRP 467

July, 1997

ISSN 1350-3162

UNIVERSITY OF



SUSSEX
AT BRIGHTON

Cognitive Science
Research Papers

Acknowledgements

Thanks to all who have helped me to get this far. First, I would like to thank both of my supervisors, Harry Barrow and David Young, for their support throughout the thesis. Inspiration for this work has also come from Alistair Bray, Julian Budd, Geoff Goodhill, Jim Stone and David Willshaw. I would particularly like to thank Julian for his constant support in all manner of ways, without which I would have faced a much more difficult task. Both my examiners, Peter Földiák and Hilary Buxton, also provided very useful feedback on the thesis.

Many fellow students have helped me on my way. First, thanks to Jo Brook for all her help with \LaTeX and postscript. Guillaume Barreau, Clive Cox, Philip Jones, Jason Noble, Sara Parsowith and Rosemary Tate have all made the many hours in the lab so much more enjoyable.

For their emotional support, I would like to thank Nosheen Choudhary, Esther Leung and Rafael Pérez Y Pérez. Bernie Lee, Tunde Olatunji and Nick Walker should all be thanked for pushing me in the early days. Most of all, thanks to Valéria Judice and my family.

Financial support for this thesis was provided by the School of Cognitive and Computing Sciences, for which I am grateful.

Modelling the development of the retinogeniculate pathway

Stephen Eglen

Summary

How does the visual system develop before the onset of visually-driven activity? By the time photoreceptors can respond to visual stimulation, some pathways, including the retinogeniculate pathway, have already reached a near-adult form. This rules out visually-driven activity guiding pathway development. During this period however, spontaneous waves of activity travel across the retina, correlating the activity of neighbouring retinal cells. Activity-dependent mechanisms can exploit these correlations to guide retinogeniculate refinement. In this thesis I investigate, by means of computer simulation, the role of spontaneous retinal activity upon the development of ocular dominance and topography in the retinogeniculate pathway.

Keesing, Stork, and Shatz (1992) produced an initial model of retinogeniculate development driven by retinal waves. In this thesis, in addition to replicating their initial results, several new results are presented. First, the importance of presynaptic normalisation is highlighted. This is in contrast to most previous work on ocular dominance requiring postsynaptic normalisation. Second, the covariance rule is adapted so that development can occur under conditions of sparse input activity. Third, the model is shown to replicate development under conditions of monocular deprivation. Fourth, model development is analysed using different spatio-temporal inputs including anticorrelations between on- and off-centre retinal units.

The layered pattern of ocular dominance in the LGN is quite different to the stripe patterns found in the cortex. The factors controlling the patterns of ocular dominance are investigated using a feature-based model of map formation (Obermayer, Ritter, & Schulten, 1991). In common with other models, variance of the ocularity feature controls the pattern of stripes. The model is extended to a three-dimensional output array to show that ocular dominance layers form in this model, and that the retinotopic maps are organised into projection columns. Future work involves extending this three-dimensional model to receive retinal-based, rather than feature-based, inputs.

Submitted for the degree of D. Phil.

University of Sussex

July, 1997

Contents

1	Introduction	1
1.1	The problem of visual system development	1
1.1.1	The approach	1
1.2	Why consider the retinogeniculate pathway?	2
1.3	The importance of modelling	2
1.3.1	Testing hypotheses	3
1.3.2	Generating hypotheses	3
1.4	Aims of the thesis	4
1.5	Outline of the thesis	4
2	Development of the retinogeniculate pathway	6
2.1	The mammalian visual pathway	6
2.2	The retina	7
2.2.1	Formation of RGC receptive fields	7
2.3	The LGN	8
2.3.1	LGN cell types	9
2.3.2	Receptive fields of LGN relay cells	10
2.3.3	The retinotopic mapping in the LGN	10
2.3.4	Ocular segregation in the LGN	11
2.3.5	Sublayering in the LGN	12
2.4	The functions of the LGN	12
2.4.1	The corticogeniculate pathway	13
2.4.2	Other theories of LGN function	15
2.5	Factors influencing neural development	15
2.5.1	Marker theories and genetic specification	15
2.5.2	The role of experience in development	17
2.6	Development of the LGN	18
2.6.1	Normal development	18
2.7	Neural activity in the developing retina	19
2.7.1	Waves of spontaneous activity	20
2.8	The central hypothesis	21
2.8.1	Propagation of retinal waves to the LGN	21
2.8.2	Abnormal development of the LGN	22
2.8.3	Activity blockade experiments	23
2.8.4	Retinotopic map formation	24
2.8.5	The need for activity-independent processes	25

2.9	Summary	26
3	Mechanisms for modelling retinogeniculate development	27
3.1	Introduction	27
3.2	Characterisation of topography and ocular dominance	27
3.2.1	Topography	27
3.2.2	Ocularity	28
3.3	An overview of common mechanisms in developmental models	29
3.3.1	Network architecture	29
3.3.2	Unit activations	30
3.3.3	Common mechanisms underlying models	30
3.4	Correlated input	31
3.4.1	Chemical markers	31
3.4.2	Correlated presynaptic activity	32
3.4.3	Discriminating input from different eyes	32
3.4.4	Representing retinal activity within a model	32
3.5	Correlated output	34
3.6	Normalisation methods and resource limitations	35
3.6.1	Divisive and subtractive normalisation	36
3.6.2	Biological evidence for normalisation at different sites	37
3.7	Weight modification rules	38
3.7.1	Correlational rules	38
3.7.2	Competitive rules	41
3.8	Modelling the retinogeniculate pathway	42
3.8.1	Models of the LGN	43
3.8.2	Related adaptive models	44
3.9	Models of topography	44
3.10	Models of ocular dominance	46
3.11	Models of the joint development of topography and ocular dominance	51
3.12	Discussion	53
3.12.1	Inputs	53
3.12.2	Outputs	53
3.12.3	Modification rules	53
3.12.4	The effect of different weight modification rules and normalisation schemes on topography and ocular dominance	53
4	An initial model of the retinogeniculate pathway	55
4.1	Introduction	55
4.2	The initial model	55
4.2.1	Weight adaptation mechanisms	56
4.3	Retinal inputs	57
4.4	Results	58
4.4.1	Visualisation of results	58

4.5	The role of initial weight biases in development	63
4.5.1	No initial weight bias	63
4.5.2	Earlier arrival of contralateral axons	67
4.5.3	Earlier arrival of contralateral axons and topographic bias	70
4.6	The importance of normalisation	76
4.6.1	The effect of normalisation upon topography	76
4.6.2	The effect of normalisation upon ocular dominance	77
4.6.3	The importance of the ordering of the normalisation techniques	84
4.6.4	Satisfying both forms of normalisation	84
4.6.5	Capping subtractive presynaptic normalisation	87
4.7	Summary	90
5	The influence of the spatial and temporal wave properties upon development	91
5.1	Introduction	91
5.2	Exploring the temporal properties of retinal inputs	91
5.2.1	The relationship between the rate of wave generation and the probability of activity in both eyes	92
5.2.2	The effect of wave generation rate upon ocular dominance and topography	93
5.2.3	Introducing the active-covariance rule	94
5.3	Development of the LGN under conditions of monocular deprivation	99
5.3.1	A weaker form of divisive normalisation	100
5.3.2	The nature of the retinotopic mapping	100
5.4	The effect of wave width upon development	104
5.5	The effect of polarity-specific wave patterns upon LGN sublamination	107
5.5.1	Representation of on- and off-centre retinal units	107
5.5.2	Measuring the degree of polarity dominance	108
5.5.3	Anticorrelated polarity-specific waves	109
5.5.4	Weakening the anticorrelations between on- and off-centre inputs	114
5.6	Comparison with biological data	116
5.6.1	Temporal characteristics of retinal waves	116
5.6.2	Spatial characteristics of retinal waves	116
5.6.3	The nature of correlations between on- and off-centre retinal units	117
5.7	Discussion	118
5.7.1	New mechanisms introduced	118
5.7.2	The polarity experiments	119
6	Factors affecting ocular dominance stripe formation	122
6.1	Introduction	122
6.2	Reasons for ocular dominance stripe formation in models	122
6.2.1	Lateral interactions	122
6.2.2	Competition between topography and ocularity	123
6.2.3	Primary/secondary feature argument	123
6.3	Obermayer's model for stripe formation	124

6.3.1	Implementation details of Obermayer experiments	124
6.3.2	Visualisation of maps formed in Obermayer experiments	124
6.3.3	How the variance of ocularity affects stripe width	124
6.4	A new model of retinotopic map formation in the LGN	129
6.5	Correspondence with topography and ocularity in the LGN	129
6.6	Discussion	136
7	Conclusions	138
7.1	Discussion of the main results	138
7.1.1	Replication of Keesing model results	138
7.1.2	Normalisation	138
7.1.3	Spatio-temporal retinal wave properties	139
7.1.4	Monocular deprivation	139
7.1.5	Segregation of on- and off-centre inputs	139
7.1.6	Models of ocular dominance stripe formation	139
7.2	Future work	140
7.2.1	Modelling of the retina	140
7.2.2	Scaling up the network	140
7.2.3	Polarity segregation	140
7.2.4	Temporal correlations	141
7.3	Outstanding biological issues	141
7.3.1	Retinal waves	141
7.3.2	Development of topography	141
7.3.3	Development of on- and off-centre units	142
7.3.4	Propagation of waves to other areas	142
	Bibliography	143
A	Mathematical details	157
A.1	The derivation of $\frac{dw}{dt} = \mathbf{C}w$ from the covariance rule	157
A.2	Why eigenvectors dominate development in correlational-based modification rules	158
A.3	Implementation of subtractive normalisation	159

List of Figures

2.1	The structure of the retina.	7
2.2	The visual pathway from the retina to the geniculate.	9
2.3	The topography and layering of the cat LGN — parasagittal and coronal views.	11
2.4	The sublayering of on- and off-centre cells in the ferret LGN.	13
2.5	Summary of the main inputs to cat LGN X relay cells and their approximate magnitudes.	14
2.6	The retinotopic mapping of the left retina onto the right optic tectum.	16
2.7	The retinotectal mapping following removal of the caudal half of the tectum.	16
3.1	Common forms of topographic mappings.	28
3.2	Common patterns of ocular dominance.	29
3.3	The definition of presynaptic and postsynaptic normalisation within a two-layer network.	36
3.4	A simple comparison of divisive and subtractive normalisation.	37
3.5	Covariance modification rules.	40
4.1	Architecture of the Keesing network.	56
4.2	Structure of the weight matrix.	59
4.3	The refinement of receptive field width for a single postsynaptic unit.	62
4.4	Network development with no initial weight bias – 1: topography and ocularity.	65
4.5	Network development with no initial weight bias – 2: refinement of receptive fields, monocularly and final projection columns.	66
4.6	The effect of contralateral bias in initial weights – 1: topography and ocularity.	68
4.7	The effect of contralateral bias in initial weights – 2: refinement of receptive fields, monocularly and final projection columns.	69
4.8	The nature of the topographic bias in the initial weights of one row of LGN units.	71
4.9	The effect of ocular and topographic weight bias on development – 1: left eye innervates two rows, right eye innervates four rows (L2R4).	73
4.10	The effect of ocular and topographic weight bias on development – 2: left eye innervates four rows, right eye innervates eight rows (L4R8).	74
4.11	Comparison of the L2R4 and L4R8 experiments with respect to the development of ocularity and receptive field width, and the final projection columns.	75
4.12	The effect of different normalisation methods – 1: weight matrices.	78
4.13	The effect of different normalisation methods – 2: ocular dominance plots.	79
4.14	The effect of different normalisation methods – 3: topographic plots.	80
4.15	The effect of different normalisation methods – 4: projection columns.	81

4.16	Sum of weights to each postsynaptic unit in a network with divisive presynaptic normalisation and no postsynaptic normalisation.	82
4.17	Projective field for one presynaptic unit from the right eye using either divisive (solid line) or subtractive (dashed line) presynaptic normalisation.	82
4.18	The effect of the growth rule upon development in the absence of presynaptic normalisation.	83
4.19	Effect of varying the order of presynaptic and postsynaptic normalisation.	86
4.20	Sum of weights from each presynaptic unit under relaxed normalisation conditions.	88
4.21	The effect of imposing maximum values on individual weights upon development with subtractive presynaptic normalisation (and no postsynaptic normalisation).	89
5.1	Overall probabilities of joint retinal activity as a function of the probability of wave generation.	93
5.2	Correlation between presynaptic units as a function of the rate of wave generation.	95
5.3	Effect of varying p_w upon topography and ocular dominance – 1: development with the covariance rule.	97
5.4	Effect of varying p_w upon topography and ocular dominance – 2: development with the active-covariance rule.	98
5.5	Network development with reduced activity levels for the left eye – 1: development with multiplicative and divisive presynaptic normalisation.	101
5.6	Network development with reduced activity levels for the left eye – 2: development with only divisive presynaptic normalisation, ignoring multiplicative normalisation.	102
5.7	Final projection columns for a network with no waves generated within the left eye.	103
5.8	Correlations between presynaptic units as a function of wave width.	105
5.9	The effect of wave width upon topography and ocular dominance.	106
5.10	The effect of wave width upon receptive field width.	107
5.11	Slices of the input correlation matrix before and after the onset of polarity-specific wave patterns.	111
5.12	Time course of polarity dominance using anticorrelated on- and off-centre waves.	112
5.13	Development of polarity dominance in postsynaptic units.	113
5.14	Final projection columns for the network with polarity-specific waves as input.	113
5.15	Polarity segregation for different correlations between on- and off-centre presynaptic units.	115
5.16	Chemical gradients to bias the formation of on- and off-centre sublamina in the LGN.	121
6.1	Final topography and ocularity maps in a two-dimensional Kohonen network with feature vectors $(x, y, \pm z_{pat})$ as input and $z_{pat} = 0.2, 0.6, 1.0$	127
6.2	Final topography and ocularity maps in a two-dimensional Kohonen network with feature vectors $(x, y, \pm z_{pat})$ as input and $z_{pat} = 2.0$	128
6.3	Three-dimensional arrangement of postsynaptic units in the Kohonen network.	130

6.4	Final topography and ocularity plots for three-dimensional feature vector (x, y, z) mapping into a three-dimensional postsynaptic block for $z_{pat} = 0.2$ and 2.0	132
6.5	Final topography plots analysed in the X-planes for a three-dimensional network with $z_{pat} = 0.2$ and $Z_{post} = 4$	133
6.6	Final topography plots analysed in X-planes and Y-planes for a three-dimensional network with $z_{pat} = 2.0$ and $Z_{post} = 4$	134
6.7	Final topography and ocularity maps in a three-dimensional network with $z_{pat} = 1.0$ and $Z_{post} = 3$	135

List of Tables

2.1	Class of retinal input to each layer of the cat LGN.	12
2.2	Timescale of cat retinogeniculate development.	19
2.3	Timescale of ferret retinogeniculate development.	20
3.1	Terms used for describing model elements.	30
3.2	The four components of the covariance rule.	38
3.3	Summary of the main models of retinotopic map development.	47
3.4	Summary of the main models of ocular dominance development.	50
3.5	Summary of the main models of the joint development of topography and ocular dominance.	52
4.1	Parameters and typical values used in the Keesing network simulations.	58
4.2	Normalisation error measures E_{pre} and E_{post} for different values of p_n at different epochs during development.	85
5.1	Theoretical probability of eye activity as a function of p_w	94
5.2	Mean receptive field width for different values of p_w using either the covariance or active-covariance rule.	96
5.3	The frequency of use and the amount of weight change per epoch for each case of the covariance and active-covariance rule.	96
5.4	Number of postsynaptic units responsive to the left, right or neither eye at the end of development under reduced activity conditions for the left eye.	103
5.5	Rules for generating polarity-specific waves.	108
5.6	Parameters for the polarity-specific wave experiments.	108
5.7	Burst and inter-burst intervals at various postnatal stages in the ferret.	117
6.1	Network equations and parameters used for Obermayer experiments.	125
6.2	Mean values of z for the two- and three-dimensional Obermayer experiments.	131

Chapter 1

Introduction

1.1 The problem of visual system development

How do mammalian visual pathways develop in the absence of patterned vision? By the time cats are born, the retinogeniculate pathway, the pathway between the retina and the lateral geniculate nucleus (LGN), has already developed into a near adult form. It is possible that each retinal cell is told exactly which LGN cells it should connect to by some genetic plan (for example, the chemospecificity hypothesis (Sperry, 1963)). However, given the large number of retinal and geniculate cells, it is more likely that some other more general guiding principles are at work (von der Malsburg & Singer, 1988). One such principle is that neural activity generated by visual stimulation drives development (Blakemore & Cooper, 1970; Hirsch & Spinelli, 1970). At this stage however, the photoreceptors in the retina are still developing and the eyelids remain closed until around ten days after birth, ruling out any influence of patterned vision on the development of the retinogeniculate pathway.

Although the retina is not yet converting visual stimulation into electrical activity, ganglion cells within the retina are spontaneously active, producing waves of activity across the ganglion cell layer (Galli & Maffei, 1988; Meister, Wong, Baylor, & Shatz, 1991). A popular hypothesis is that these waves of activity provide local correlations which are used by some form of Hebbian synaptic modification process to refine the retinogeniculate pathway during prenatal development (Maffei & Galli, 1990; Wong, Meister, & Shatz, 1993). This hypothesis is the subject of this thesis.

1.1.1 The approach

The approach of this thesis is to investigate the processes of development in the retinogeniculate pathway before the onset of patterned vision by building a simplified computer model of the retina, the LGN and the axons connecting the retina to the LGN. This model encapsulates the retinal waves of activity and a Hebbian-based self-organising process to test if retinal waves can refine the retinogeniculate pathway as has been suggested by various experimenters (Maffei & Galli, 1990; Wong et al., 1993). The two aspects of the pathway which are investigated here are the development of retinotopic mappings in the LGN and the segregation of the LGN into eye-specific

layers:

1. Retinotopic mappings — neighbouring retinal cells tend to connect to neighbouring geniculate cells. Initially this mapping is quite coarse, and refines during development.
2. Ocular segregation — LGN cells initially receive inputs from both eyes before becoming selective to only one eye. Cells responding to the same eye are grouped into eye-specific layers within the LGN.

1.2 Why consider the retinogeniculate pathway?

There are many visual pathways in the mammalian brain. There are several reasons why this thesis focuses on just the retinogeniculate pathway:

- Although the retinogeniculate pathway is one of the simplest visual pathways, it is highly likely that there are general principles of development at work. By studying these principles in a simple pathway, it is hoped that they might be applicable to other more complicated pathways. For example, it is suggested that the retinal waves of activity generate action potentials in LGN cells which could in turn drive geniculocortical development (Penn, Gallego, Mooney, & Shatz, 1995).
- There is a considerable amount of biological literature on the retinogeniculate pathway in several different mammalian species. This is because the pathway is relatively simple and easy to access in comparison with other (mostly cortical) pathways.
- Much of the retinogeniculate development occurs before the onset of patterned vision. Most previous modelling studies have not explicitly considered this stage, either assuming inputs are spontaneous (Willshaw & von der Malsburg, 1976), Gaussian correlated (Goodhill, 1992) or natural images (Barrow, 1987).
- The projection of visual space from the retina to the LGN is different to the projections found in other commonly modelled pathways, such as the geniculocortical pathway. The mapping from the retina to the LGN is a point to column map, whereas the mapping from LGN to cortex is normally considered as a point to point map.
- There is no feedback pathway from the LGN to the retina. Most other pathways (especially cortical pathways) include some element of feedback which makes the modelling effort more complicated. Additionally, although there is inhibitory lateral connectivity in the adult LGN, it is not thought to develop until after the feedforward pathway and so can be ignored (Shatz & Kirkwood, 1984).

1.3 The importance of modelling

Models have two main roles in neuroscience — to test existing hypotheses that neuroscientists have proposed, and to generate new hypotheses that can be taken back to the neuroscientists and investigated in the natural systems.

1.3.1 Testing hypotheses

Sometimes a hypothesis can be written down in a concise mathematical form which describes the natural system in a fairly detailed manner. The mathematical model can then be tested to see if it generates the behaviour that is seen in the natural system. Success is normally measured by the similarity between the natural system's and the model's behaviour. An example is the model by Hodgkin and Huxley (1952) showing how action potentials in a squid giant axon are generated. Quite often the mathematical models are fairly complicated and so they are translated into a computer model to allow rigorous testing.

However, it is more often the case that the hypothesis cannot be directly translated into a set of mathematical statements, due to a lack of detail in the hypothesis. In this case, the modeller has to make certain assumptions to build a complete working model. Hence, building a model forces someone to think about all of the details of a hypothesis, rather than just the central components.

During this process, the modeller may have to introduce different kinds of assumptions. There are those which are regarded as "implementation details", such as needing a certain kind of integration technique to prevent numerical instabilities in the model. These details are of little relevance to the neuroscientist. However, there are other assumptions that may be of more relevance to the neuroscientist in refining a hypothesis. For example, the Hebb rule (Hebb, 1949) states that if a presynaptic cell and a postsynaptic cell are active at the same time, the connection strength between the two should be increased. When this rule is implemented in a computer model it becomes obvious that some other mechanism must be in place to prevent the connection from growing without bounds. In this situation, building a model of an existing hypothesis not only tests the hypothesis by seeing if the model's behaviour matches the natural behavior, it also suggests what extra components may be needed in the hypothesis to make it more realistic.

Finally, models require ranges for parameter values. Quite often these parameters are unknown quantities in the natural system. In this situation, the modeller must search through a range of parameter values in order for the model to work. These parameter values can then be compared with the natural system to see if they are realistic, or if they raise any inconsistencies in the hypothesis.

1.3.2 Generating hypotheses

As well as testing and refining existing hypotheses, computer models can also be used to generate new hypotheses. One advantage of working with a computer model over a natural system is that it is very easy to change the way the model works to test how it behaves under different situations. In comparison, performing the same experiment with a natural system may be difficult or not feasible with current experimental methods, such as lesioning certain pathways or removing a certain mechanism (Sejnowski, Koch, & Churchland, 1988). The results from the model in these situations generate predictions about what should happen in the natural system. These predictions can then be tested by neuroscientists to further develop the hypothesis.

However, to date, there have only been a few examples when computer models have been able to generate predictions which have subsequently been tested by experimenters. Probably the clearest example of these is the recent prediction that decreased between-eye correlations increase the width of ocular dominance stripes in the cortex (Goodhill, 1993). This prediction was later confirmed when examining strabismic cats (Goodhill & Löwel, 1995).

1.4 Aims of the thesis

This thesis aims to cover the following issues:

- How can topography and ocular segregation develop in the LGN? Can it develop by way of just the one mechanism, in the same way as was shown for ocular dominance and topography in the cortex (Goodhill, 1992)?
- How are retinotopic maps affected by the dimensionality of the output structure? Most models only consider a two-dimensional output structure, whereas this thesis is concerned with the three-dimensional structure of the LGN.
- What are the within-eye and between-eye correlations that the model requires for development and are they similar to the correlations that exist in the developing retina?
- Can the segregation of on- and off-centre cells be explained by the same mechanisms as those needed for topography and ocular segregation, assuming certain correlations between on- and off-centre cells?

1.5 Outline of the thesis

Chapter 2 provides a review of the relevant biological literature for this thesis. It contains a brief introduction to the early stages of the mammalian visual pathway, concentrating on the properties of cells in the retina and the LGN. The factors influencing neural development of visual pathways, especially the retinogeniculate pathway, are discussed. The chapter concludes by introducing the central hypothesis of the thesis: spontaneous retinal waves drive development of the retinogeniculate pathway.

Chapter 3 reviews the previous computer models of retinotopic map formation and ocular dominance. Although there are many models which superficially seem quite different, they share many common mechanisms. Despite the large number of previous models, only one model has explicitly examined the role of spontaneous activity in retinogeniculate development (Keesing et al., 1992). A crucial feature missing from all models however is the lack of information on how retinotopic maps can form projection columns similar to those found in the LGN (Sanderson, 1971a).

Chapter 4 describes the initial model of retinogeniculate development produced by Keesing et al. (1992). In the first half of this chapter we replicate the results from this model. Once the basic results have been replicated, the model is then analysed with respect to the role of the various weight biases in the initial network and the form of the normalisation used.

In Chapter 5, we explore the spatio-temporal properties of the retinal waves and how they affect development of ocular dominance and topography. We find that the rate of wave generation can disturb topographic map formation, although the monocular status of units is not affected as much. We introduce a modified covariance rule which can restore normal network development when retinal waves are infrequently generated. The spatial wave properties also affect development: wider waves produce wider receptive fields. Finally, in this chapter we show that polarity

specific waves can cause postsynaptic units to become responsive to either on- or off-centre presynaptic units, but the input correlations required by the model are at odds with the current biological data.

Chapter 6 investigates the forms of ocular dominance patterns that can be produced by cortical models. Most cortical models generate ocular dominance stripes similar to those found in area 17 of the cortex, but can these models also be applied to the pattern of eye-specific layers found in the LGN? In this chapter we replicate the model described by Obermayer et al. (1991) which shows how ocularity variance controls stripe width. When this model is extended to use a three-dimensional arrangement of postsynaptic units, the model reproduces the eye-specific layers for small values of the ocularity variance. If ocularity variance is related to between-eye correlations, these results predict that retinal waves can drive eye-specific segregation such that the eye-specific layers are organised along the shortest axis, the dorso-ventral axis, of the LGN.

Chapter 7 summarises the major conclusions from this thesis:

- Postsynaptic normalisation is redundant in the model of retinogeniculate development. This is in contrast to previous models of ocular dominance. Presynaptic normalisation must be implemented divisively for projection columns to form in the model.
- The spatio-temporal properties of the retinal waves affect development. A new weight modification rule, the active-covariance rule, is created which allows normal network development when the probability of wave generation is low. Using this rule, the network can also account for the effects of monocular deprivation in the LGN.
- The model only predicts segregation of on- and off-centre retinal cells when the cells are anticorrelated. This is at odds with the current biological data on polarity-specific firing in retinal waves.
- The Obermayer et al. (1991) model of ocular dominance that normally produce stripes can also be used to model development of eye-specific layers when postsynaptic units are arranged in a three-dimensional block. Ocularity maps onto the smallest dimension of the postsynaptic block when the variance of ocularity is low.

This chapter also provides directions for further research and mentions some outstanding biological issues that are of relevance to this work.

Chapter 2

Development of the retinogeniculate pathway

This chapter summarises the relevant biological data about the development of the retinogeniculate pathway. It is composed of two main sections. The first section introduces the retinogeniculate pathway, describing its main components and how they function. Since cross-species comparisons are difficult, most of the data is taken either from the cat or ferret, since they share similar developmental timetables. The purpose of the first half of this chapter is to introduce the relevant concepts of the early visual pathway. The second half of the chapter, beginning with Section 2.5, considers the role of both activity-independent and activity-dependent processes in the development of the visual system in general and then the retinogeniculate pathway in particular.

2.1 The mammalian visual pathway

The visual pathway in mammals is composed of around 30 or so different visual areas, each contributing to a different aspect of visual function (Felleman & Van Essen, 1991). The starting point of the mammalian visual pathway is the retina, which converts light into electrical signals. These signals are then passed out of the retina and into two different regions: the LGN and the superior colliculus (SC) (a structure homologous to the optic tectum in other animals). The LGN is located in the thalamus, whereas the SC is located in the midbrain. The LGN and the SC send the processed retinal signals to different regions. Retinal signals to the LGN are destined mainly for cortical areas, whereas signals to the SC are destined for the midbrain. The role of these two different pathways in visual processing is not clear, although it is thought that the cortical pathway is concerned with detailed spatial vision whereas the sub-cortical pathway handles processes of attention and visually guided behaviour.

Although the visual pathway is not strictly hierarchical (Felleman & Van Essen, 1991), there is a progression of complexity of cell responses in the visual areas. In the early parts of the pathway, such as the retina and the LGN, cells have relatively simple receptive field properties, such as the centre-surround response (where cells fire maximally to a light spot on a dark background). In the primary visual cortex (the major recipient of geniculate axons) the cells are much more selective, showing highly tuned responses to stimuli of certain orientations or disparity and wavelength. Higher up in the cortical visual pathway, the complexity of cell response selectivity keeps increas-

ing. For example, around 10% of cells in area IT of the monkey respond maximally to images of either monkey or human faces rather than other stimuli (Perrett, Rolls, & Caan, 1982). In this view, each level of the visual pathway produces receptive field properties that are elaborated on by other areas.

Given the vast size and complexity of the mammalian visual pathway, it is common to study just a small subset of the pathway. The emphasis in this thesis is on retinogeniculate development, and so the next sections will mainly discuss the pathway from the retina to the LGN and up to the cortex. Later in this chapter however, some reference will be made to the sub-cortical pathways, especially the retinotectal pathway.

2.2 The retina

The retina is the starting point of the visual pathway. The retina is a multilayered structure, with both vertical and horizontal pathways connecting together five basic types of retinal cell, as shown in Figure 2.1. Light enters the eye and hits the photoreceptors, where the light is converted into electrical activity. This activity is propagated down (and across) the layers of cells to the final layer of cells, the retinal ganglion cells (RGCs). The RGC axons are then bundled together to form the optic nerve, which carries the visual information into the brain.

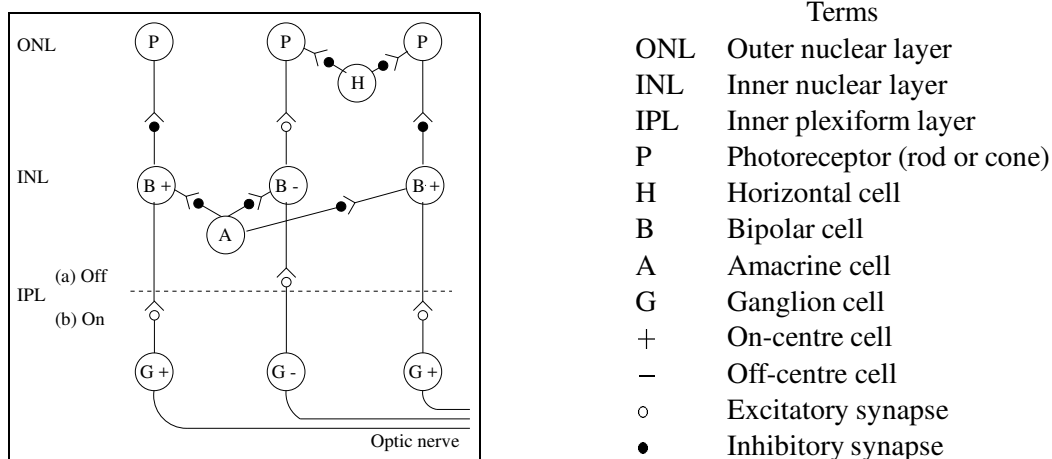


Figure 2.1: The structure of the retina. The retina contains both a vertical pathway (photoreceptors to bipolars to ganglions) and two horizontal pathways (the horizontal cells and the amacrine cells). This diagram does not show all synaptic interactions nor the large degree of convergence of photoreceptors onto bipolars or bipolar cells onto ganglion cells. Figure adapted from Dowling (1987, p118).

2.2.1 Formation of RGC receptive fields

All photoreceptors hyperpolarise (become more negative) in response to visual stimulation. Photoreceptors then converge onto a bipolar cell by either excitatory or inhibitory synapses. If the synapses are all excitatory, then the bipolar cell also hyperpolarises in response to visual stimulation of the photoreceptors. On the other hand, if the synapses are all inhibitory, the bipolar cell will depolarise (become less negative). As well as inputs from photoreceptors, bipolar cells also

receive antagonistic input from horizontal cells, which normally receive synapses from photoreceptors over a wider area.

This combination of direct input from a narrow range of photoreceptors and antagonistic indirect input from a wider range of photoreceptors produces a centre-surround type receptive field profile for the bipolar cells. A bipolar cell with inhibitory input from central photoreceptors is called an on-centre cell since it depolarises in response to stimulation of the central photoreceptors and hyperpolarises in response to stimulation of more peripheral photoreceptors. By similar reasoning, bipolar cells that receive excitatory inputs from central photoreceptors are called off-centre. This convergence of photoreceptors onto bipolars is very high, since over 90% of retinal cells are photoreceptors (Masland, 1996).

The retinal ganglion cells receive their input from two sources: bipolar and amacrine cells. Typically, ganglion cells receive input from a pair of bipolars in a push-pull manner: excitation from one bipolar coincides with removal of inhibition from another bipolar. In this way, ganglion cells are classified as on- or off-centre, according to the centre signs of the excitatory bipolar. The sign of the ganglion cell can also be determined morphologically by examining the dendritic branching point of the ganglion cell: dendrites stratifying in sublamina A of the IPL are off-centre, whereas those dendrites stratifying in sublamina B belong to on-centre ganglion cells. (Nelson, Famiglietti, & Kolb, 1978). The additional input to ganglions from amacrine inputs work in an antagonistic manner, similar to horizontal cells, although the function of amacrine cells is not clear.

As well as classification of RGCs into on- or off-centre, according to the sign of the receptive field, RGCs are further divided into three main categories according to both physiological and anatomical criteria (Stone, 1978, 1983). X cells have medium-sized cell bodies, and comprise around 40% of the ganglion cell population. Y cells have large-sized cell bodies, but make up less than 10% of the number of ganglion cells. The remaining 50% or so of cells have a range of cell-sizes and physiological properties, and are classed as W cells. The different classes of RGC have different responses to the same visual stimuli, meaning that the stimuli are processed in different ways. For example, X cells have a sustained response to stimuli which is useful for fine-detailed spatial processing, whereas Y cells have more transient responses making them much more suitable for motion analysis. The W class of cells is a heterogeneous class, and although the percentage of W cells in cats and ferrets is quite high, comparatively little is known about them since they project mainly to sub-cortical rather than cortical structures.

The retina does not uniformly sample the visual world. The central part of the visual field is sampled at a much higher density than peripheral parts of the visual field. Additionally, the size of receptive fields of retinal cells varies with the eccentricity of the cell — cells in the fovea have small receptive fields, whereas cells in the periphery have large receptive fields.

2.3 The LGN

The LGN is the primary recipient of retinal axons. The pathway from the retina to the LGN is illustrated in Figure 2.2. The optic nerve from each eye meets at the optic chiasm, at which point axons from the nasal part of the retina cross over to the other side of the brain, whereas axons from the temporal retina remain on the same side. (It is thought that molecular markers are responsible

for causing the temporal retinal axons to stay on the same side of the brain, whereas the nasal retinal axons cross over to the other side (Wingate & Thompson, 1995.) The axons leave the optic chiasm to form an optic tract on each side of the brain to innervate the appropriate LGN. In this way, each LGN receives inputs from areas of the two retinae which correspond to the same part of the visual field.

For most animals, the size of the inputs to the LGN from the two retinae are different: the contralateral retinal input (that coming from the retina on the opposite side of the brain to the LGN) is larger than the ipsilateral (same side) input. The difference in size of the contralateral and ipsilateral input is related to the positioning of the animal's eyes on its head, and so varies from species to species. In animals with laterally placed eyes, such as the grey squirrel, the size of the monocular input is very large. In comparison, animals with forward placed eyes (such as cats and monkeys) tend to have a large binocular field and hence the contralateral and ipsilateral inputs are more balanced (Kaas, Guillery, & Allman, 1972), although the contralateral input is still larger than the ipsilateral input.

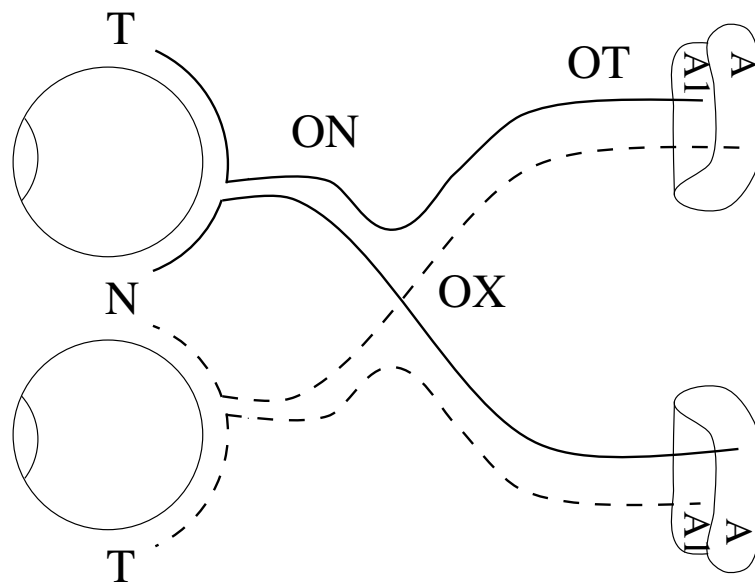


Figure 2.2: The visual pathway from the retina to the geniculate. Key: T – temporal retina. N – nasal retina. ON – optic nerve. OX – optic chiasm. OT – optic tract. A – LGN layer A. A1 – LGN layer A1.

2.3.1 LGN cell types

There are two classes of cells found in the geniculate: relay cells and interneurons. Approximately 75% of LGN cells are relay cells, the other 25% are inhibitory interneurons (LeVay & Ferster, 1979). The relay cells receive excitatory input from between one and five retinal cells (Cleland & Dubin, 1971, 1976). Relay cell axons leave the LGN to form the geniculocortical pathway, connecting the LGN to layer IV of area 17 of the cortex. The interneurons only make local inhibitory connections within the geniculate and do not project to other areas.

Relay cells can be classified into several subtypes. First, they can be labelled on- or off-centre, depending on their centre-surround response to light. Second, they can also be classified as X- or

Y-like, depending on whether they show a null-response to sinusoidal gratings, as an analogue of the retinal X and Y classes (Bowling & Wieniawa-Nakiewicz, 1986).

2.3.2 Receptive fields of LGN relay cells

The spatial receptive field profiles of relay cells are similar to RGCs, although the inhibitory surround region tends to be stronger (Hubel & Wiesel, 1961), probably due to the effect of the inhibitory geniculate interneurons. This similarity of receptive fields is due to the low convergence of RGCs onto relay cells: typically a geniculate cell receives its inputs from between 1–5 retinal cells which tend to be from the same eye and of the same polarity (on- or off-centre) (Mastrorarde, 1987a).

This similarity of retinal and geniculate spatial receptive fields has led to the notion of the geniculate as a “relay station”, transferring retinal information without significant processing to the cortex. For example, geniculate cells show a slight orientation selectivity, but this property originates in the retina, and is formed by the oriented dendritic fields of retinal cells rather than being generated in the retinogeniculate pathway (Levick & Thibos, 1982; Leventhal & Schall, 1983). Similarly, the direction-selectivity found in some geniculate cells (Thompson, Zhou, & Leventhal, 1994) is most likely a reflection of the sensitivity of their retinal inputs (Shou, Leventhal, Thompson, & Zhou, 1995).

In the temporal domain however, geniculate cell responses differ quite strongly from retinal cells. Around two-thirds of geniculate cells exhibit a lagged response to retinal stimuli (Mastrorarde, 1987a, 1987b). This temporal distinction has led to geniculate cells being further categorised as either lagged or non-lagged (Humphrey & Weller, 1988). Although the function of this lagged response is not clear, it could be used to generate various cortical responses such as direction selectivity (Saul & Humphrey, 1990).

Finally, X and Y relay cells can fire in one of two modes: tonic or bursting (Guido, Lu, & Sherman, 1992; Sherman & Guillery, 1996). In tonic firing, the cell responds with multiple action potentials that correspond in a linear fashion to the stimuli. In comparison, in burst mode the response to the stimulus is highly non-linear, comprising of one or two action potentials to stimulus onset and then no further action potentials for the rest of the stimulus presentation. It is thought that the two modes are suitable for different functions. Bursting mode is better suited for the initial detection of stimuli as a way of directing attention to a particular stimulus. Tonic mode however is more suited for faithful transmission of the retinal signal.

The subject of this thesis is the neural pathway from the retina to the LGN, which is commonly termed the retinogeniculate pathway. In particular, this thesis is concerned with the development of topography and ocular segregation in this pathway. These two properties will now be discussed in more detail.

2.3.3 The retinotopic mapping in the LGN

When the optic tract reaches the LGN, the axons innervate the LGN in such a way that neighbouring retinal cells tend to project to neighbouring geniculate cells. This is known as a retinotopic or topographic mapping. Such maps are commonly found not only within visual areas, such as area 17, but also in other sensory areas such as the somatosensory cortex (Kaas, Nelson, Sur, &

Lin, 1979). However, the projection of retinal space into the geniculate is not just a conventional topographic mapping. Most maps project an input space into an output space of either the same or lower dimensionality. In the LGN however, the dimensionality of the spaces is reversed: although the retina is regarded as a two-dimensional sheet of cells, the LGN is a three-dimensional block of cells. The mapping is organised such that a point of visual space in the retina projects to a column of cells within the LGN. This group of LGN cells is termed a ‘projection column’, which is normally defined as a column of LGN cells which contains 90% of all cells with receptive fields (RFs) responding to the same region of visual space (Sanderson, 1971b). Typically, these columns are oriented perpendicular to the LGN layers.

Although cells within a projection column receive input from the same region of visual space, they can respond in different ways to the same stimulus. For example, X relay cells at different depths within a column respond to the same stimuli with different timing latencies. These latencies could be useful for creating certain selectivities of cortical cells, such as direction selectivity (Bowling, 1989a).

The mapping of retinal space into the cat LGN is shown in Figure 2.3. The layout of retinal space in the ferret follows similar retinotopic principles, although the ferret LGN is positioned differently (Zahs & Stryker, 1985). Just as the retina has more cells representing the central visual field, so does the LGN, as can be seen from the amount of geniculate covering the central $\pm 2^\circ$ in comparison with the amount covering $10\text{--}20^\circ$. However, equal numbers of RGCs project to equal volumes in the LGN, and so the gradient of visual field representation in the geniculate is a reflection of RGC density gradient (Sanderson, 1971b).




Figure 2.3: The topography and layering of the cat LGN — parasagittal and coronal views. Parasagittal view adapted from Figure 9 of (Sanderson, 1971a). Coronal view adapted from Figure 14 of (Sanderson, 1971a). The C layers have been omitted here for clarity, but lie just below layer A1. The region of layer A to the right of the 40° isoazimuth line is known as the monocular segment, since there is no corresponding region for layer A1. As can be seen from the lines of isoazimuth and isoelevation, the central visual field is overrepresented in the LGN. These lines of isoelevation and isoazimuth can be regarded as projection columns.

2.3.4 Ocular segregation in the LGN

Both the cat and ferret LGN are divided into three layers — A, A1 and C. (The C layer is commonly divided further into sublayers — C, C1, C2 and C3. When considering the retinogeniculate pathway the C laminae are normally ignored, since it mainly receives input from W ganglion cells and mostly projects onto subcortical structures.) These layers are defined according to several criteria, including its cytoarchitecture (the grouping of cells into cell-rich strata and cell-poor spaces,

and their size distribution), the morphology of individual cells, and by the segregation of the retinal inputs (Casagrande & Brunso-Bechtold, 1985). In this thesis, we are mainly concerned with this last criteria, the segregation of retinal inputs: in the mature LGN, contralateral inputs innervate layer A, whereas ipsilateral inputs innervate layer A1. Details of the segregation of retinal inputs according to eye of origin and class are given in Table 2.1.

LGN layer	Eye of input	Class of retinal cell
A	Contra	X/Y
A1	Ipsi	X/Y
C	Contra	Y
C1	Ipsi	W
C2	Contra	W
C3	Non-retinal	—

Table 2.1: Class of retinal input to each layer of the cat LGN, showing the eye of origin of retinal inputs and physiological class of retinal input (Sherman, 1985). There are no retinal inputs to layer C3 — inputs to this layer come from midbrain structures.

2.3.5 Sublayering in the LGN

Within a geniculate layer, there is evidence for further segregation of retinal inputs according to the class of retinal cell. The strongest form of sublayering is the segregation of on- and off-centre relay cells. In the ferret, both layers A and A1 divide into on and off leaflets (Stryker & Zahs, 1983), as shown in Figure 2.4. In the cat, the sublayering is not so strict, showing more of a tendency for on-centre cells to be found dorsally within a layer and off-centre cells to be found ventrally (Bowling & Caverhill, 1989; Thurlow, Bowling, & Cooper, 1993). It has also been proposed that the on- and off-centre cells are organised into columns perpendicular to the layers (Berman, Naporn, & Payne, 1987; Berman & Payne, 1989). Later studies however could not verify such an organisation, suggesting that the initial results were a product of the multi-unit recording technique (Bowling & Caverhill, 1989).

As well as cells clustering according to centre sign, X and Y RGCs also tend to arborise differently in the cat, with around 60% of X cell boutons clustering at the top of a layer, and about 70% of Y cell boutons clustering in the lower half of a layer (Bowling & Michael, 1984) — although see Sur, Esguerra, Garraghty, Kritzer, and Sherman (1987) for some qualifications. This clustering of boutons produces a slightly unexpected organisation of geniculate cells: geniculate Y-like cells tend to be found at the borders of the layers, with X-like cells more concentrated in the centre of the layers (Bowling & Wieniawa-Nakiewicz, 1986). Finally, there is also a tendency for relay cells with similar orientation preference to be clustered together (Shou & Leventhal, 1989). This does not produce any sublayering however, since the orientation preference of neighbouring cells varies smoothly throughout the depth (both perpendicularly and obliquely) of the LGN.

2.4 The functions of the LGN

Given that the spatial receptive field profiles of geniculate cells are fairly similar to their retinal inputs, it is commonly assumed that the LGN is merely a relay station sending its retinal inputs

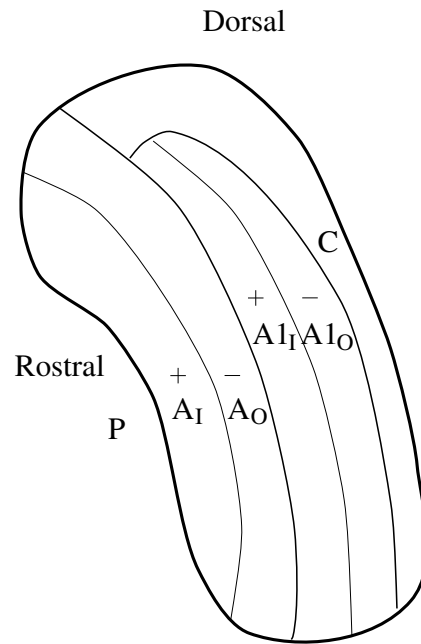


Figure 2.4: The sublayering of on- and off-centre cells in the ferret LGN. The figure shows a parasagittal view of the LGN. Inner leaflets of a layer are indicated with an I subscript and outer leaflets with an O subscript. The polarity of cells within the leaflets are indicated by + (on-centre) and - (off-centre). The P indicates the location of the perigeniculate nucleus (PGN). Figure adapted from (Stryker & Zahs, 1983).

onto cortical areas without significant processing. In this thesis, the focus is placed on the way the retinogeniculate mapping develops, rather than on the role of the LGN in visual processing. It is worth briefly reviewing however some of the possible roles that the LGN may play in the processing of visual information.

2.4.1 The corticogeniculate pathway

Although the LGN is the primary recipient of retinal axons, these axons only account for between 10–20% of excitatory input to geniculate relay cells — the most dominant source of input to relay cells is from cells in layer VI of the visual cortex, accounting for around 50% of the input (Sherman & Koch, 1986). (The remaining 30–40% of relay cell input comes from LGN interneurons and perigeniculate cells). Cortical cells also project back onto LGN interneurons, although relay cells receive roughly twice as many synapses from cortical cells as interneurons (Montero, 1991). This projection is topographic, such that cortical cells responding to contralateral (or ipsilateral) input tend to synapse back onto relay cells in layer A (or A1) (Murphy & Sillito, 1996).

This pathway, the corticogeniculate pathway, is often termed a feedback pathway, although the term ‘feedback’ is slightly misleading due to its large size. In terms of the numbers of axons, it is the biggest geniculate pathway, as shown in Figure 2.5. Given the evolutionary pressure to reduce the amount of wiring in the brain, the LGN must be doing something significant to justify such a large corticogeniculate pathway (Jones, 1985).

The purpose of this pathway is as yet unclear, although it is commonly assumed that it gates the flow of retinal information to the cortical areas in some way. Another common view is that

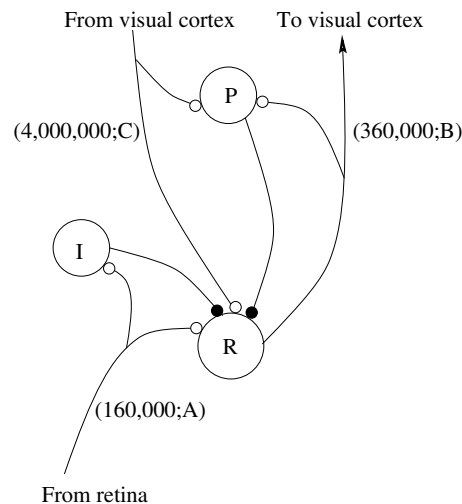


Figure 2.5: Summary of the main inputs to cat LGN X relay cells and their approximate magnitudes. The Y pathway is simpler than the X pathway, since LGN interneurons are not involved. Figure adapted from (Sherman & Koch, 1990). Numbers in parentheses indicate the approximate size of the pathway along with the references: A — (Williams et al., 1986), B — (LeVay & Ferster, 1979), C — (Sherman & Koch, 1986). Excitatory synapses are open, inhibitory synapses are filled. Cell types: R — LGN relay cell, I — LGN interneuron, P — perigeniculate cell.

the LGN is an “internal retina” which, unlike the real retina, can have its activity altered by inputs coming from other (mostly cortical) areas (Singer, 1977). This view has been developed into a general notion that thalamic areas are “active blackboards” for cortical areas to use as places to collect information and update hypotheses (Mumford, 1991). The corticogeniculate pathway has also been suggested to form a crucial role in shaping cortical cell responses such as “end-stopping” (preference for short stimuli rather than long stimuli) (Murphy & Sillito, 1987). Another possibility is that the corticogeniculate pathway acts to synchronize the activity of geniculate relay cells that are responding to a certain stimulus. This in turn helps cortical cells to respond to one particular stimulus since the relay cells are now more likely to be firing coherently, making the stimulus more salient (Sillito, Jones, Gerstein, & West, 1994; Singer, 1994). Finally, cortical inputs to geniculate relay cells can affect the low-threshold Ca^{2+} conductance of the cell and thus switch the cell from a bursting mode to a tonic mode (Sherman & Guillery, 1996), which dramatically affects its response to stimuli.

Regardless of the function of the corticogeniculate pathway, the gating of signals is assumed to be performed by interaction with cells in the reticular nucleus formation (RNT). This lies dorsally and anterior to the thalamus. Axons from each thalamic area, including the LGN, send side-branches out to an appropriate area of the RNT on their way to the cortex. Additionally, cortical axons also make side branches into the RNT on their way back to the appropriate thalamic area. The area of the RNT linked to LGN relay cells is known as the perigeniculate nucleus (Sherman & Koch, 1990), and is illustrated in Figure 2.5.

2.4.2 Other theories of LGN function

The main theories regarding the role of the LGN in visual processing are dominated by the existence of the corticogeniculate pathway and have been outlined above. In this section we briefly consider some of the other theories.

One role of the LGN may be to organise the retinal inputs in a way that can help cortical processing by grouping together certain cells so that they are close together. For example, the segregation of inputs in the LGN is maintained in the projection to the cortex. In the cat, central regions of the A laminae project to layer IVb of cortical area 17, whereas border regions project to layer IVa (Bowling, 1989b).

Alternatively, since there are around four times as many geniculate relay cells than RGCs (Sanderson, 1971b), there is considerable divergence in the retinogeniculate pathway. This divergence could give the cortex a fairer representation of the X and Y cell pathways. In the retina, there are around five times as many X RGCs as Y RGCs. In the geniculate, this ratio is reduced to around 2:1, since Y-RGCs innervate many more LGN relay cells than X RGCs (Sur et al., 1987; Sherman & Koch, 1990). This increase in the proportion of Y-like cells in the geniculate reduces the dominance of the X pathway that exists in the retina.

Another role of the LGN may be captured in the way that geniculate cells respond over time. Lagged and non-lagged geniculate cells can transform retinal responses in the temporal domain. It has been hypothesised that just as the retina removes spatial correlations using centre-surround receptive fields (Atick & Redlich, 1992), the LGN removes temporal correlations using the temporal response characteristics of lagged and non-lagged cells (Dan, Atick, & Reid, 1996). Finally, it has also been suggested that geniculate cells adjust the variance of retinal signals on their way to the cortex, transforming the non-linear variance of retinal cell firing rates to mean firing rates into linear variances. (Levine, Cleland, Mukherjee, & Kaplan, 1996).

2.5 Factors influencing neural development

The first half of this chapter has focused on the properties of the visual pathway from the retina, through the LGN, up to the cortex, and back down to the LGN. The second half of this chapter considers the central question of how this pathway comes into existence during the early stages of the animal's development. In general, we are concerned with how one set of cells, the presynaptic cells, connects to another group of cells, the postsynaptic cells, to produce the mature pattern of ocular segregation and topographic mapping. For the specific problem of retinogeniculate development, the presynaptic cells are the retinal cells, and the postsynaptic cells are the geniculate cells (or tectal cells, depending on the species). In this section we also consider the development of similar sensory pathways, both visual and non-visual.

2.5.1 Marker theories and genetic specification

The early investigations into the development of visual pathways were commonly performed in amphibians, rather than mammals. Such pathways commonly show no intermixing of signals from different eyes, although the retinotectal mapping is retinotopic, as shown in Figure 2.6.

To look at the development of this retinotectal mapping, experiments by Sperry and colleagues through the 1940–1960s examined the regrowth of retinal cells into the optic tectum in amphibians

after the optic nerve had been severed. Their results showed that when retinal ganglion cells grew back into the tectum, they always seemed to know which part of the tectum to innervate. On the basis of this and similar experiments, Sperry proposed the ‘chemoaffinity hypothesis’ (Sperry, 1963). Applied to the retinotectal pathway, it claims that each retinal cell has a unique marker, and the tectal cell to which the retinal cell should be wired to carries the same unique marker. During development, each retinal cell searches for the tectal cell with the same (or closest) marker and then connects to it.

The chemoaffinity hypothesis assumes that the connectivity of the cells is rigidly determined by the distribution of markers in the retinal axons and in the tectum. However, several mismatch experiments showed that abnormal retinotectal mappings could develop during regrowth. For example, Gaze and Sharma (1970) showed that a whole retina could re-grow onto half a tectum in a topographic fashion, as shown in Figure 2.7. These and other regrowth experiments showed that there is not a fixed affinity between retinal and tectal cells. (A thorough review of the various mismatch experiments, along with a computer model accounting for most of the biological evidence, is given by Willshaw and von der Malsburg (1979), showing that in many cases the regenerated maps are very different to the maps that would be expected from a fixed chemoaffinity process.)

Despite this evidence from regeneration experiments, it is highly likely that chemical markers of some form are used early in pathway development (Goodman & Shatz, 1993), but by themselves they do not account for the full development of the visual system. So, what other processes are crucial in development?

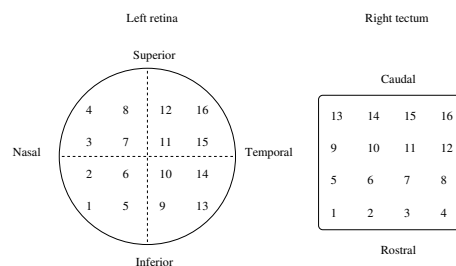


Figure 2.6: The retinotopic mapping of the left retina onto the right optic tectum. Numbered retinal cells normally innervate the appropriately numbered area of the tectum. The topography of the retina is preserved in the tectum, such that neighbouring retinal cells innervate neighbouring areas of the tectum. Figure adapted from (Gaze & Sharma, 1970).

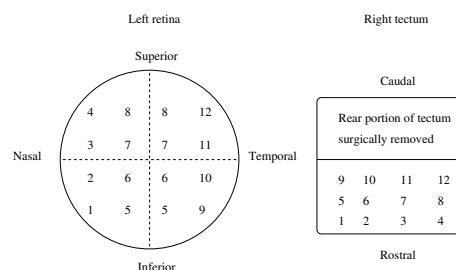


Figure 2.7: The retinotectal mapping following removal of the caudal half of the tectum. The retinal mapping is compressed so that the remaining half of tectum still shows an orderly mapping of retinal space. Although the map in the tectum is still orderly, some tectal cells (marked 5–8) can be stimulated by more than one region of the retina as a result of the compression. Figure adapted from (Gaze & Sharma, 1970).

2.5.2 The role of experience in development

The role that visual experience plays in the development of visual pathways was first assessed by a series of experiments looking at the development of ocular dominance in cat visual cortex (Wiesel & Hubel, 1963a, 1963b; Hubel & Wiesel, 1963). In a newborn kitten, most cortical cells in the primary visual cortex (area 17) respond to stimulation of either eye, although there is a bias favouring stimulation of one of the eyes. After around six weeks of normal visual experience, most cells have adapted to respond to stimulation of only one eye, ignoring stimulation from the other eye. Overall, both eyes normally innervate an equal number of cortical cells. However, if one of the eyes is covered from birth for six weeks so that it does not receive any visual input, most (if not all) of the cortical cells do not respond to stimulation of the eye that was deprived of vision. Wiesel and Hubel (1963a) repeated these experiments with kittens of different ages. They found that the older the kitten was at the onset of deprivation, the weaker the effects of deprivation. Furthermore, depriving one eye of vision during adulthood did not affect the pattern of ocular dominance at all. These results indicated a critical period of development, during which visual deprivation radically alters development, and after which the effects of deprivation are weakened.

Other deprivation experiments examined the effects of visual deprivation on the orientation selectivity of cortical cells. The cortical cells of normal cats have a preference for stimuli of a particular orientation. The distribution of these orientation preferences is complete: taking a population of cells, all possible preferred orientations are found. If a newborn kitten is raised in a restricted visual environment consisting of either horizontal or vertical lines, the effect on the distribution of preferred orientations is severe: most cells developed a preference for either horizontal or vertical lines and cells preferring oblique lines were rare (Hirsch & Spinelli, 1970; Blakemore & Cooper, 1970).

A different approach to studying development has been to reroute visual inputs into non-visual thalamic areas (Sur, Garraghty, & Roe, 1988). Sur and his colleagues rerouted retinal signals of newborn ferrets into the auditory thalamus. The ferrets were then raised to maturity. After maturation, cells in the auditory thalamus were found to be responsive to visual stimulation. Furthermore, cells in the auditory cortex were also found to be responsive to visual stimuli, exhibiting orientation and directional selective receptive field properties. These properties are similar to those found in normal ferret visual cortex, although the receptive fields tended to be broader and more sluggish than normal. Sur et al. conclude by suggesting that “the modality of a sensory thalamic nucleus or cortical area may be specified by its inputs during development”. In other words, it could be that there is nothing intrinsically special about the visual cortex: it becomes the main area for visual processing just because it receives visual input.

Similar experiments have since been performed in other animals. Métin and Frost (1989) rerouted newborn hamster visual inputs into somatosensory cortex and then reared the hamsters to adulthood. They found that somatosensory cells developed visual receptive field properties similar to those found in normal hamsters. A slightly different approach was taken by Schlaggar and O’Leary (1991), transplanting a section of embryonic rat occipital cortex (normally the area which develops to become the visual cortex) to somatosensory cortex. Normal somatosensory cortex develops “barrels”, which are not found elsewhere in the cortex. In the transplanted occipital cortex however, barrels similar in shape and distribution to normal somatosensory barrels were

found. Hence the ability to generate barrels is not due to some intrinsic property of somatosensory cortex, but instead could be generated by the thalamic inputs to cortex.

Taken together, there is a wide range of evidence showing that development of visual pathways does not just follow a strict genetic blueprint: visual experience also plays an important role. A popular principle based on the results from the experiments just outlined is that the cortex makes use of activity-driven processes — processes that are driven by the neural activity in the visual pathway (Goodman & Shatz, 1993). Quite often, these processes are self organising, in that the neurons adapt according to the pattern of activity in the system, rather than developing according to some predefined program (von der Malsburg & Singer, 1988). For example, Blakemore and Cooper (1970) hypothesised that:

the visual cortex may adjust itself during maturation to the nature of its visual experience. Cells may even change their preferred orientation towards that of the commonest type of stimulus; so perhaps the nervous system adapts to match the probability of occurrence of features in its visual input. (Blakemore & Cooper, 1970, p478)

In this framework, the deprivation experiments alter cortical cell responses by altering the statistics of the visual environment, and therefore the levels and patterns of neural activity in the system. The results of the re-routing and transplant experiments can also be explained using an activity-dependent process: each part of the cortex adapts via general modification principles to its inputs, regardless of the sensory modality of the inputs.

2.6 Development of the LGN

Are activity-driven processes responsible for any development of the LGN? The previous section has shown that many aspects of postnatal development of the visual system rely on activity-dependent processes, with the neural activity being generated by visual stimulation of the retina. However, in the case of the ferret and cat LGN, the development of ocular segregation and topography has begun before eye opening. This rules out any processes driven by visually-evoked neural activity. This section considers the development of the LGN before the onset of visual stimulation and considers if activity-driven processes are responsible for any aspects of LGN development.

2.6.1 Normal development

The normal sequence of development for the cat and ferret retinogeniculate pathway are summarised in Tables 2.2 and 2.3. The cat and ferret share a common sequence of neural development, although the ferret is born about three weeks earlier than the cat. Most developmental experiments are therefore performed on the ferret postnatally rather than in the cat prenatally since it is easier to perform the experiments in postnatal animals.

The sequence of normal development in the cat

In the cat, retinal axons begin to enter the LGN at around embryonic day 32 (E32). The contralateral eye has a competitive advantage because its axons innervates the LGN around three days before axons from the ipsilateral eye (Shatz, 1983). When an retinal axon innervates the LGN, it tends to give off side-branches throughout both layers A and A1 of the LGN, so that most regions

Day	Event
E23	Retinal axons are yet to arrive at the optic chiasm (Sherman, 1985).
E30–E40	Correlated bursting activity begins (R. Wong 1994, personal communication).
E32	First retinal axons reach the LGN. During next two weeks, most cells in A and A1 remain binocular (Sherman, 1985).
E38	Functional synaptic transmission exists between ganglion cell axon and LGN neuron (Shatz, 1994).
E47	Segregation of retinal inputs into laminae begins.
E48	Loss of ganglion cells begins (Sherman, 1985).
E60	Retinal OPL appears (Greiner & Weidman, 1980).
E63	Birth.
P5	RGC dendrites begin to unistratify into one of the two IPL sublayers of the retina (Maslim & Stone, 1988).
P6	Rod photoreceptors still developing (Sidman & Wislocki, 1954).
P10	Eye opening (Greiner & Weidman, 1980).

Table 2.2: Timescale of cat retinogeniculate development. Days are given as either embryonic (E) or postnatal (P).

of the LGN initially receive input from both eyes (although some regions are solely innervated by contralateral axons throughout development).

By E48, the segregation into eye-specific layers has begun. The majority of axons from the contralateral eye solely innervate layer A, and most of the ipsilateral axons only innervate layer A1. This is achieved by the elimination of inappropriate side-branches and the elaboration of branches in the appropriate layer. Cellular lamination cannot yet be seen at E48, but occurs later at around E62, just before birth (Shatz, 1983). By birth, the LGN has clearly segregated into eye-specific layers. At this stage however, the development of the retinal circuitry is lagging far behind that of the geniculate. The retinal photoreceptors and most of the rest of the retinal circuitry are still developing early after birth (Greiner & Weidman, 1980). Only by around P7–10 do the eyelids open — long after the onset of eye-specific segregation in the geniculate.

The sequence of normal development in the ferret

The normal pattern of development in the ferret is similar to cat development, with the difference that most of the development occurs postnatally. In addition, the segregation of on- and off-centre inputs occurs sometime in the third postnatal week, P14–21 (Linden, Guillery, & Cucchiari, 1981). Shortly before this time, the RGCs dendrites restrict themselves to the appropriate sublayer in the IPL of the retina (P9, Wingate & Thompson, 1995).

2.7 Neural activity in the developing retina

Activity-dependent processes require neural activity in order to work — is there any such activity present at the time when the LGN is developing? In the case of LGN development, any activity

Day	Event
E28	Retinal axons have not yet innervated LGN (Guillery, Lamantia, Robson, & Huang, 1985).
E41	Birth. Contralateral axons are extended further into the LGN than ipsilateral axons (Linden et al., 1981).
P4–7	Eye specific layers form (Linden et al., 1981).
P9	On- and off-centre RGCs restricted within the IPL (Wingate & Thompson, 1995). On- and off-centre RGCs fire synchronously (Wong & Oakley, 1996).
P14–21	Axon terminals segregate into on and off sublayers in LGN (Linden et al., 1981; Hahm, Langda, & Sur, 1991).
P17	Asynchronous firing of on- and off-centre RGCs present (Wong & Oakley, 1996).
P30	Waves are no longer present (Wong et al., 1993). Eyes open (same time as cat in terms of days post-conception) (Linden et al., 1981; Wingate & Thompson, 1995).

Table 2.3: Timescale of ferret retinogeniculate development. Days are given as either embryonic (E) or postnatal (P).

cannot be the result of visual stimulation for two reasons. First, the eyes are not open in the cat until around P7–10 (Greiner & Weidman, 1980), by which time the LGN has considerably developed. Second, the retinal circuitry is still developing and the photoreceptors are not yet mature (Sidman & Wislocki, 1954; Greiner & Weidman, 1980). It is therefore highly unlikely that any activity in the retina can be generated by visual stimulation. Is there any retinal activity in the absence of visual stimulation?

Since the early 1980's, it has been known that retinal ganglion cells in the adult cat have a significant maintained discharge even in the absence of visual stimulation. Furthermore, neighbouring retinal ganglion cells show correlated activity patterns, whereas more distant pairs of cells do not. Finally, the correlation respects the polarity of the cells: neighbouring like-centred pairs of cells are correlated, whereas neighbouring opposite-centred pairs of cells are anticorrelated (Mastrorade, 1983a, 1983b, 1983c). However, such measurements were only performed on adult cats, leaving open the question of whether such correlations are present in developing animals.

2.7.1 Waves of spontaneous activity

The possibility of correlated retinal activity before the onset of patterned vision has recently been addressed by a set of experiments recording from the developing retina. The existence of neural activity in the developing retina was first reported by Galli and Maffei (1988). In vivo recordings of rat retinal ganglion cells aged E17–E21 showed individual cells were spontaneously active. (The term *spontaneous* was used to emphasise that the activity was not evoked by visual or other external stimulation.) Further studies in the prenatal rat using pairs of electrodes to record from two or three neighbouring ganglion cells showed that neighbouring ganglion cells are correlated in their activity patterns: they tend to fire together, and be silent together (Maffei & Galli, 1990).

These experiments confirmed the existence of spontaneous activity, correlated across neighbouring retinal cells, in the developing retina.

Since then, two developments in recording techniques have revealed more details about the nature of this spontaneous activity. First, the development of the multielectrode array has allowed the activity of around one hundred neighbouring retinal ganglion cells to be recorded simultaneously (Meister et al., 1991). These recordings show that the cells fire synchronously, producing waves of activity that move across the retina. Typically, these waves last for a few seconds, and then the patch of retina being recorded is quiet for around a minute before another wave begins. As the retina develops, the waves tend to be present for longer and are broader. By P30, the waves are no longer detectable and have been replaced by normal adult maintained activity levels (Mastrorarde, 1983a; Wong et al., 1993).

Second, the use of fluorescent dyes in optical imaging have revealed fluctuations in levels of intracellular calcium in the retinal ganglion cell layer (Wong, Chernjavsky, Smith, & Shatz, 1995). This imaging technique has allowed the mechanisms generating the waves to be investigated in more detail, suggesting that the horizontal pathway in the retina (amacrine–ganglion) is responsible for generating the waves, rather than the vertical pathway (photoreceptors–bipolar–ganglion). This optical imaging technique has also been used to demonstrate that the correlations between on- and off-centre retinal cells changes over time. In the ferret, neighbouring on- and off-centre ganglion cells have similar firing patterns for the first two weeks of postnatal life. However, at around P17, the firing patterns of on- and off-centre cells differs: off-centre cells tend to fire much more often than on-centre cells, such that when an on-centre cell fires, so does an off-centre cell, but not vice-versa (Wong & Oakley, 1996). This distinct firing pattern occurs at around the same time as when LGN layers are segregating into on- and off-centre leaflets (Hahm et al., 1991).

2.8 The central hypothesis

The existence of correlated activity in the retina during the period of retinogeniculate development has led many people to the hypothesis at the centre of this thesis (Maffei & Galli, 1990; Jeffery, 1989; Meister et al., 1991; Wong et al., 1993):

The correlated spontaneous activity in the retina drives Hebbian-based synaptic modification processes to develop the patterns of ocular segregation and topography in the retinogeniculate pathway.

Although this hypothesis is specifically concerned with retinogeniculate development, it is also thought to apply to the refinement of other sensory pathways (Wong, 1993; Katz & Shatz, 1996), such as the avian auditory pathway (Lippe, 1994), and rat somatosensory cortex (Yuste, Peinado, & Katz, 1992).

2.8.1 Propagation of retinal waves to the LGN

A pre-condition for the waves to drive retinogeniculate development is that retinal axons must be making functional connections to the geniculate during development. This seems to be the case: cat geniculate cells can be activated by electric stimulation of one or both of the optic nerves from E39 onwards (Shatz & Kirkwood, 1984). More recent work in the neonatal mouse (P0–4)

has shown that LGN neurons are driven to generate an action potential (or ‘spike’), and that the period of spiking (every 30–90 seconds) matched the period of firing in the retina (Penn et al., 1995; Mooney, Penn, Gallego, & Shatz, 1996). Hence, not only are there functional synapses between the retinal and geniculate cells, but retinal waves can drive geniculate neurons to spike. Furthermore, groups of LGN neurons were correlated in their firing patterns, making it likely that the spontaneous activity may propagate into, and modify, the geniculocortical pathway (Mooney et al., 1996).

2.8.2 Abnormal development of the LGN

Does the LGN develop according to some predefined genetic program, or can its development be influenced by environmental factors? This section describes some of the ways in which LGN development is affected by various conditions, both natural and experimental.

Evidence that the LGN is not purely hard-wired comes from Siamese cats (Shatz, 1977). In these cats, a genetic mutation causes temporal ipsilateral retinogeniculate axons to be misrouted at the optic chiasm. This causes the axons to end in layer A1 of the contralateral LGN, rather than in their normal target, layer A1 of the ipsilateral LGN. (However, it is possible that the genetic mutation also directs the temporal ipsilateral retina to the contralateral geniculate.)

Experimental manipulations to the retinogeniculate pathway also affects its development. In the ferret, for example, the removal of both eyes before the retinal axons reach the LGN (E28) prevents development of the eye-specific laminae (Guillery et al., 1985). When one eye is removed from the ferret at birth, the remaining eye innervates its normal region, as well as some other geniculate regions which are normally innervated by the eye which has been removed. For example, in the LGN ipsilateral to the surviving eye, layer A1 is thicker than normal and includes part of the nucleus that is normally the monocular segment (Guillery et al., 1985). Furthermore, when one eye in the ferret is removed at P0, the remaining eye innervates the whole of both LGNs, producing an on-centre and off-centre layer to replace the eye-specific layers (Morgan & Thompson, 1993). The on-layer is anterior and the off-layer is posterior, indicating that although the absolute positioning of cells is not genetically encoded, there is probably some bias (either through time of arrival or chemical gradients) to ensure that the on-leaflets are always anterior to the off-leaflets, in the same way that the contralateral eye always innervates layer A.

Similar enucleation experiments have been performed in the prenatal cat. After monocular enucleation at E44, both LGNs develop only two layers instead of the normal five. The dorsal layer is a so-called “magnocellular” layer, since it contains medium sized cells, and a ventral “parvocellular” layer with small sized cells. Also, the interlaminar gap between the two layers occurs in a novel position, rather than in the same position as one of the normal layers (Garrahy, Shatz, & Sur, 1988).

All of the previous experiments mentioned indicate that the retinogeniculate pathway is not simply defined by some genetic blueprint, but instead its development can be affected by various environmental factors. If the LGN can be affected by environmental factors, and if the retina is spontaneously active, can this activity also affect the development of the LGN?

2.8.3 Activity blockade experiments

The role of activity in the development of the retinogeniculate pathway has typically been investigated by blockade experiments, whereby the retinal activity is silenced in some way. The first set of experiments were performed in the neonatal kitten, blocking action potentials from one eye using intraocular injections of TTX (tetrodotoxin — a sodium-channel blocker) (Archer, Dubin, & Stark, 1982; Dubin, Stark, & Archer, 1986). Action potential blockade was maintained in one eye for the first twenty weeks of life whilst the other eye remained active. At the end of this period, individual relay cells from both LGN were examined. LGN cells in the deprived layer (layer A1 (or A) of the ipsilateral (or contralateral) LGN to the deprived eye) received mixed retinal inputs from both on- and off-centre retinal cells and both X and Y retinal cells. In comparison, cells in the non-deprived layer received inputs from only on- or off-centre cells and from only X or Y retinal cells. The blockade also caused the size of the geniculate layers to change: the non-deprived layer in each LGN increased in size (from 500 μm deep to around 800 μm deep) at the expense of the deprived layer (from 500 μm to 200 μm). The effect of action potential blockade was therefore to prevent the refinement of synaptic connectivity of the retinogeniculate pathway, causing geniculate cells to receive abnormal retinal input from mixed classes and for the non-deprived eye to innervate geniculate territory that should have been innervated by the deprived eye. Examination of retinogeniculate synapses showed that action potential blockade prevents their structural development, leaving them in an immature newborn state (Kalil, Dubin, Scott, & Stark, 1986).

The study by Dubin et al. (1986) only looked at early postnatal development, by which time eye-specific layers have already formed. More binocular cells were found in the experimental animals than the control animals, although they were mainly located at the border between layers A and A1 rather than throughout the whole lamina. A more appropriate experiment was performed by Shatz and Stryker (1988), blocking activity in both eyes by infusion of TTX above the optic chiasm during the period E42–58 when eye-specific layers form (Shatz, 1983). The effect of TTX on individual retinal axons was two-fold. First, axons from both eyes did not refine into their appropriate geniculate layer, but instead spread throughout much of the LGN preventing the development of eye-specific layers. Second, the terminal arbors of the axons were much wider than the normal thin column-like arbors. The existence of wide-branching arbors indicated that the retinotopic mapping was distorted (Sretavan, Shatz, & Stryker, 1988).

Similar blockade experiments have also been performed in the ferret to assess the role of activity in LGN development. The approach in these experiments was to block NMDA receptors in geniculate cells, since NMDA receptors have been hypothesised to play a role in detecting correlated retinal and geniculate activity (Constantine-Paton, Cline, & Debski, 1990). NMDA receptor blockade during the period P14–21 prevented the segregation of the geniculate into on- and off-leaflets (Hahm et al., 1991). Blockade of afferent activity with TTX application during postnatal weeks three to four also prevented the segregation into on- and off-leaflets (Cramer & Sur, 1997). The diffusible messenger nitric oxide is also required for segregation during this period (Cramer, Angelucci, Hahm, Bogdanov, & Sur, 1996). NMDA blockade was repeated in ferrets aged P1–7, the period when the LGN normally segregates into eye-specific layers (Linden et al., 1981). However, eye-specific layers still formed in the absence of NMDA receptor activation. Rather than ruling out the role of activity-dependent processes in eye-specific segregation,

Smetters, Hahm, and Sur (1994) rightly concluded that there could be other activity-dependent mechanisms at work, or that the NMDA receptors are used to detect correlated activity for other purposes, such as topographic map refinement (Cline & Constantine-Paton, 1989).

2.8.4 Retinotopic map formation

The previous experiments have shown that many aspects of retinogeniculate development rely on neural activity. However, these experiments have only focused on the development of single cells within the pathway, without considering how the overall topography is affected. Unfortunately, experiments considering how the overall topography in mammalian visual system is affected by activity blockade are rare, looking instead at the development of single axon arbors (Sretavan et al., 1988; Simon, Prusky, O'Leary, & Constantine-Paton, 1992) and inferring from single axons how the overall topography might be affected.

Several experiments have been performed however in vertebrates to examine the role of activity-driven processes in retinotopic map formation. In developing frogs, NMDA receptor blockade prevents the refinement of the retinotectal map to such an extent that an area of tectum receives input from about four times the normal area of the retina (Cline & Constantine-Paton, 1989). This indicates that the role of activity may be to eliminate inappropriate synaptic connections, and that without the detection of correlated activity between retinal and tectal cells, mediated by NMDA receptors, the inappropriate retinotectal connections remain and the appropriate connections do not get strengthened.

Similar experiments have been performed in goldfish, following the regeneration of the retinotopic map after severing the optic nerve. After nerve crush, the optic nerve normally grows back into the tectum to reproduce a normal retinotopic map. TTX infusion during this regeneration period did not affect the overall retinotopic arrangement, but instead it dramatically increased the size of individual tectal receptive fields by around a factor of three. There was no difference in RGC receptive field size between the control and experimental animals, indicating that the larger tectal receptive field sizes are due to tectal cells receiving many more retinal inputs from a larger region of retina (Schmidt & Edwards, 1983). This again supports the notion that the inappropriate synaptic connections are not being eliminated by activity-driven processes.

If these principles can be applied to mammalian visual systems, it indicates that there are a combination of processes working together to form retinotopic maps. Since the overall topographic mapping in goldfish is not destroyed by activity blockade, some activity-independent process, such as chemical markers, must be guiding retinal axons to their targets. However, this can only produce a coarse-grained topographic mapping. The role of the activity-dependent processes is to refine this mapping by the selective reinforcement of synapses: inappropriate synapses are removed and appropriate synapses are strengthened (Goodman & Shatz, 1993).

Mechanisms of activity-driven processes

First of all, is the pattern of activity crucial for development, or is it just the presence (or level) of activity that initiates some other process, such as the release of nerve growth factors, that is responsible for development (Frank, 1987)? This hypothesis was tested by (Stryker & Harris, 1986) by replacing visually-induced activity with electrical stimulation of the optic nerves in kittens at a time when cortical ocular dominance is developing. When both nerves were stimulated simulta-

neously, ocular dominance columns in the cortex failed to develop. In comparison, asynchronous stimulation of the nerves induced the normal cortical ocular dominance patterns. The temporal pattern of activity, and not just its mere presence, is therefore important. The spatial pattern of activity is also important: exposure to stroboscopic illumination simultaneously activates most retinal cells and prevents the refinement of the topographic maps in goldfish, leaving tectal cells with abnormally large receptive fields (Schmidt & Buzzard, 1993).

Most ideas of activity-driven processes are centred around the principle that “cells which fire together should wire together”. This is a common summary of the idea proposed by Hebb (1949) that if a presynaptic cell is connected to a postsynaptic cell and they tend to be active at the same time, the synaptic strength between them should be increased. Whether this principle is actually implemented in the nervous system is yet to be determined, although some evidence of this kind of plasticity (such as long-term potentiation and long-term depression) has been found in the hippocampus (Bliss & Lømo, 1973) and neocortex (Artola, Bröcher, & Singer, 1990; Friedlander, Frégnac, & Burke, 1993). Recent work in the developing retinogeniculate pathway has also shown that short bursts of stimulation can cause long lasting enhancement of synaptic transmission (Mooney, Madison, & Shatz, 1993; Mooney, Penn, & Shatz, 1995; Shatz, 1996). It is therefore likely that some form of Hebbian synaptic modification is taking place, although the fine details are yet to be worked out. In Chapter 3, the mathematical formulations of the Hebbian principle are discussed in more detail.

2.8.5 The need for activity-independent processes

Various experiments have just been described which summarise the evidence in favour of development being driven by activity-dependent processes. Are these processes alone sufficient to account for retinogeniculate development, and neural development in general? It is highly unlikely that everything is driven by activity-mediated processes for two reasons. First, activity blockade has been shown to have no effect on certain developmental processes. For example, TTX blockade in goldfish during the time when the optic nerve is finding its way to the optic tectum does not prevent the optic nerve reinnervating the tectum in the normal retinotopic manner (Schmidt & Edwards, 1983). Second, there are some forms of organisation which are unlikely to be solely formed by activity-driven processes. In the case of LGN development, the contralateral eye always innervates layer A and the ipsilateral eye always innervates layer A1. This is due to the relative advantage that the contralateral eye has, since it innervates the geniculate several days before the ipsilateral eye and can therefore innervate deeper into the geniculate. The difference in arrival times for axons from different eyes is an example of an activity-independent mechanism. At the sublaminal level, the on-leaflet is always anterior to the off-leaflet in the ferret. (In the case of segregation of on- and off-centre inputs, the mechanism responsible for the consistent ordering of the sublayers between animals is still unknown.)

It is therefore likely that a combination of activity-independent and activity-dependent processes are at work during different stages of development. The early development of the retinogeniculate pathway could be performed using activity-independent mechanisms, such as positional markers or chemical gradients (Yuasa, Hirano, Yamagata, & Noda, 1996; Rétaux & Harris, 1996), to create a coarse retinotopic mapping. This coarse mapping could then be refined or sharpened

using activity-dependent processes which eliminate inappropriate synapses and strengthens topographically correct synapses. It is likely that this combination of early activity-independent processes followed by later activity-dependent processes apply to other pathways, both visual and non-visual (Goodman & Shatz, 1993).

2.9 Summary

This chapter has introduced the main elements of the early visual pathway in mammals from the retina to the geniculate, and from the visual cortex back down to the geniculate. The central hypothesis of this thesis has been introduced, stating that the LGN develops to a near-adult state using a combination of activity-independent and activity-dependent processes:

- Retinal axons from both eyes meet at the optic chiasm and project into the appropriate LGN using a combination of activity-independent markers (Hankin & Lund, 1991; Goodman & Shatz, 1993). The initial retinotopic map is quite coarse and most geniculate cells are binocularly driven.
- Spontaneous waves of retinal activity appear soon after the initial map is formed, and jointly refine the retinotopic map and cause geniculate cells to become monocular through Hebbian-like activity-dependent processes (Shatz & Stryker, 1988; Wong, Herrmann, & Shatz, 1991).

The main aims of the thesis are to investigate the form of these activity-independent and activity-dependent processes to see if they are sufficient to account for the refinement of ocular segregation and topography in the LGN. These processes are investigated using computer simulations of retinogeniculate development. In the next chapter, we review the main computer models of visual system development.

Chapter 3

Mechanisms for modelling retinogeniculate development

3.1 Introduction

The last chapter described the two main features within the LGN that are the focus of this thesis. The problems of ocularity and topography development have been well studied, mainly in models of the retinotectal pathway or the retinocortical pathway. Only one model has directly addressed the problem of the joint development of ocularity and topography in the LGN, which will be discussed in more detail in the next chapter (Keesing et al., 1992). Here, we review most of the other influential models and their common components.

3.2 Characterisation of topography and ocular dominance

3.2.1 Topography

The issue of topographic development within a neural pathway can be addressed at two levels. At the level of the single cell, a postsynaptic cell may initially receive diffuse inputs from a wide range of presynaptic cells. During development, the cell refines its connections to receive input from a much smaller and more localised area of the presynaptic sheet. For the postsynaptic unit to respond only to a small region of the presynaptic sheet, it needs to receive input from a few neighbouring presynaptic units. If however the unit receives input from different areas of the presynaptic sheet, it is unlikely that the global map will be topographic.

At the level of the cell population, neighbouring postsynaptic cells may initially receive inputs from different areas of the presynaptic sheet of cells. Gradually during development the postsynaptic cells adapt their connections so that neighbouring postsynaptic cells receive input from neighbouring regions of the presynaptic sheet. When a group of presynaptic units innervate a target set of units, there are various ways in which the units can be organised. Using the terminology of Nelson and Bower (1990), the following types of map are commonly found in different brain areas:

- continuous maps — Figure 3.1(a). These maps exhibit smooth spatial organisation on both a local and global scale. These maps are fairly common, especially within the visual system. Examples include the retinotopic maps in superior colliculus and primary visual cortex.

- patchy maps — Figure 3.1(b). These maps show local smooth organisation, but no global organisation. Such maps are called ‘piece-wise continuous’ because the continuity is only local (Willshaw & von der Malsburg, 1976). The tactile projections to cerebellum is an example of a patchy map.
- scattered maps — Figure 3.1(c). No spatial organisation of neighbouring units exists on either a local or global scale. Such maps can be found in the olfactory cortex.

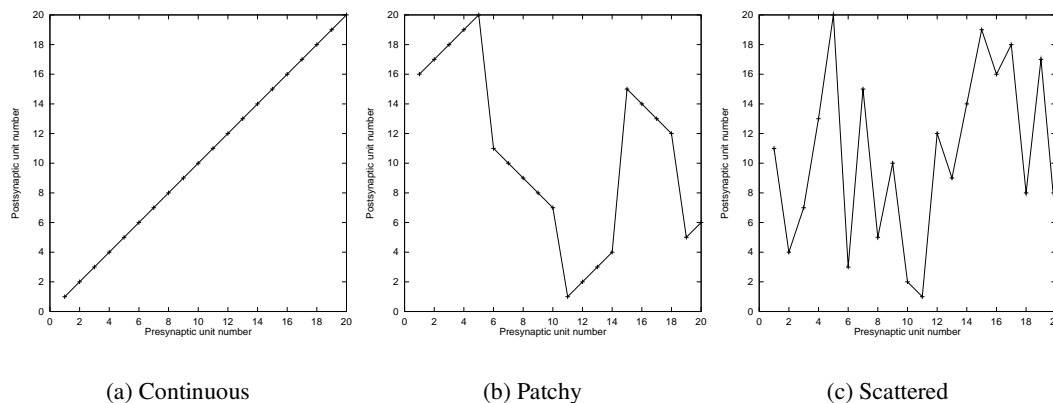


Figure 3.1: Common forms of topographic mappings. A group of twenty presynaptic units is connected to a group of twenty postsynaptic units, each arranged in a one-dimensional line. It is assumed that each presynaptic unit connects to just one postsynaptic unit. Each data point in a graph represents a connection between a presynaptic and postsynaptic unit, and points representing neighbouring presynaptic units are connected by lines. (a): Continuous mapping — both local and global order. (b): Patchy mapping — local ordering but no overall order. (c): Scattered mapping — neither local nor global ordering of neighbouring postsynaptic units. (Diagrams are illustrative only and do not correspond to any simulation results.)

3.2.2 Ocularity

In a similar fashion to the development of topography, ocularity development can be looked at from both the single-cell and multi-cell level. At the single-cell level, we are concerned with how the unit refines its connections from receiving inputs from both eyes to the mature state of receiving input from only one eye. At the multi-unit level, we are concerned with the pattern of ocularity across the whole population of cells. When inputs from two different eyes innervate a group of cells, there are several ways that the patterns can be organised, including the following common forms:

- patches — Figure 3.2(a). Although inputs from two eyes innervate the map, one of the eyes is much larger than the other eye and innervates more postsynaptic units. This occurs in the superior colliculus, where the ipsilateral retinal input is much smaller than the contralateral projection (Graybiel, 1975; Hubel, Levay, & Wiesel, 1975). Patches can also form if one eye is active more often than the other eye.

- stripes — Figure 3.2(b). In comparison to patches, when the inputs from both eyes are of the same size, stripes tend to arise. The best-known example of stripes is the pattern of ocular dominance in area 17 of visual cortex. (Hubel & Wiesel, 1972; LeVay, Hubel, & Wiesel, 1975).
- layers — Figure 3.2(c). The output map can organise so that all of the inputs from one eye are gathered together into one large region. This is closest to the form found in the LGN (Sanderson, 1971a).




Figure 3.2: Common patterns of ocular dominance. Each postsynaptic unit is coloured black if it receives most of its input from the left eye, and white if most input comes from the right eye. The size of the square is proportional to the strength of the dominance. (a): Patchy mapping. The left eye input is much larger than the right eye input, and so occupies more of the output map. (b): Stripes. Both left and right eye occupy roughly the same amount of output space. (c): Layered mapping. Input from the left eye goes mainly to the bottom half of the postsynaptic units, whereas right eye goes mainly to the top half. (Figures are simulation results from the Obermayer model, using different input stimuli. Details of this model are given in Chapter 6.)

3.3 An overview of common mechanisms in developmental models

Many developmental models have been produced to demonstrate how neural activity within the presynaptic sheet can drive refinement of either topography or ocularity of postsynaptic units. Although the models are different in many ways, they all share several common elements which are now described. To enable easy comparison between the models, the notation listed in Table 3.1 will be used rather than the notation from the original references. We first describe the typical architecture of the models and the encoding of activity within the models. After this, we then discuss each of the main mechanisms used in models to simulate neural development.

3.3.1 Network architecture

The problem of modelling a visual pathway is conventionally approached by assuming two layers of units, an input layer (or presynaptic layer) and an output layer (or postsynaptic layer). Units in the presynaptic layer are connected to units in the postsynaptic layer via weighted connections, or weights. These weights are an idealisation of the strength of the synaptic connection between the

Term	Meaning
i, l	Index for presynaptic unit
j, k	Index for postsynaptic unit
w_{ij}	Feedforward weight connecting presynaptic unit i to postsynaptic unit j
l_{jk}	Lateral weight connecting postsynaptic unit j to postsynaptic unit k
N_{pre}	Number of presynaptic units
N_{post}	Number of postsynaptic units
x_i	Activity of presynaptic unit i
y_j	Activity of postsynaptic unit j
σ	Size of neighbourhood
ε	Rate of weight-update

Table 3.1: Terms used for describing model elements.

presynaptic axon and the postsynaptic dendrite. Additionally, some models have lateral weights between units in the postsynaptic layer to account for lateral connectivity between units. In the case of the retinogeniculate pathway, the presynaptic layer represents the two retinae, and the postsynaptic layer represents the LGN.

3.3.2 Unit activations

The membrane potential of a cell is commonly modelled by assigning a real-valued activation level to each unit. For presynaptic units, the activation levels correspond to the activity of retinal units. For models concerned with the influence of the visual world upon neural development, the activity of retinal units corresponds to visual stimulation of the retina. In the case of retinogeniculate development however, presynaptic activity reflects the spontaneous retinal waves described in Chapter 2.

Postsynaptic unit activity is usually a linear function of the presynaptic activity modulated by the value of the weights between the pre- and postsynaptic units. However, some models also include non-linearities and a history of past activation levels to make postsynaptic activation more complex.

3.3.3 Common mechanisms underlying models

The task of modelling neural development is to hypothesise a set of mechanisms to explain how the weights can evolve from an initial unsorted state to a highly refined state. This final state should mimic some aspect of connectivity found in the mature visual system, such as the overall topographic layout of the group of units or the orientation selectivity of an individual unit. Although there are many different models which seem considerably different on a first glance, there are several common underlying principles found in models of development of topography and ocular dominance:

- Correlated input.

Neighbouring retinal units are identified by having similar chemical markers or activation levels. As the distance between the retinal units increases, the similarity of the markers or activity decreases.

- Correlated output.

In a topographic map, neighbouring postsynaptic units have similar responses to the same input stimuli. Neighbouring postsynaptic units therefore need to interact in some way in order to develop similar response properties.

- Modification rules.

Modification rules determine how weights between units adapt over time. Although there are a variety of modification rules, most assume that weights should be increased when the pre- and postsynaptic unit activity is correlated (Hebb, 1949).

- Normalisation / resource limitation constraints.

Most modification rules also need extra constraints to prevent weights growing without bounds. These constraints are often implemented using weight normalisation techniques.

Each of these mechanisms is now discussed in more detail.

3.4 Correlated input

One problem facing the network is how it can discover which presynaptic units are close to each other in the presynaptic sheet. Presynaptic units close to each other should connect to the same postsynaptic unit, whereas distant presynaptic units should connect to different postsynaptic units. This problem has been solved by using either chemical markers or neural activity.

3.4.1 Chemical markers

The first method is to assume that each presynaptic unit contains a set of chemical markers (also known as labels) which codes for the unit's affinity to a postsynaptic unit. Neighbouring presynaptic units have similar affinities for a given postsynaptic unit, and the similarity in affinity decays with distance between presynaptic units. Prestige and Willshaw (1975) introduced two types of labelling schemes. In group I schemes, presynaptic unit j has a maximum affinity for postsynaptic unit j , and an affinity to other postsynaptic units that decreases with distance from postsynaptic unit j . (This assumes that there are an equal number of presynaptic and postsynaptic units.) Group II schemes are more general and simply state that each presynaptic unit has a maximum affinity for the leftmost postsynaptic unit, and the affinity to other postsynaptic units decreases with distance from the leftmost unit (Prestige & Willshaw, 1975). In both of these schemes, it is assumed that the markers are fixed and do not change during development. This raises the problem of how the same set of markers arise in both pre- and postsynaptic sheet in the first instance. An additional problem for chemical markers is how they can regenerate after experimental manipulations, as mentioned in Section 2.5.1.

A different approach to the marker problem is to assume that various sources of markers exist throughout the presynaptic sheet which are induced via the weighted connections to the postsynaptic units. By assuming that these markers can diffuse throughout the presynaptic sheet as well as into the postsynaptic sheet, nearby presynaptic units will tend to have similar sets of chemical markers. In contrast, distant presynaptic units will have quite different sets of markers (von

der Malsburg & Willshaw, 1977; Willshaw & von der Malsburg, 1979). In this way, the initial distribution of markers in the presynaptic sheet can start out random and gradually settle into a organised distribution in both pre- and postsynaptic layers.

3.4.2 Correlated presynaptic activity

The second method of indicating which presynaptic units are neighbours is to encode the proximity of the units through the activation level of units. Neighbouring presynaptic units will tend to be active at the same time, producing correlated activity patterns, whereas distant presynaptic cells will either be uncorrelated or anticorrelated.

Most modern models of visual system development encode proximity of cells within the presynaptic sheet using neural activity. In the specific case of retinogeniculate development, the presynaptic sheet corresponds to the retina, and so the neural activity refers either to visual-evoked activity or the spontaneous waves of activity. For either of these two sources of retinal activity, the correlation between retinal units of the same eye decreases as the distance between the units increases (Field, 1987; Wong et al., 1993).

3.4.3 Discriminating input from different eyes

Within an eye, distances between retinal units can be encoded with either chemical markers or neural activity. Both approaches can also be extended to the problem of discriminating presynaptic units from two different eyes.

The chemical marker approach is to assume that each eye has a unique marker, so that each presynaptic unit has two sets of markers: one set coding for topography and the other set coding for ocularity. So, even if two presynaptic units in the same position within the two eyes have the same topographic markers, they will have different ocularity markers. This provides a basis for a model to discriminate between units from different eyes.

Using neural activity to discriminate units from different eyes is more complicated depending on the source of the retinal activity. In the case of spontaneous waves, the waves are generated independently within each eye. Any pair of retinal units from the two eyes have uncorrelated activity patterns, regardless of the position of the units within each eye. If, however, the retinal activity is visually-evoked, then it is expected that retinal units from different eyes but at the same position within the eye will be correlated, since they respond to the same part of the visual world. Therefore, two units at the same retinal location but from different eyes will be more correlated than two units from the same eye separated by a modest distance.

3.4.4 Representing retinal activity within a model

There are two main ways of representing retinal activity (either driven by visual stimulation or spontaneous) in computer models:

1. Explicit — the model is explicitly presented with a series of input vectors (von der Malsburg, 1973; Barrow, 1987; Durbin & Willshaw, 1987). Each input vector corresponds to one state of the visual world or retinal activity. For example, in the model by von der Malsburg (1973), a sequence of oriented lines were used as input to the network.

2. Implicit — a correlation function is used which captures the long-term correlations between presynaptic units (Swindale, 1980; Linsker, 1986a; Miller, Keller, & Stryker, 1989). These correlational functions usually have relatively simple form (such as Gaussian functions of the distance between the retinal units (Miller et al., 1989)) rather than being calculated empirically by measuring the correlations amongst a representative set of input vectors.

The choice of representation of retinal inputs affects the rest of the model. First, it affects the way in which the computer model runs. With explicit inputs, one input vector is presented to the network, and then the weights are updated before presenting the next input vector. With implicit inputs, the same correlation matrix is used repeatedly to update the weights of the network. Second, using a correlation function is normally much quicker computationally than using a set of explicit input vectors. Third, the use of implicit inputs allows for an eigenvector analysis of the modification rule, since the modification rule is normally of the form $\Delta \mathbf{w} = \mathbf{C} \mathbf{w}$ (see Appendix A.2). This analysis predicts that the shape of receptive fields is dominated by the principle eigenvector of the correlation matrix.

However, this approach has two main disadvantages. First, the correlation matrices tend to be of a fairly simple form and are not usually complex enough to encapsulate the higher-order statistical properties of the visual world. Second, since the matrices describe the long-term correlations between retinal units and uncorrelated noise in the retinal inputs does not change long-term correlations, the correlation matrix is unaffected by any amount of uncorrelated noise. In contrast, the addition of noise to explicit retinal inputs can severely affect development, as would be expected.

Types of explicit input vector

The alternative to using implicit inputs is to use a set of input vectors that explicitly describe the activation levels for a set of presynaptic units. There are two kinds of explicit input vectors that are commonly used in models:

- (low-dimensional) feature vectors — each element of the input vector codes for a different feature of the input stimuli. Typical features include centre of mass, ocularity or orientation (Durbin & Mitchison, 1990; Obermayer, Blasdel, & Schulten, 1991). Neighbouring elements of the input vector typically represent different (rather than similar) properties of the input stimuli.
- (high-dimensional) retinal inputs — each element of the vector stores the activity level of a retinal unit (von der Malsburg, 1973; Barrow, 1987; Obermayer, Ritter, & Schulten, 1990). Neighbouring input vector elements represent the intensity of neighbouring parts of the input stimuli, preserving neighbourhood relations of the input stimuli.

The low-dimensional feature vectors have two advantages over high-dimensional inputs. First, they are of much smaller dimensionality: low-dimensional feature vectors typically code for between two to five different features, whereas high-dimensional retinal inputs can typically store the activity of 32×32 retinal units. Models using low-dimensional feature vectors therefore tend to be much smaller and faster than comparable models using high-dimensional retinal inputs. Second, they simplify the problem by specifying in advance the key input features of the stimuli. In contrast, models using high-dimensional inputs typically face the extra problem of having to extract

salient features from retinal activity. However, this means that a network using low-dimensional inputs cannot discover new features — for example, if the feature vector codes only for centre of mass and orientation, the network cannot extract other salient properties, such as depth, from the feature vector. It can also be argued that a lot of pre-processing is required to extract relevant features from a high-dimensional input vector, and so presenting the network with a feature vector considerably simplifies the task facing the network. Also, the low-dimensional feature vectors can constrain the nature of inputs that the network receives. For example, in the model by Obermayer et al. (1991), the ocular parameter z was varied only to represent monocular, rather than binocular, stimuli. Finally, the feature vectors cannot be directly compared with elements of the biological system, unlike high-dimensional input vectors.

3.5 Correlated output

A common feature of mature neural maps is that neighbouring postsynaptic units produce similar output in response to an input. This feature is not present in the immature network, but is instead acquired during development. To acquire this feature, some form of interaction between neighbouring postsynaptic units must take place. In natural systems, lateral interconnections between postsynaptic cells are likely to ensure that neighbouring cell responses are similar. These lateral interconnections have been introduced into models in several different ways:

- Lateral interaction connections. As well as including feedforward weights (from the presynaptic to postsynaptic layer), lateral connections l_{jk} can be placed between postsynaptic units. The activation of a postsynaptic unit is then defined not only by its inputs from presynaptic units, but also from neighbouring postsynaptic units.

The lateral connections are normally fixed in advance by the modeller to produce short-range excitatory connections and long-range inhibitory connections (von der Malsburg, 1973; Willshaw & von der Malsburg, 1976; Miller et al., 1989). More recent models have also shown that initial random lateral connections can also be adapted over time to develop smooth lateral interaction profiles (Da Silva Filho, 1992; Sirosh & Miikkulainen, 1994).

- Modification rules. The effect of lateral interconnections can often be approximated by assuming that the weights of only the postsynaptic units within a given distance from a “winning unit” are updated (Kohonen, 1982). Often, the winning unit is described as the postsynaptic unit responding strongest to the input stimuli. The modification rule changes the weights of each unit so that the unit will become more responsive to that input stimuli the next time it is presented. In this way, neighbouring postsynaptic units will become selective for similar input stimuli. A similar technique is used by the elastic net in its regularisation term (Durbin & Willshaw, 1987).
- Growth rules. Other weight modification rules can be used to encourage a presynaptic unit to make branches into neighbouring postsynaptic units according to how well it already connects to the postsynaptic units, as used for example by Keesing et al. (1992). Alternatively, a sprouting term can be introduced which makes tentative connections from a presynaptic unit into neighbouring postsynaptic units (Willshaw & von der Malsburg, 1976).

3.6 Normalisation methods and resource limitations

Neural systems work within several constraints, including restrictions on firing rates and synaptic strengths which must be both positive and bounded at some upper limit. These constraints are often introduced into models for two reasons. First, many Hebbian-based modification rules are unstable, adapting weights without bounds. Normalisation constraints ensure that weights remain within limits. Second, the normalisation constraints introduce competition between weights: as one connection strength increases, others must decrease to keep the normalisation sum constant. It can also be argued that constraints are introduced simply because they are present in the natural system: synaptic strengths cannot increase without bounds, and cells have upper limits on firing rates. However, it is often better to keep the model simple and introduce extra mechanisms only when they perform a specific function.

Normalisation can be implemented in several ways. A simple method is to constrain individual weights to lie within a lower and upper bound. Usually, once a weight reaches either of these bounds, the weight is frozen so that it remains at the same value for the rest of the simulation. Weight freezing is often used as a way of ensuring that the network eventually converges to a stable solution (Linsker, 1986a; Miller et al., 1989; Goodhill, 1992) but also means that most weights are forced to either the minimum or maximum value.

A more common form of normalisation is to constrain the sum of a set of weights to be a fixed value. Any increase in one of the weights means that other weights must be decreased to keep the sum constant. This constraint is normally applied separately to all of the weights connected to a given unit, rather than to all of the weights in a network. This provides a form of competition for each unit, so that if one weight connected to a unit is increased by modification, some of the other weights also connected to that unit must decrease.

For models with more than one postsynaptic unit, the normalisation can be applied to each of the presynaptic or postsynaptic units, or to both sets of units. To make it clear which normalisation is being referred to, the following two terms will be used:

Presynaptic normalisation The weights *from* each presynaptic unit are constrained to sum to a constant value. This is also known as “fan-out” normalisation.

Postsynaptic normalisation The weights *to* each postsynaptic unit are constrained to sum to a constant value. This is also known as “fan-in” normalisation.

See Figure 3.3 for an example of both presynaptic and postsynaptic normalisation. Although either presynaptic or postsynaptic normalisation ensures that weights remain bounded, postsynaptic normalisation is more common in developmental models. Some models choose to apply both presynaptic and postsynaptic normalisation (Miller et al., 1989; Goodhill, 1992; Keesing et al., 1992). In principle, it is not possible to satisfy both presynaptic and postsynaptic normalisation — enforcing the second constraint will break the first constraint. However, in practice, both constraints can usually be satisfied (Goodhill, 1992, see also Section 4.6.4 of this thesis).

Using two normalisation techniques usually means that one of the normalisations is being used for its intrinsic properties rather than to keep the weights bounded. First, normalisation can be used to ensure that a presynaptic unit remains connected to the sheet of postsynaptic units, and vice-versa. Otherwise, it is possible that some units will “disconnect” from the network. Second, and

more subtle, the different normalisation techniques modify the weight vector in different ways, which can affect the final outcome quite dramatically. The properties of different normalisation techniques are now discussed.

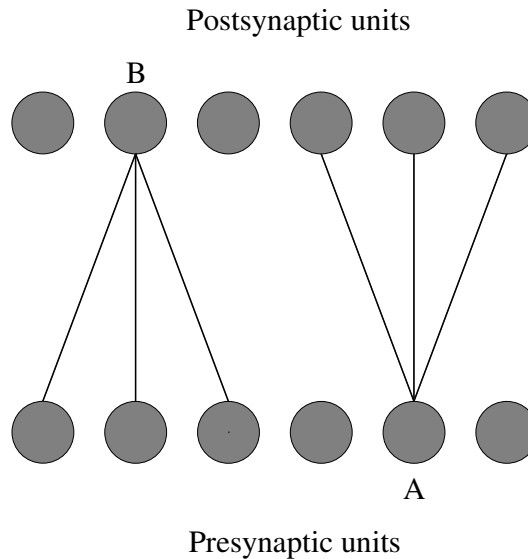


Figure 3.3: The definition of presynaptic and postsynaptic normalisation within a two-layer network. Presynaptic normalisation constrains the sum of the weights from each presynaptic unit (such as unit A). Postsynaptic normalisation constrains the sum of the weights to each postsynaptic unit (such as unit B). To simplify the figure, only a subset of the possible connections between units have been drawn.

3.6.1 Divisive and subtractive normalisation

The two common methods for implementing normalisation across a group of weights are by divisive (also known as multiplicative) or subtractive (also known as additive) means. If we assume that the target value for the sum S of some weights is T , either of the following methods will ensure $S = T$ after normalisation:

1. Divisive: *divide* each element of the weight vector by $\frac{S}{T}$. This is a one-step procedure because no positive weights can turn negative through divisive normalisation.
2. Subtractive: *subtract* $\frac{S-T}{n}$ from each element of the weight vector. This is an iterative method to prevent weights turning from positive to negative (see Appendix A.3 for details).

When the weights are constrained to be non-negative, the sum S is normally defined to be the sum of the weights. In other models when weights are free to vary between negative and positive values (such as Linsker, 1986a) (violating Dale's law (Dale, 1935)), S is normally taken to be the sum of the square of each weight. For presynaptic normalisation, the set of weights from each presynaptic unit i are normalised using $S_i = \sum_{j=1}^{N_{\text{post}}} w_{ij}$. For postsynaptic normalisation, the weights into each postsynaptic unit j are normalised using $S_j = \sum_{i=1}^{N_{\text{pre}}} w_{ij}$.

Divisive and subtractive normalisation modify the weight vector in different ways. Divisive normalisation changes the magnitude of the weight vector whilst maintaining its direction. In

contrast, subtractive normalisation changes both the direction and magnitude of the weight vector. This difference is demonstrated for a two-dimensional vector in Figure 3.4. Divisive normalisation maintains the direction of the weight vector, whereas subtractive normalisation will tend to push elements of the weight vectors to extreme values. When these normalisation schemes are used in conjunction with modification rules, they can drastically affect the development of a postsynaptic unit's properties such as ocular dominance (Miller & Mackay, 1994; Goodhill & Barrow, 1994).




Figure 3.4: A simple comparison of divisive and subtractive normalisation. The original weight vector $(4\ 2)^T$ is to be normalised such that $w_1 + w_2 = 3$. Under divisive normalisation, the vector becomes $(2\ 1)^T$, which lies in the same direction as the original vector. Under subtractive normalisation however, the vector becomes $(2.5\ 0.5)^T$, which moves the vector closer to the w_1 axis.

These normalisation techniques can be criticised for being too rigid (Swindale, 1996): when normalising the sum of the weights, the weight vectors are always bounded to lie on the surface of a hypercube. Other normalisation techniques are available which are not so strict. For example, as long as the sum S is initially well below the target value T , the weights can be allowed to adapt until $S \geq T$ for a unit. Once the limit is reached, all of the weights for the unit are then frozen at this saturation point (Prestige & Willshaw, 1975). Alternatively, it is possible for the modification rule to incorporate extra terms to account for normalisation (Oja, 1982; Häussler & von der Malsburg, 1983). In these rules, weights are not constrained to a certain sum, but instead gradually converge to a target value. Other modification rules ensure that weights remain bounded simply by making sure that the weights are similar to their inputs (Kohonen, 1982; Durbin & Willshaw, 1987). At each time step, weights are adapted so that they move closer to the current input vector. In this way, as long as the input vectors are bounded, the weights will also remain bounded.

3.6.2 Biological evidence for normalisation at different sites

Biological evidence supports both pre- and postsynaptic normalisation of synapses in neural maps. If a retina innervates a tectum which is larger than normal, each tectal cell still receives retinal inputs, although the density of input to each tectal cell is lower than normal. This indicates limits on the number of connections that each presynaptic cell can make (Hayes & Meyer, 1988a). In

contrast, other experiments have shown that there are limits on the number of contacts each postsynaptic cell can make. First, the addition of an extra eye to innervate the optic tecta of tadpoles did not affect the number and size of connections to tectal cells when compared with normal two-eyed tadpoles (Norden & Constantine-Paton, 1994). Second, the size of retinal axonal arbors varied in accordance with changes in the number of retinal or tectal cells (Xiong, Pallas, Lim, & Finlay, 1994). For details of other related experiments, see (Hayes & Meyer, 1988b; Sabel & Schneider, 1988; Pallas & Finlay, 1991). These experiments show that there are limits on the number of contacts that a cell (either presynaptic or postsynaptic) can make, although whether constraints simultaneously exist at both sites within one system is still unknown.

3.7 Weight modification rules

Modification rules specify how the weights should be updated in response to the state of elements of the network. Such elements include the activation level of units and the value of other weights. To keep the rules fairly plausible with respect to biological principles, the rules should only make reference to quantities that are simple to calculate and available locally to the weights being changed. Many such rules have been created for modelling neural development. In this section, the two main types of modification rule used for modelling the development of topography and ocular dominance are described.

3.7.1 Correlational rules

Most modification rules capture the basic principle of the Hebbian synapse (Hebb, 1949). This states that if a presynaptic cell and postsynaptic cell are both firing at the same time, the strength of the connection between the two cells should be increased. On its own, this rule is not stable for two reasons. First, it will allow weights to increase without bound. It must therefore work in conjunction with normalisation mechanisms (Rochester, Holland, Haibt, & Duda, 1956). Second, it does not explicitly specify when connection strengths should be decreased. Sejnowski (1977) expanded the notion of the Hebbian synapse into a covariance rule to indicate under what conditions the connections should increase or decrease. This rule is shown in Table 3.2, where the rule is broken down into four cases, according to the activation levels of the pre- and postsynaptic units.

$$\Delta w_{ij} = \varepsilon(x_i - \langle x_i \rangle)(y_j - \langle y_j \rangle)$$

Case	Presynaptic activity ($x_i - \langle x_i \rangle$)	Postsynaptic activity ($y_j - \langle y_j \rangle$)	Weight change Δw_{ij}	Biological correlate
1	+	+	+	Hebbian / LTP
2	+	-	-	Homosynaptic LTD
3	-	+	-	Heterosynaptic LTD
4	-	-	+	None

Table 3.2: The four components of the covariance rule. The four cases of the rule are listed according to changes in sign of the presynaptic and postsynaptic components. A possible biological correlate for each case (except case four) is also given (Miller, 1996). Key: LTP – long-term potentiation. LTD – long-term depression.

Biological evidence for the covariance rule

Three out of the four cases of the covariance rule are supported by biological evidence of synaptic plasticity. Case one (active input and output) is verified by studies of long-term potentiation (LTP) (Bliss & Lømo, 1973; Markram, Lubke, Frotscher, & Sakmann, 1997). Case three (inactive input and active output) is illustrated in studies of monocular deprivation: when one eye is closed, the number of cortical cells receiving input from the closed eye is much smaller in comparison to the number of units receiving input from the open eye (Hubel & Wiesel, 1977; Hubel, Wiesel, & LeVay, 1977; Shatz & Stryker, 1978). Finally, evidence for either case two or case four comes from a study which combined monocular deprivation with blockade of postsynaptic activity (Reiter & Stryker, 1988). This caused a reversal to the normal monocular deprivation results: the number of cortical units responding to the closed eye was much greater than the number responding to the open eye.

Variations of the covariance rule

A new rule, developed from a variant of the covariance rule with extra constants, was used in a series of simulations by Linsker (1986a, 1986b, 1986c). These simulations investigated the development of centre-surround and orientation-selective cells within a multi-layer network. If it is assumed that the weights change on a much slower time-scale than the presentation of new inputs, the covariance rule can be averaged over time to derive a modification rule of the form:

$$\frac{d\mathbf{w}}{dt} = \mathbf{C}\mathbf{w}$$

For details of this derivation, see Appendix A.1. In this rule, \mathbf{C}_{ij} is a correlation matrix describing the correlation between pairs of presynaptic units i and j . Hence the covariance rule, which measures correlations between pre- and postsynaptic units, has been reduced to a rule which only depends on correlations between presynaptic units. Using this formulation of the covariance rule is simpler than the original since it needs fewer parameters and does not require an explicit representation of input vectors. Having the rule in this format also allows an eigenvector analysis to predict the development of receptive fields (MacKay & Miller, 1990a, 1990b). (Appendix A.2 describes why an eigenvector analysis is useful for predicting receptive field development.)

Other covariance-like rules

The approach of Montague, Gally, and Edelman (1991) The covariance rule adjusts weights in proportion to the correlation between pre- and postsynaptic activities. A slightly different approach was taken by Montague et al. (1991), using a modification rule which measured the correlation between presynaptic unit activity and the level of concentration of a postsynaptic molecule. Each postsynaptic unit released the postsynaptic molecule in proportion to its activity. This makes the rule very similar to the covariance rule, except for two points. First, the postsynaptic molecule was allowed to diffuse across the postsynaptic sheet. In this way, neighbouring postsynaptic units developed similar concentration levels of the molecule. (This is similar to using lateral connections in the postsynaptic sheet encouraging neighbouring units to have similar activations, although diffusion occurred on a longer time-scale.) Second, the modification rule only specified weight changes for three out of the four cases listed in Table 3.2: when both presynaptic unit activity and postsynaptic level of molecules were both below their thresholds (case four of the table), no

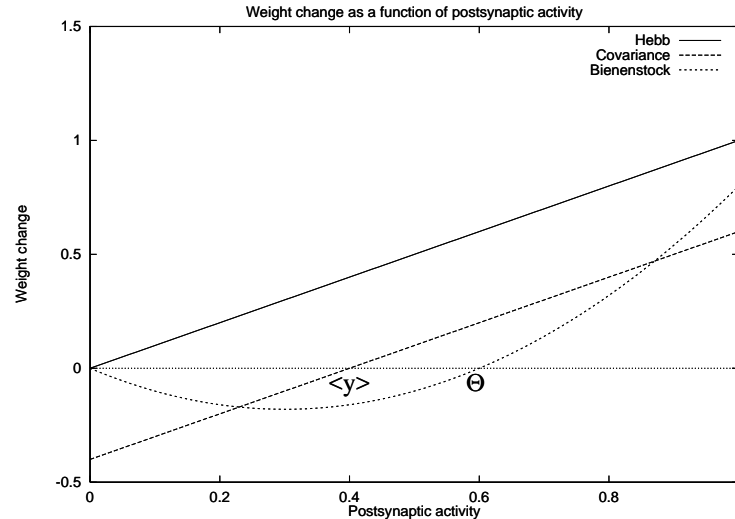


Figure 3.5: Covariance modification rules. The amount of weight change as a function of the postsynaptic activity is plotted for three different correlational rules assuming fixed presynaptic activity ($x_i > \langle x_i \rangle$). The Hebbian rule only specifies weight increments (no negative values for weight change), whereas the other two rules can produce weight decrements if the postsynaptic activity is below threshold ($\langle y \rangle$ or θ). Both Hebbian and covariance rules are linear, whereas the BCM rule is non-linear.

weight change was specified.

The BCM rule Other similar correlational-based modification rules have also been proposed. Unlike the Hebbian and covariance rules, the BCM rule (named after its authors) is non-linear, as shown in Figure 3.5 (Bienenstock, Cooper, & Munro, 1982). The weight change depends not only on a non-linear function of both presynaptic and postsynaptic activity, but also an extra time-varying threshold term θ :

$$\Delta w_{ij} = \epsilon x_i \phi(y_j, \theta_j)$$

A common choice for the postsynaptic function ϕ is $\phi(y, \theta) = y(y - \theta)$, although others can be used as long as they satisfy certain constraints (Bienenstock et al., 1982). The original formulation of the rule also set the threshold θ to be a long-term exponential average of the postsynaptic activity. This has the effect of raising the threshold as the unit becomes more selective to some of its inputs. An alternative scheme for setting the threshold is to start with an initial low value and slowly increase the threshold during development (A. Bray, personal communication). This threshold implements a form of temporal competition between inputs presented at different times, since the amount of weight update is dependent not only on the postsynaptic response, but also on this threshold. As the unit becomes more selective, weights will be increased only for stimuli that make the unit fire above threshold.

The ABS rule As more fine-level details of the nature of synaptic plasticity are being uncovered, modification rules accounting for these details are being created. For example, (Artola et al., 1990) discovered a new aspect of plasticity in rat visual cortex: long-term depression would only occur if the level of postsynaptic activity fell between a low threshold and a high threshold. If the level of

postsynaptic activity was beneath the low threshold, neither LTD nor LTP occurred. This threshold for LTD has been incorporated into a BCM-like rule to create a rule with two thresholds, known as the ABS model (Artola et al., 1990; Hancock, Smith, & Phillips, 1991).

3.7.2 Competitive rules

Competitive rules work in a different manner to correlational rules. First, after each iteration, the rule only updates the weights of a subset of, rather than all, postsynaptic units. At each iteration, the set of units to be updated is often decided by selecting postsynaptic units within a given distance of the ‘winning’ unit. Second, rather than measuring correlations between unit activities, competitive rules work by iteratively making the vector of weights for a postsynaptic unit more similar to the current input vector.

The Kohonen rule

In the Kohonen network, the activation of a postsynaptic unit is defined as the Euclidean distance between the current input vector and the unit’s weight vector. The postsynaptic unit with the smallest activation is declared the winner. All of the units within a certain geometrical distance of the winner have their weights updated with the modification rule.

$$\begin{aligned}
 y_j &= \sum_{i=1}^{N_{\text{pre}}} (w_{ij} - x_i)^2 && \text{activation rule} \\
 \text{Find } m: & \quad y_m \leq y_j \quad \forall j && \text{winning unit} \\
 \Delta w_{ij} &= g(|j - m|, \sigma) \epsilon (x_i - w_{ij}) && \text{weight modification rule} \\
 g(x, \sigma) &= \exp(-x^2 / 2\sigma^2) && \text{neighbourhood function}
 \end{aligned}$$

The size of the neighbourhood around the winning unit is defined by σ , which decays over time to allow the receptive fields to refine gradually during the simulation. As long as the inputs are finite, each weight will always remain within the range of values for the presynaptic unit to which it is connected. The Kohonen rule therefore does not require extra normalisation terms to prevent unbounded weight growth.

Obermayer/Goodhill

A modified Kohonen rule for modelling visual system development was independently developed by Obermayer et al. (1990) and Goodhill (1992). In this rule, activation levels are set using the dot product rule, the winning unit is the postsynaptic unit with the highest output, and the weights are adapted by adding some fraction of the current input vector to the weight vector.

$$\begin{aligned}
 y_j &= \sum_{i=1}^{N_{\text{pre}}} w_{ij} x_i && \text{activation rule} \\
 \text{Find } m: & \quad y_m \geq y_j \quad \forall j && \text{winning unit} \\
 \Delta w_{ij} &= g(|j - m|, \sigma) \epsilon x_i && \text{weight modification rule}
 \end{aligned}$$

Again, the value of σ decays over time to allow for a gradual refinement of receptive fields. This weight update rule specifies only how weights are incremented, since elements of the input vector are always positive. Normalisation constraints across either presynaptic or postsynaptic

units (or both) are therefore required to specify weight decrements which prevent weights increasing without bounds. Both the Kohonen and Obermayer/Goodhill rules are “hard-competitive” rules because at each iteration, only the weights of a subset of postsynaptic units are updated.

The elastic net

The elastic net (Durbin & Willshaw, 1987) combines both a regularisation and a matching term into one rule which is similar to the Kohonen rule. In the elastic net, each postsynaptic unit j has a receptive field of Gaussian shape with width k centred on the input space at w_j . For each input vector \mathbf{x} presented to the network, the activation y of each postsynaptic unit is calculated. The level of activation of a postsynaptic unit is set to be an exponential function of the Euclidean distance between the input vector and the weight vector for the unit. The weight modification rule is made up of two terms: a matching term to make the weights more similar to the inputs, and a regularisation term to ensure neighbouring units have similar weight vectors.

In comparison to the hard competitive rules, the elastic net rule is a “soft-competitive” rule since all weights in the network are updated, although the amount of update is modulated by the normalised activation levels. The equations describing the elastic net are given as:

$$y_j = \exp\left(\frac{-(\sum_{i=1}^{N_{\text{pre}}}(w_{ij} - x_i)^2)^{0.5}}{2\sigma^2}\right)$$

$$\Delta w_{ij} = \alpha \frac{y_j}{\sum_{k \in A} y_k} (x_i - w_{ij}) + \beta k \sum_{k \in N_j} (w_{ik} - w_{ij})$$

N_j = set of nearest neighbours of postsynaptic unit j

A = set of all postsynaptic units

Within the modification rule, the first term is the regularisation term, pushing weight vectors towards input vectors, and the second term is the matching term to encourage neighbouring postsynaptic units to have similar weight vectors. σ is a control parameter of the model which decreases over time to allow the gradual refinement of the weights once the overall order has been established. Like the Kohonen rule, no extra normalisation constraints are required because the weights adapt to become similar to the inputs.

Biological evidence for competitive rules

In comparison to the covariance rules, there is less biological support for the competitive rules. This is because they were devised as more abstract formulations of how neural maps develop (Kohonen, 1982, 1988). However, given the abstractions and assumptions used for other parts of the model (such as approximating synaptic connections between cells with a single number), the competitive rules are no less plausible than the other model components.

3.8 Modelling the retinogeniculate pathway

The first half of this chapter has considered the common mechanisms used in models for replicating various aspects of neural development. For the second half of this chapter, we now focus on previous models that are relevant to the problem of modelling development of the retinogeniculate pathway. We first describe the specific models of the retinogeniculate pathway, before describing

other relevant models which investigate the development of topography and ocular dominance in other visual areas.

3.8.1 Models of the LGN

The retinogeniculate pathway has been simulated before in several models of the mammalian visual pathway. The models generally break down into two main categories: those that are adaptive, and those that are non-adaptive. These two classes of model are now reviewed.

Adaptive models of the LGN

To date, only two adaptive models of the LGN have been produced. The first was by Keesing et al. (1992), which forms a major part of this thesis, and so is discussed in more detail in Chapter 4. The only other model investigating the development of the LGN is a morphogenesis model (Lee & Malpeli, 1994; Tzonev, Schulten, & Malpeli, 1995). This model specifically investigated the role of the blind spot in the the positioning of the laminar transition of the LGN. (Towards the posterior end of the the macaque monkey LGN, the number of layers changes from six to four.)

In the morphogenesis model, each LGN unit was initially assumed to receive input from one retinal cell of random eye (ipsilateral/contralateral) polarity (on/off) and position (distance from fovea). At each iteration, LGN units made tentative movements in random directions within the LGN. If the tentative move did not decrease the energy level of the system, the unit moved back to its old location. Otherwise, the tentative move was probabilistically made. The overall energy of the system was described by equations that described positional, retinotopic and correlational energies. The positional energy favoured units to be located at a certain depth within the LGN. The retinotopic energy encouraged units to be positioned a certain distance along the anterior-posterior axis of the LGN. The correlational energy encouraged retinal units of the same type to be close together and units of different types to be distant. The probability of moving units was slowly reduced over time, which eventually stabilised the position of the units. A retinal blind spot was also added to the LGN model by preventing a column of LGN units from taking invalid ocularity input.

As long as the initial retinotopy was coarsely defined, the LGN settled down into a mature state showing a segregation of units into eye-specific and polarity layers, as well as a smooth retinotopic map. Without the blind spot, the LGN settled showed a clean transition from six to four layers, although the position of the transition varied along the anterior-posterior axis of the model. With the introduction of the blind spot into the model, the laminar transition still occurred, but with much less variation in the anterior-posterior position.

Although this model developed the desired refinement of topography and ocularity within the LGN, most of this refinement was built in to the model through the retinotopic and positional energy functions which measured how close the units were to their desired positions within the LGN. The main focus of the model was instead to show how the blind spot influences the positioning of the laminar transition from six to four layers. Importantly, the model did not investigate the role of neural activity within the retina, an idea central to this thesis.

Non-adaptive models

Several models of cortical development have included a retinogeniculate pathway as well as a geniculocortical pathway. However, the connections between the retina and the geniculate are

usually fixed in advance and do not adapt. For example, Barrow (1987) fixed the retinogeniculate connections to perform difference of Gaussian convolution to remove the DC component from the retinal signal. Wehmeier, Dong, Koch, and Van Essen (1989) used the retinogeniculate pathway to simply provide a “fan-out” of information by connecting each retinal unit to four geniculate units, with no convergence of retinal units onto geniculate units and no transformation of the retinal signal. Suarez, Koch, and Douglas (1995) used a series of low-pass and high-pass filters to model LGN unit responses, but again the weighted connections between the retina and the LGN were fixed.

3.8.2 Related adaptive models

As indicated by the previous two sections, adaptive models of the the retinogeniculate pathway are not common. However, there are many other models which are relevant to this work. Most of these models consider the development of either the retinotectal or geniculocortical pathway. It is likely that the developmental principles underlying the various different visual pathways are similar, and therefore these other models are worth considering.

Retinotectal mappings

The development of topography within the retinotectal pathway has been commonly investigated using computer models (Prestige & Willshaw, 1975; Willshaw & von der Malsburg, 1976; Whitelaw & Cowan, 1981). The main difference between the retinotectal pathway and the retinogeniculate pathway however is the nature of the topographic mapping. Both the retina and the tectum are regarded as two-dimensional sheets, with a point in the retina projecting to a point within the tectum. The retinogeniculate pathway however maps a point within the retina to a column of cells within the geniculate, as described in Section 2.3.3. A further difference between the tectum and the geniculate is that the tectum normally receives input from just one eye, and hence there is no ocular segregation.

Geniculocortical mappings

Issues of ocular dominance and topography have frequently been investigated theoretically by modelling the geniculocortical pathway (Swindale, 1980; Miller et al., 1989; Goodhill, 1992). The main difference again, however, is the nature of the pathway: although the cortex is a three-dimensional sheet of cells with six anatomically defined layers, it is often modelled as a two-dimensional sheet of units. This two-dimensional cortex assumption again produces point-to-point mappings between the pre- and postsynaptic sheets. Some cortical models rely heavily on extensive lateral cortical connections and are not considered further (Somers, Nelson, & Sur, 1995; Suarez et al., 1995).

3.9 Models of topography

In the following three sections the major models of topography, ocularity and their joint development will be discussed. The main mechanisms underlying each model are listed, along with the conditions required for normal development.

In this section, we consider the major developmental models of retinotopic map formation. Most of the early models focused on accounting for certain experimental manipulations to the

visual pathway (such as tectal grafts and mismatch experiments). Some of the relevant biological manipulations will be mentioned, although for more details refer back to the original papers.

Table 3.3 describes the components of the major models for the development of topographic mappings. As can be seen from this table, the various model components can be combined in different ways to investigate retinotopic map development. All models produce a refinement of topography at the single-unit level. For those models not using chemical markers to code for proximity, extra mechanisms are needed to choose between alternative global layouts of the maps, such as whether the left-hand side of the presynaptic sheet should connect to the left-hand or right-hand side of the postsynaptic sheet.

One of the earliest models of topographic map formation was presented by Prestige and Willshaw (1975). The model investigated both group I and group II labels for map development. Group I labels produced an ordered retinotopic map, which was expected given the nature of the affinities between presynaptic and postsynaptic units. Group II labels could also produce an ordered mapping, but only in conjunction with weight saturation to prevent all presynaptic units connecting to postsynaptic units with the highest affinity.

The arrow model of Hope, Hammond, and Gaze (1976) adopted a sorting approach to the problem of retinotopic map formation. It assumed that each presynaptic unit connected to only one postsynaptic unit and that two postsynaptic units could swap presynaptic units if the presynaptic units were in the wrong position relative to the postsynaptic units. Connections between units were also occasionally swapped at random. In the original presentation of the model, Hope et al. (1976) showed that the model does not fully account for map formation by presenting experimental evidence of the reciprocal translocation of two tectal grafts which cannot be reproduced by the model. Explicit group II labels, as well as the capacity for presynaptic units to innervate more than one postsynaptic unit, were introduced in an extension to this model by Overton and Arbib (1982).

Willshaw and von der Malsburg (1976) produced the first model of retinotopic map formation (the neural activity model) using neural activity to encode the distances between presynaptic units. The model used similar principles to those of an earlier model for the development of orientation selectivity (von der Malsburg, 1973). To solve the problem of the map getting caught in local optima (generating piece-wise continuous maps with no overall order), weight modification was initially restricted to a small subset of postsynaptic units known as polarity markers (not to be confused with chemical markers). These polarity markers imposed an overall orientation on the map, to ensure the map always developed with a specific global orientation.

As with the arrow model, the neural activity model did not account for data from graft translocation experiments (Willshaw & von der Malsburg, 1979). This placed the emphasis back on marker-based theories. Marker-based theories suffer from two main criticisms however. First, it is assumed that the pre- and postsynaptic units develop the same set of markers independently. Second, if any presynaptic or postsynaptic units are removed, the markers must be regenerated (or regulated) to allow the map to redevelop.

Both of these criticisms were addressed by the tea trade model (von der Malsburg & Willshaw, 1977; Willshaw & von der Malsburg, 1979), by assuming presynaptic markers can continuously diffuse throughout the presynaptic sheet and also into the postsynaptic sheet. In the initial map

development, there are no markers in the postsynaptic sheet, and so the global orientation of the map is dominated by the initial connection strengths. In regeneration experiments however, postsynaptic units already have markers as a result of the initial development before regeneration. These regeneration experiments produced novel map formations if the markers preferred a different global map orientation to the orientation specified in the regenerated weights.

The tea trade model was abstracted into a set of equations suitable for theoretical analysis by Häussler and von der Malsburg (1983). An eigenvector analysis showed that by suitable selection of model parameters, the eigenvalues of all non-diagonal eigenvectors were negative, leaving just the two diagonal eigenvectors with positive eigenvalues. The diagonal eigenvectors therefore dominate development to produce an ordered retinotopic mapping.

Whitelaw and Cowan (1981) took a different approach by using both neural activity and chemical markers to encode proximity of presynaptic units. In the model, group II labels were interpreted as adhesion coefficients to represent a tendency for presynaptic units to bind to certain postsynaptic units. Weight update was Hebbian based modulated by the adhesion coefficients. This model was later updated to account for more biological data (Cowan & Friedman, 1991). First, it introduced a tendency for neighbouring presynaptic units to stick together to account for the polarity mismatch experiments (Meyer, 1979). Second, it introduced random depolarisation of synapses so that in the absence of any retinal activity, a rough retinotopic map would still form.

The last two models in Table 3.3 were originally developed as models for producing topographic mappings between abstract presynaptic and postsynaptic sheets of differing dimensionalities, not just the one- or two-dimensional mappings normally studied in visual system development (Kohonen, 1982; Durbin & Willshaw, 1987). There are several examples of application of both the Kohonen rule (Cottrell & Fort, 1986; Obermayer et al., 1990; Goodhill, 1992) and the elastic net (Durbin & Mitchison, 1990; Goodhill & Willshaw, 1990) to topographic maps in the visual system, some of which are described in more detail in Section 3.11.

3.10 Models of ocular dominance

Table 3.4 lists the main models of ocular dominance that have been proposed. The first two models in this table are applications of topographic models to the problem of ocular dominance.

The neural activity model (Willshaw & von der Malsburg, 1976) was also applied to the problem of ocular dominance (von der Malsburg & Willshaw, 1976). Common to most other early models of ocular dominance, it was (unrealistically) assumed that units in different eyes were never active at the same time. In this way, two presynaptic units from different eyes should never have their weights simultaneously increased to the same postsynaptic unit. The Hebbian principle of “cells that fire together should wire together” encouraged units from the same eye to connect to the same postsynaptic unit. Postsynaptic normalisation ensured that as postsynaptic units received more weight strength from one eye, weights from the other eye were reduced.

von der Malsburg (1979) adapted the tea trade model to the problem of ocular dominance. In addition to the topographic markers, each presynaptic unit also had ocularity markers that were unique to each eye. All presynaptic markers were fixed but could induce markers into the postsynaptic sheet. In the postsynaptic sheet, only the ocularity markers could diffuse amongst neigh-

Models of topographic map development

Model	Encoding of topography	Correlated output	Modification rule	Constraints	Global topography
Prestige (75)	Group I or II labels.	Group I or II labels.	Affinity between presynaptic and postsynaptic labels.	Saturation of inputs to a postsynaptic unit.	Labels.
Hope (76)	Implicit graded labels impose gradient on presynaptic units.	Implicit in modification rule.	Swapping of connections by order and at random.	Each presynaptic unit innervates only one postsynaptic unit.	Implicit labels.
Willshaw (76)	Retinal inputs.	Mexican-hat fixed lateral interactions †	Hebbian (if postsynaptic unit above threshold).	Postsynaptic normalisation of weights.	Polarity markers and bias in initial weights.
Malsburg (77)	Diffusion of chemical markers through presynaptic sheet.	Induction of markers into postsynaptic sheet and sprouting.	Similarity between markers of presynaptic and postsynaptic units.	Presynaptic normalisation of weights. Normalisation of marker concentrations in each unit.	Bias in initial weights.
Whitelaw (81)	Group II labels and retinal inputs.	Group II labels and nearest-neighbour lateral interactions.	Hebbian (modulated by graded labels).	Pre- and postsynaptic normalisation of weights.	Labels.
Kohonen (82)	Retinal inputs.	Neighbourhood weight update.	Kohonen rule.	None.	Bias in initial weights.
Durbin (87)	Feature vector.	Regularisation term.	Elastic net rule.	None.	Bias in initial weights.

Table 3.3: Summary of the main models of retinotopic map development. †: The term “Mexican-hat fixed lateral connections” describe fixed short-range excitatory lateral connections and long-range inhibitory connections. Full references for each model can be found in the text.

bouring units, under the constraint that the amount of ocularity markers in each postsynaptic unit were fixed, in a similar fashion to postsynaptic weight normalisation. Malsburg noted that in the absence of any postsynaptic weight normalisation, some postsynaptic units would disconnect from the presynaptic sheet (all of the weights to the postsynaptic unit would be zero).

Swindale (1980) used a synaptic density function to quantify how neighbouring synapses from both the same and opposite eye influenced the development of a synapse. The exact form of the function was not crucial, as long as the same-eye function was short-range excitatory and long-range inhibitory (and vice-versa for the opposite-eye function). Fourier analysis of the initial development of the model found that stripe width was determined by the peak frequency in the Fourier transform of the synaptic density function. The model also accounted for the difference between monkey and cat ocular dominance patterns (monkey stripes tend to be more regular than in the cat) by varying the relative strengths of the short- and long-range interactions (Swindale, 1981).

The model by Bienenstock et al. (1982) was primarily applied to orientation selectivity, but was also applied to the development of monocularity of a single postsynaptic unit (it did not consider a group of postsynaptic units). Several different visual environments were simulated. Of most relevance to this thesis, the uncorrelated rearing paradigm (where inputs from each eye are independent) caused the postsynaptic unit to become monocular. In other environments with positive between-eye correlations, the postsynaptic unit remained binocular.

The model developed by Miller et al. (1989) used a form of correlational rule similar to the rule developed by Linsker (1986a). The model investigated various different functions describing the correlations both within and between eyes, and showed that monocular units would develop if the between-eye correlations were either zero or negative. Various normalisation constraints were used in the model, including both postsynaptic normalisation and freezing of weights. Presynaptic normalisation was only necessary when the lateral interaction function was purely short-range excitatory — when the lateral interactions included long-range inhibition, no presynaptic normalisation was required. This model has since been applied to the development of orientation selectivity and the joint development of ocular dominance and orientation selectivity (Miller, 1994; Erwin, Obermayer, & Schulten, 1995).

Analysis of this model showed that the form of normalisation was critical to the outcome of the model. If the eyes were anticorrelated, divisive postsynaptic normalisation was sufficient to produce monocular units. On the other hand, if the eyes were even slightly correlated, then subtractive normalisation must be used to enable the unit to become monocular (Miller & MacKay, 1992; Miller & Mackay, 1994). A perturbation analysis of the correlational model showed however that the addition of small positive correlations between the two eyes was sufficient to prevent the development of monocular units when using a correlational rule (Dayan & Goodhill, 1992; Goodhill, 1992).

The basis of the weight modification rule used by Montague et al. (1991) has already been described on page 39. Diffusion of the postsynaptic molecule is the main mechanism by which neighbouring postsynaptic units interact to create ocular dominance stripes. Inputs were assumed to come from spontaneous waves of activity, although when a wave was present in one eye, the other eye was quiet. Inputs were therefore anticorrelated, and the development of monocular units

was not surprising given the results from earlier models using anticorrelated between-eye inputs (Miller et al., 1989).

Elliot, Howarth, and Shadbolt (1996) adopted a rather different approach to ocular dominance by using a network which allowed constant sprouting and retraction of connections. At each time step, sprouting or retraction of presynaptic axons was performed probabilistically depending on how the sprouting or retraction affected the energy of the system. The energy function was constructed to directly introduce competition between inputs so that no extra normalisation terms were required. In the original model, the two eyes were never correlated, although later work showed that segregation of inputs from different eyes can occur in this model even in the presence of strong between-eye correlations (Elliot & Shadbolt, 1996).

The last entry in Table 3.4 describes a mathematical analysis and simulation of ocular dominance within a competitive network (Bauer, Brockmann, & Geisel, 1997). This was not the first model to use a competitive network for ocular dominance, but the other models also investigated the development of topography, and so are discussed in the next section. Analysis of the Kohonen rule showed that ocular dominance can develop in the presence of between-eye correlations, and that the larger the between-eye correlations, the narrower the width of the ocular dominance stripes. Numerical simulation of the model showed a close fit to the analytical results, verifying the analysis. Additionally, the simulation results were unaffected by the choice of postsynaptic normalisation (divisive or subtractive), in comparison to earlier work by Goodhill (1992).

Models of ocular dominance development

Model	Encoding of ocularity	Correlated output	Modification rule	Constraints	Stripe width
Malsburg (76)	Retinal inputs: correlated within-eye and anticorrelated between-eyes.	Fixed Mexican-hat lateral interactions.	Correlational rule.	Postsynaptic normalisation of weights.	Range of inhibitory lateral connections.
Malsburg (79)	Fixed ocularity markers unique to each eye in presynaptic units. Topographic markers fixed for all units.	Lateral diffusion of ocularity markers in postsynaptic sheet.	Similarity between presynaptic and postsynaptic markers.	Normalisation of markers at each postsynaptic unit. Divisive presynaptic normalisation of weights.	Compromise between topographic and ocularity markers.
Swindale (80)	Synaptic density functions with Mexican-hat profile for synapses from same-eye and inverse Mexican-hat profile for synapses from different eyes.	Modification rule sums over neighbouring synapses using synaptic density functions.	Weight change proportional to convolution of neighbouring synapses and synaptic density functions within and between eyes.	Implicit postsynaptic normalisation of weights. Upper bound on each weight.	Extent of within-eye and between-eye synaptic density functions.
Bienenstock (82)	Retinal inputs: uncorrelated between eyes.	Not considered.	BCM rule.	Time varying threshold.	Not considered.
Miller (89)	Correlation matrix: Gaussian within-eye correlations and zero (or inverse Gaussian) between-eye correlations.	Fixed lateral interactions.	Correlational rule.	Postsynaptic normalisation.	Extent of within-eye correlations and lateral interactions.
Montague (91)	Retinal inputs: anticorrelated waves.	Lateral diffusion of postsynaptic molecule.	Covariance-based rule.	Bounds on each weight. Other normalisations not mentioned.	Not considered.
Elliot (96)	Retinal inputs: anticorrelated between eyes.	Energy function sums over nearest postsynaptic neighbours.	Probabilistic sprouting/retraction according to energy of system.	None (sprouting and retraction are independent).	Not considered.
Bauer (97)	Retinal inputs: correlated within and between eyes.	Neighbourhood weight update.	Kohonen rule.	None.	Between-eye correlations and width of neighbourhood.

Table 3.4: Summary of the main models of ocular dominance development. Items not investigated by the models are marked “Not considered”. Full references for each model can be found in the text.

3.11 Models of the joint development of topography and ocular dominance

The models described in Sections 3.9 and 3.10 consider only the development of either topography or ocular dominance, but not their joint development. In the last few years, several models have been developed that combine both features into one model. These are described in Table 3.5.

Obermayer et al. (1991) combined retinotopy, ocularity and orientation features into a feature vector to study their joint development in a cortical model. This work followed on from earlier work looking at the joint development of retinotopy and orientation (Obermayer et al., 1990). The inputs to the model were a five-dimensional feature vector (x, y, ϕ, q, z) to represent the neural activity upon two retinae. (x and y coded for the centre of mass, z for the ocularity of the stimuli, and ϕ, q for orientation.) The final pattern of the map depended on the relative variance of each feature. When the variance of z was small compared to x and y , the map coded visual space as the primary feature and ocularity was not represented in the map. As the variance of z increased, ocularity became a secondary feature of the map, with space still as the primary feature. This type of map corresponds to the normal pattern of topography and ocular dominance found in the cortex. Finally, when the variance of z was sufficiently high, ocularity became the primary feature in the map, with visual space regulated to a secondary feature.

Goodhill (1992) devised a model similar to the neural activity model (Willshaw & von der Malsburg, 1976), but with a new competitive rule and using inputs with positive between-eye correlations. In the presence of positive between-eye correlations, units would become monocular as long as the postsynaptic normalisation was subtractive. (When divisive postsynaptic normalisation was used, most postsynaptic units remained binocular.) Goodhill also found that as the between-eye correlations decreased, stripe width increased. This prediction was later verified by results from strabismic cats showing wider ocular dominance stripes (Löwel, 1994; Goodhill & Löwel, 1995). Similar results were also obtained from an elastic net model (Goodhill & Willshaw, 1990).

The LISSOM model (laterally interconnected synergetically self-organizing map) extended the previous competitive models by introducing explicit short-range excitatory and long-range inhibitory lateral connections between postsynaptic units and allowing these lateral connections to adapt during development. Using patterns of retinal activity which produced positive between-eye correlations, the LISSOM model developed the usual ocular dominance stripes (Sirosh & Miikkulainen, 1995). In addition, the mature excitatory lateral connections for monocular postsynaptic units typically connected to other units that responded to the same eye, in correspondence with evidence from visual cortex of cats reared with squints (Löwel & Singer, 1992). Finally, this model also showed that stripe width varies inversely with the strength of between-eye correlations.

The analytical description of the tea trade model for forming retinotopic maps (Häussler & von der Malsburg, 1983) has also been applied to the problem of ocular dominance and topographic development (Andrade & Moràn, 1996). Diffusion of activity in both pre- and postsynaptic sheets was responsible for correlating neighbouring pre- and postsynaptic units, with the extent of the diffusion affecting stripe width within the model. Although the model developed topographic mappings, the amount of bias in the initial weights was quite high in comparison to other models and was not the main focus of the model.

Models of the joint development of ocular dominance and topography

Model	Encoding of ocularity and topography	Correlated output	Modification rule	Constraints	Stripe width
Obermayer (91)	Feature vector (x, y, ϕ, q, z) .	Neighbourhood weight update.	Kohonen rule.	None.	z (ocularity feature).
Goodhill (92)	Retinal inputs: positive correlations both within and between eyes.	Neighbourhood weight update.	Obermayer/Goodhill rule.	Divisive presynaptic normalisation and subtractive postsynaptic normalisation. Weight freezing.	Extent of between-eye correlations.
Sirosh (95)	Retinal inputs: positive correlations both within and between eyes.	Adaptive lateral connections.	Hebbian based.	Divisive postsynaptic normalisation of both feedforward and lateral weights.	Extent of between-eye correlations.
Andrade (96)	Correlation matrix: lateral diffusion between uncorrelated presynaptic units.	Lateral diffusion between postsynaptic units.	Correlational rule.	Conservation term in weight-update rule, jointly enforcing pre- and postsynaptic normalisation.	Extent of lateral diffusion.

Table 3.5: Summary of the main models of the joint development of topography and ocular dominance. All models used biases in the initial weights to enforce a global map orientation. Full references for each model can be found in the text.

3.12 Discussion

This chapter has reviewed the main mechanisms and models used to theoretically investigate the problems of topographic map and ocular dominance development. Although the models use a range of different techniques, most of the underlying principles are similar.

3.12.1 Inputs

First, distances between presynaptic units need to be encoded. This is done either through the use of chemical markers within units or by their patterns of neural activity. Units that are close to each other tend to have similar markers/activity patterns, and as the distance between units increases, the similarity decreases. Of particular importance, the network must find a way of discriminating presynaptic units from different eyes. When the two eyes are independently active, units in different eyes are uncorrelated. For the problem of retinogeniculate development, the waves are generated independently in each eye, and so the waves are uncorrelated. However, for those models using visual stimulation as the source of neural activity, units in both the same and different eyes are positively correlated.

3.12.2 Outputs

Second, neighbouring postsynaptic units need to interact with each other so that they can respond to similar regions of the presynaptic sheet. This can be achieved by several different mechanisms, including lateral connections within the postsynaptic sheet, diffusion of activity, or by neighbourhood weight updating in the modification rule.

3.12.3 Modification rules

Third, modification rules adapt the weighted connections between pre- and postsynaptic units according to the basic Hebbian notion that if the two units are firing together, the connection strength between the units should be increased. Some rules only specify weight increments, whereas other rules specify both increments and decrements. However, most modification rules are not solely sufficient to account for the development of neural mappings. Other constraints are additionally required to introduce competition between a set of weights and to prevent unbounded weight growth.

3.12.4 The effect of different weight modification rules and normalisation schemes on topography and ocular dominance

All of the models of retinotopic map formation reviewed in this chapter show that ordered maps form as long as there is an overall bias imposed on the map (either through the initial weights or using chemical markers) and if neighbouring postsynaptic units can interact with each other. Both of these mechanisms are also used in the Keesing model of retinotopic map formation in the LGN (Keesing et al., 1992). However, none of the models reviewed here consider map development when the postsynaptic units are arranged in a higher dimension than the presynaptic units. Although the Keesing model uses a one-dimensional retina and two-dimensional LGN, it only investigates the nature of the retinotopic map within one row of units in the LGN. In the

next chapter, we therefore consider the form of the retinotopic map produced across the whole postsynaptic sheet of the Keesing model.

Most of the previous modelling work has demonstrated that development of monocular units depends on the correlations between eyes, the type of weight modification rule, and the type of normalisation schemes used. First of all, when the two eyes are anticorrelated, all of the models can produce monocular postsynaptic units. This is expected since presynaptic units from different eyes are never simultaneously active. For the correlational rule, an eigenvector analysis of a correlation matrix with between-eye anticorrelations predicts that monocular receptive fields will dominate development (Miller & MacKay, 1992; Miller & Mackay, 1994). Most models additionally require some form of normalisation to keep weights within bounds and to introduce competition. Out of these models, only the tea trade model used presynaptic, rather than postsynaptic, weight normalisation, but even this model also used postsynaptic normalisation of ocular markers (von der Malsburg, 1979). Presynaptic normalisation, when used, is only required to keep presynaptic units connected to the same postsynaptic sheet.

When there are positive between-eye correlations present, the correlational models are not suitable for modelling ocular dominance development. Miller and Mackay (1994) show that correlational rules must use postsynaptic subtractive normalisation, rather than divisive normalisation, but even this does not guarantee the development of monocular units. Perturbation analysis of the correlation rule showed that the addition of small positive between-eye correlations is sufficient to prevent the development of monocular receptive fields with a correlation rule (Dayan & Goodhill, 1992; Goodhill, 1992).

Better results are obtained from models using competitive rules (Goodhill, 1992; Sirosh & Miikkulainen, 1995; Bauer et al., 1997). The first model to demonstrate that monocular receptive fields develop in the presence of positive between-eye correlations also made the important prediction that stripe width varied inversely with the strength of positive between-eye correlations (Goodhill, 1992). However, this model required subtractive postsynaptic normalisation, which was not found to be a requirement in a subsequent competitive model (Bauer et al., 1997). (As Bauer et al. note, this difference in dependence on normalisation results could be a consequence of the size of input stimuli used: whereas Goodhill used stimuli covering the entire retina, Bauer et al. used Gaussian activity patterns covering a much smaller part of the retina.)

Since the spontaneous retinal waves produce no positive between-eye correlations, the choice of a correlational rule by Keesing et al. (1992) seems suitable for the task of modelling retinogeniculate development. In the next chapter we review this model in detail and investigate whether, as with most of the other models reviewed here, postsynaptic normalisation is required for the development of monocular receptive fields.

Chapter 4

An initial model of the retinogeniculate pathway

4.1 Introduction

This chapter introduces a model of retinogeniculate development proposed by Keesing et al. (1992). The first half of this chapter introduces the main components of the model and the visualisation techniques used to analyse network performance. We then present our replication of the results from the original publication. We also investigate the role of the different forms of the initial weight bias upon development.

In the second half of this chapter, the role of the different weight normalisation schemes in the network are investigated. The original model used both pre- and postsynaptic normalisation, although the results presented here indicate that only presynaptic normalisation is required for normal development. This is in contrast to the models reviewed in Chapter 3, where all of the models required postsynaptic normalisation for ocular dominance development. Given the nature of the projection columns within the LGN, we demonstrate that the presynaptic normalisation should be divisively implemented. However, results are presented showing that after certain modifications, subtractive presynaptic normalisation can also generate the required projection columns.

4.2 The initial model

The Keesing model consists of two one-dimensional retina fully connected to a two-dimensional LGN, as shown in Figure 4.1. The reduction in dimensionality of both the retina and LGN was made to keep the computational demands low.

The retinogeniculate model consists of only feedforward connections. There is no modelling of either the feedback pathway from the cortex or of the local inhibitory circuitry within the LGN (both described in Chapter 2). Both the inhibitory circuit and feedback pathway develop later than the feedforward connections, and so are likely to either be non-existent or have little effect in comparison with the feedforward circuitry. First, the lateral inhibition in the cat LGN is first measurable only at day E59, well after the start of segregation. Second, although corticogeniculate synapses are present in the newborn cat, they are not very well defined (Shatz & Kirkwood, 1984), and it is thought that in most mammals (excluding ferrets) the corticogeniculate pathway reaches its target several days later than the geniculocortical pathway (Johnson & Casagrande, 1993).




Figure 4.1: Architecture of the Keesing network. Each retina is one-dimensional and fully connects to the two-dimensional LGN. This represents a coronal slice of the LGN rather than the whole volume of the LGN. Each retina consists of X_{pre} units (normally 50); the LGN consists of $X_{\text{post}} \times Y_{\text{post}}$ units (normally 10×8).

4.2.1 Weight adaptation mechanisms

The aim of the model is to adjust the set of connection strengths in the network to replicate the pattern of connectivity found in the developing retinogeniculate pathway. These connection strengths (or weights) between presynaptic unit i and postsynaptic unit j are represented by the value w_{ij} . Postsynaptic units are assigned a unique index j as shown in Figure 4.2. Several mechanisms are used to adapt the strength of these weights during development:

1. Postsynaptic activations: $y_j = \sum_{i=1}^{N_{\text{pre}}} w_{ij} x_i$

The activity of postsynaptic units is the dot product of the input vector and the unit's weight vector.

2. Covariance rule: $\Delta w_{ij} = \epsilon(x_i - \alpha)(y_j - \beta)$

The covariance rule uses constants (α, β) rather than long-term averages of unit activations as used in the original formulation of the rule (Sejnowski, 1977).

3. Normalisation:

$$w'_{ij} = \frac{w_{ij} \cdot T_{\text{pre}}}{\sum_{k=1}^{N_{\text{post}}} w_{ik}} \quad \text{Presynaptic normalisation}$$

$$w'_{ij} = w_{ij} - \frac{r_s}{N_{\text{pre}}} (T_{\text{post}} - \sum_{k=1}^{N_{\text{pre}}} w_{kj}) \quad \text{Postsynaptic normalisation}$$

In the Keesing model, the presynaptic normalisation is divisive and the postsynaptic normalisation subtractive. The values T_{pre} and T_{post} are the target values for presynaptic and postsynaptic normalisation respectively. For the subtractive normalisation, the parameter r_s controls the rate at which subtractive normalisation is enforced. Unless stated otherwise, this is set to 1.0 to ensure that the sum of weights after normalisation is exactly the target value. Smaller values of r_s can be used to push the weight vector towards (rather than directly onto) the constraint surface.

$$4. \text{ Growth rule: } \Delta w_{ij} = \gamma \sum_{k \in \text{neigh}(j)} w_{ik}$$

The growth rule encourages a retinal unit to increase the synaptic strength to neighbouring LGN units. The set of neighbours of a unit, $\text{neigh}(j)$, consists of those units within a certain geometrical distance of the unit j . To avoid border effects, units at the left and right edges of the LGN are considered to be neighbours using wrap-around.

The growth rule is used probabilistically after every iteration with a fixed probability p_g . The size of the neighbourhood decreases over time. It is assumed here that the radius of the neighbourhood, r , decreases by one unit every t_r epochs until the radius is zero, after which point the growth rule is not used for the remainder of the simulation.

$$5. \text{ Weight bounds: } w_{ij} \geq 0$$

All weights are bounded to be positive. If any of the rules modify the weights such that the weight will become negative, the weight is instead set to zero strength (and can be further modified). Unless explicitly stated otherwise, there is no upper bound placed on individual weights.

4.3 Retinal inputs

In the description of the original model, retinal inputs were assumed to be Gaussian functions of position, although no details were given on the temporal properties of the waves. In the earliest experiments (not reported here), it was assumed that retinal inputs were always anticorrelated (Eglen, 1995). Given that the retinal waves are independently generated within each eye, the anticorrelated waves were replaced with uncorrelated waves generated by the following procedure:

At the start of development, both eyes are quiet. At each time step, if an eye is quiet (no wave present), a new wave is initiated in the eye with probability p_w . Wavefronts are initiated at either the left- or right-hand side of the retina, and the wavefront moves by one pixel per time step to the other side of the retina. Once the wave reaches the other end of the retina (T time steps after the wave has started), the wave disappears and the retina goes into a refractory period for R time steps. After this refractory period, a new wave may then be initiated.

For these experiments, the wave is modelled as a Gaussian function of position (of standard deviation σ_w) from the wavefront centre c :

$$x_i = \exp\left(-\frac{(i-c)^2}{2\sigma_w^2}\right)$$

4.4 Results

In the following sections, we present the results of using the Keesing model under various different conditions. To provide a fair comparison between the different simulations, unless stated otherwise, all parameter values were kept the same. A list of the model parameters, along with their meaning and typical values, are given in Table 4.1.

Parameter	Meaning	Typical value
p_w	Probability of a new wave starting	0.02
σ_w	Standard deviation of the wave	0.85 – 1.00
R	Refractory period between waves	1
T	Number of time steps wave is present on retina	50 (X_{pre})
X_{pre}	Width of each retina	50
N_{pre}	Number of presynaptic units	100 ($X_{\text{pre}} \times 2$ eyes)
X_{post}	Width of the LGN	10
Y_{post}	Height of the LGN	8
N_{post}	Number of postsynaptic units	80 ($X_{\text{post}} \times Y_{\text{post}}$)
T_{pre}	Target sum for weights from each presynaptic unit	1.0
T_{post}	Target sum for weights to each postsynaptic unit	1.25 †
r_s	Enforcement rate of subtractive normalisation	1.0
ϵ	Rate of weight update in covariance rule	0.001
α	Presynaptic threshold in covariance rule	0.1
β	Postsynaptic threshold in covariance rule	0.0125
γ	Constant for growth term	0.1
r	Radius of growth rule	[2,1,0]
p_g	Probability of using growth rule	0.01
t_r	Number of epochs between radius decreasing by one	200

† There is a relationship between the target sums T_{pre} and T_{post} , since the total amount of weight strength in the network must be the same using either presynaptic or postsynaptic normalisation: $N_{\text{pre}} \times T_{\text{pre}} = N_{\text{post}} \times T_{\text{post}}$

Table 4.1: Parameters and typical values used in the Keesing network simulations.

4.4.1 Visualisation of results

Several different types of plots are used in the following sections to visualise network development. Here are brief descriptions of each type of plot:

- The raw weight matrix.

The weight matrix is displayed as a raster image. One row of the image corresponds to all of the weights for a postsynaptic unit. One column of the image corresponds to all of the weights from a presynaptic unit. This is illustrated in Figure 4.2. The intensity of each pixel in the image is proportional to the magnitude of the corresponding weight. To enable the full range of different magnitudes of weights to be displayed in different matrix plots, the brightness of each pixel is scaled according to the minimum and maximum values of weights for each plot. Although this prevents a direct comparison of images from different

networks, in practice the range of weights in most experiments is roughly constant. (The two notable exceptions to this are shown in Figure 4.12, in the absence of any normalisation, and Figure 5.4, when the probability of waves being generated is very small.)

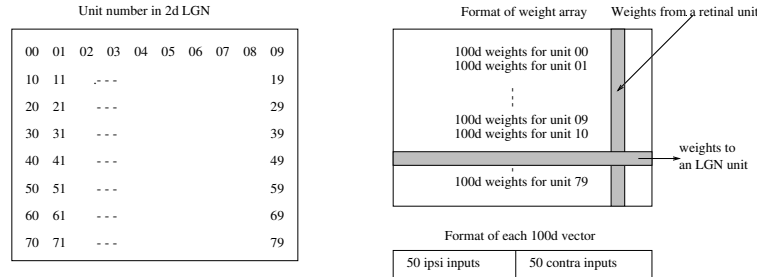


Figure 4.2: Structure of the weight matrix. Each LGN unit is numbered as shown, from 0 to 79. The right-hand figure shows how the weights are displayed: row n of the weight matrix corresponds to the 100-dimensional vector for LGN unit n . Each weight vector is subdivided into the 50 weights for the ipsilateral eye and 50 weights for the contralateral eye.

- Ocular dominance plots.

Hinton diagrams are used to indicate the degree to which each postsynaptic unit receives input from the left and right eye. Each postsynaptic unit within the two-dimensional postsynaptic sheet is represented by a box. The size of the box represents the degree of monocularity, and the colour of the box represents the dominant eye.

$$t_j^{\text{left}} = \sum_{i=1}^{X_{\text{pre}}} w_{ij} \quad t_j^{\text{right}} = \sum_{i=X_{\text{pre}}+1}^{2X_{\text{pre}}} w_{ij}$$

$$z_j = \frac{t_j^{\text{left}}}{t_j^{\text{left}} + t_j^{\text{right}}} - \frac{1}{2}$$

If z_j is greater than zero, most of the input for postsynaptic unit j comes from the left eye, and so the box is coloured black. Otherwise, most of the input comes from the right eye, and the box is coloured white. The size of the box is proportional to the absolute value of z_j , scaled such that the maximum box corresponds to $z_j = \pm 0.5$, which indicates the unit receives all of its input from one eye. If the sum of weight strengths to a postsynaptic unit is less than a certain value (0.005 for the simulations presented here), the postsynaptic unit is defined as “dead”, and is represented by a grey circle in the Hinton diagram (as shown for example in Figure 4.9).

- Topographic plots.

The topographic plots show the centre of mass and the standard deviation of the weight vector for each postsynaptic unit. This assumes that the weight vector has a unimodal distribution. This is obviously not the case when the postsynaptic unit is binocular, but during development, as units become more dominant towards one eye, weight vectors become more unimodal. This progressive development towards a unimodal distribution is illustrated for a typical weight vector in Figure 4.3.

To reduce the effect of the non-dominant eye on the centre of mass measurements, we take the centre of mass and standard deviation for the subset of the weight vector from the eye providing the dominant input to the unit (given by z_j). The centre of mass, \bar{x}_j , and the standard deviation, s_j , for each postsynaptic unit are given by:

$$s_j = \left(\frac{1}{m_j} \sum_{i=1}^{X_{\text{pre}}} ((i - \bar{x}_j)^2 w'_{ij}) \right)^{0.5}$$

where

$$m_j = \sum_{i=1}^{X_{\text{pre}}} w'_{ij} \quad \text{and} \quad \bar{x}_j = \frac{1}{m_j} \sum_{i=1}^{X_{\text{pre}}} i w'_{ij}$$

$$w'_{ij} = \begin{cases} w_{ij} & \text{if } z_j \geq 0 \\ w_{(i+X_{\text{pre}})j} & \text{if } z_j < 0 \end{cases}$$

The dummy variable w'_{ij} for the weights allows us to sum over weights from either the left or the right eye, depending on the value of z_j .

Figure 4.3 shows the development of one postsynaptic unit which demonstrates the usefulness of the centre of mass and standard deviation measures. To allow easy comparison of the topographic plots with the raw weight matrices, the postsynaptic unit number is represented on the Y-axis (increasing in value down the page) and the centre of mass is represented along the X-axis (increasing to the right of the page). Solid and dashed lines are used in the topographic plot to indicate that the postsynaptic unit is dominant for the left or right eye respectively.

The topographic plots are not ideal for the Keesing network, due to the periodic boundary conditions imposed on the X dimension of the network. Postsynaptic units that respond to presynaptic units at the edges of the presynaptic sheets produce artificially high standard deviation values. Despite this limitation, the plots are still useful and are therefore included.

- Projection column plots.

Projection column plots show how the centre of mass of the receptive field of postsynaptic units varies when moving vertically down through the layers of the LGN. (These plots are similar to those used by Sanderson (1971a) to show how the visual field maps onto the LGN.) The centre of mass of each postsynaptic unit is plotted for each unit. Lines are then drawn between points of neighbouring postsynaptic units that are within the same column of the LGN as long as the postsynaptic units receive most of their input from the same eye. These plots are drawn using the same axes as the topographic plots (postsynaptic unit along the Y-axis and centre of mass along the X-axis).

- Projective field plots.

Most of the visualisation methods focus on showing the properties of postsynaptic units. Projective field plots focus on the presynaptic units instead by showing how one presynaptic unit connects to all of the postsynaptic units. The weight vector is plotted with the index of the postsynaptic unit on the X-axis and the magnitude of each weight on the Y-axis.

In addition to these plots, which show the state of the network at one particular time during development, two types of plot are used to summarise how the network develops over the time course of the simulation:

- Refinement of receptive field size.

The standard deviation of the centre of mass for a postsynaptic unit's weight vector can be used as a rough indicator of the width of the unit's receptive field. These widths can be averaged over all postsynaptic units to produce a "mean receptive field width", which is plotted at various stages of development. Error bars for each point indicate ± 1.0 standard deviation of the receptive field width. An example is given in Figure 4.5(a).

- Development of monocularity.

This plot shows the average value of the monocularity index z_j for all postsynaptic units at various points during development for both the left ($z_j > 0$) and right ($z_j < 0$) eyes. Error bars for each point indicate ± 1.0 standard deviation of the z_j values. An example is given in Figure 4.5(b).

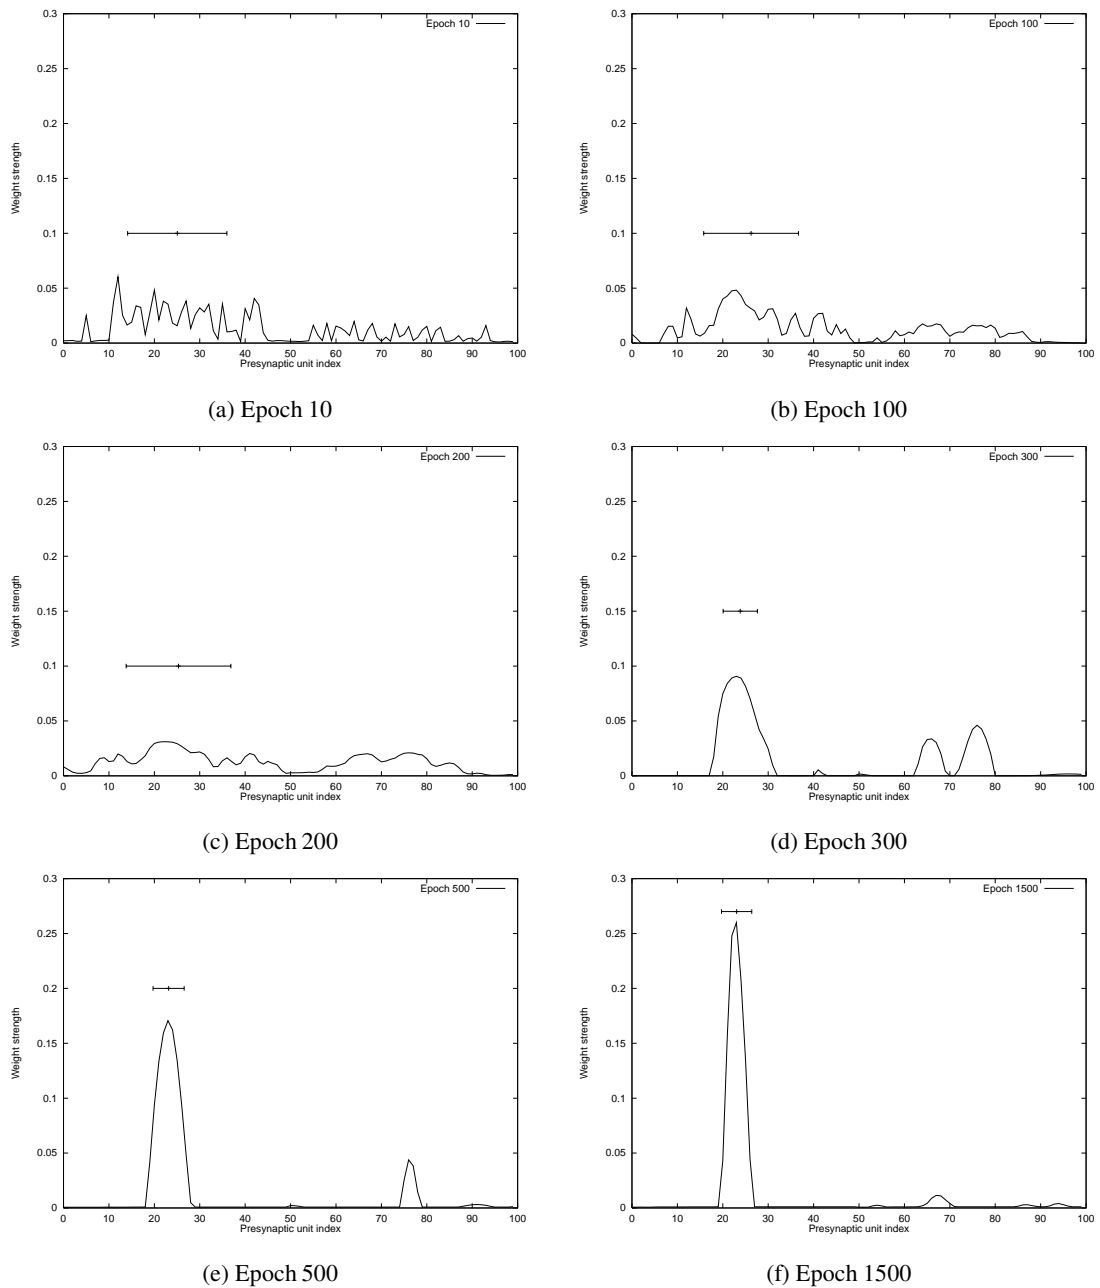


Figure 4.3: The refinement of receptive field width for a single postsynaptic unit. Plots (a) to (f) show the weights for one postsynaptic unit at different times during development. The error bar above the plot shows the centre of mass of the weight vector for the dominant eye ± 1.0 standard deviation. As time progresses, it can be seen that the receptive field refines to a small region of the ipsilateral eye.

4.5 The role of initial weight biases in development

In this section, we consider the effects of using various amounts of bias in the initial weights of the network to see how development is affected. Providing an initial bias in the network's weights is a common technique to prevent the network getting stuck in local minima and to provide a global orientation for retinotopic maps (Willshaw & von der Malsburg, 1976). In this section we present the results of three sets of experiments, gradually increasing the amount of initial weight bias in successive experiments:

1. No initial weight bias.
2. Earlier arrival of contralateral axons.
3. Earlier contralateral arrival plus retinotopic bias.

4.5.1 No initial weight bias

In the absence of any bias in the initial random weights, the modification rules have limited success in refining the weights to mirror the development of the retinogeniculate pathway. Figure 4.4 shows the network at various times during development. In the early stages, the postsynaptic units have wide receptive fields with little dominance of one eye over the other. As development continues, the receptive fields refine in two ways. First, as shown by the topographic plots, the unit receives input from a smaller part of the retina. This decrease in receptive field during development is summarised in Figure 4.5(a). Second, from the ocular dominance plots, the units also become more monocular, responding much stronger to one eye. The progressive development of monocularity is fairly smooth throughout the time course of the simulation, as shown in Figure 4.5(b)


By the end of development, most postsynaptic units have small, monocular receptive fields. However, there is little global organisation. Although nearest neighbour units have similar receptive field positions (two or three neighbouring units cluster together in the topographic plot), there is no global ordering of the receptive fields. This is shown by the many large jumps in receptive field position for neighbouring postsynaptic units in the topographic plots. It is also clearly seen in the projection columns plot shown in Figure 4.5(c): individual columns should be restricted to a small vertical region of the plot, but instead each column occupies a wide range of the input space.

As well as the absence of global topographic order, the final overall pattern of ocular dominance is very different to that found in the retinogeniculate pathway. The network has failed to fully segregate postsynaptic units into two large groups according to the eye of origin. Instead, the pattern of monocular dominance resembles the stripes found in the cortical areas, with a small group of neighbouring units responsive to one eye and then another group responsive to the other eye. Most of the units however have become strongly monocular, with only six units in the simulation presented in Figure 4.4 remaining binocular. All of these six units are surrounded by left- and right-eye dominant units. This is similar to other models of ocular dominance where binocular units are only found at the borders of stripes between alternating eyes (Goodhill, 1992).


It is interesting to compare these results with the innervation pattern seen when ferret retinal inputs are rerouted into the medial geniculate nucleus (the thalamic region which normally receives auditory input) (see Figure 10b of Roe, Garraghty, Esguerra, & Sur, 1993). The rerouted

ferret retinal inputs cluster into patches rather than layers. This evidence can be interpreted as suggesting that extra information must be provided for the formation of layers, and it is likely that this information is specified by activity independent mechanisms (Sur, 1995; Angelucci, Clasca, Bricolo, Cramer, & Sur, 1997). This is investigated in the next set of experiments.








Topography plot



Weight matrix



Ocular dominance plot




(a) Epoch 100


(b) Epoch 600

(c) Epoch 1500


Figure 4.4: Network development with no initial weight bias – 1: topography and ocularity. The state of the network is shown after (a) 100 epochs, (b) 600 epochs and (c) 1500 epochs.



(a) Refinement of receptive field width



(b) Development of monocularity



(c) Projection columns

Figure 4.5: Network development with no initial weight bias – 2: refinement of receptive fields, monocularity and final projection columns.

4.5.2 Earlier arrival of contralateral axons

Contralateral axons innervate the cat LGN a few days earlier than the ipsilateral axons (Shatz, 1983). This gives the contralateral inputs a competitive advantage over the ipsilateral inputs for innervating further into the LGN. In these simulations, it is always (arbitrarily) assumed that the right eye corresponds to the contralateral eye.

The competitive advantage for the contralateral eye is modelled by biasing the initial weights of the network so that the contralateral inputs initially innervates units higher in the LGN than the ipsilateral eye. For example, the initial weights between units of the left eye and postsynaptic units in the bottom two rows of the LGN are set to random values, with the remaining weights from the left eye initially set to zero. The weights from the right eye are initialised in a similar fashion, so that connections between the right eye and units in the bottom four layers of the LGN are set to random values, and the rest set to zero. With this initial setup, units in the bottom two rows (rows seven and eight) of the LGN initially receive input from both left and right eyes, and so are binocular units. Units in rows five and six of the LGN are initially monocular since they receive input from only the right eye. All of the units in the top four rows are initially dead since they receive no input. This is illustrated in Figure 4.9.

With this bias in favour of the contralateral eye, the network develops in a more ordered way in comparison to the networks without any initial weight bias. Figure 4.6 and Figure 4.7 show the development of topography and ocularity in this network. For this set of experiments, two slightly different networks were compared. The first network used both pre- and postsynaptic normalisation, whereas the second network used only presynaptic normalisation. This allowed us to examine network development under different normalisation conditions.

For both networks, the size of the receptive fields refined gradually throughout development. The final topographic mapping however was still only ordered on a local scale, rather than a global scale, as shown in the topography plots in Figure 4.6. In both networks, the topographic mapping is slightly more ordered for the left eye than for the right eye. The effect of the normalisation with respect to topography is most clearly seen in the projection column plots in Figure 4.7. Although some projection columns still vary over the majority of the input space, the projection columns in the network with just presynaptic normalisation tend to be more ordered than in the network with both pre- and postsynaptic normalisation.

The strongest effects of the normalisation in the network however are seen when looking at the development of monocularity in the network. First, the postsynaptic units reached a high level of monocularity much sooner during development when using just presynaptic normalisation than with normalisation at both sites. Second, the network with only presynaptic normalisation formed the desired pattern of ocular dominance: segregation into two eye-specific layers, each layer occupying four consecutive rows of the LGN. All units in this network are strongly monocular. In comparison, the network with both pre- and postsynaptic normalisation formed ocular dominance stripes. (Although the ocular dominance pattern is stripy, there is some overall layering in evidence however: most of the units responsive to the left eye are found in the bottom four rows of the LGN, and vice-versa for the right eye.) The network with both pre- and postsynaptic normalisation also has several binocular units, again at the borders of regions responding to different eyes.

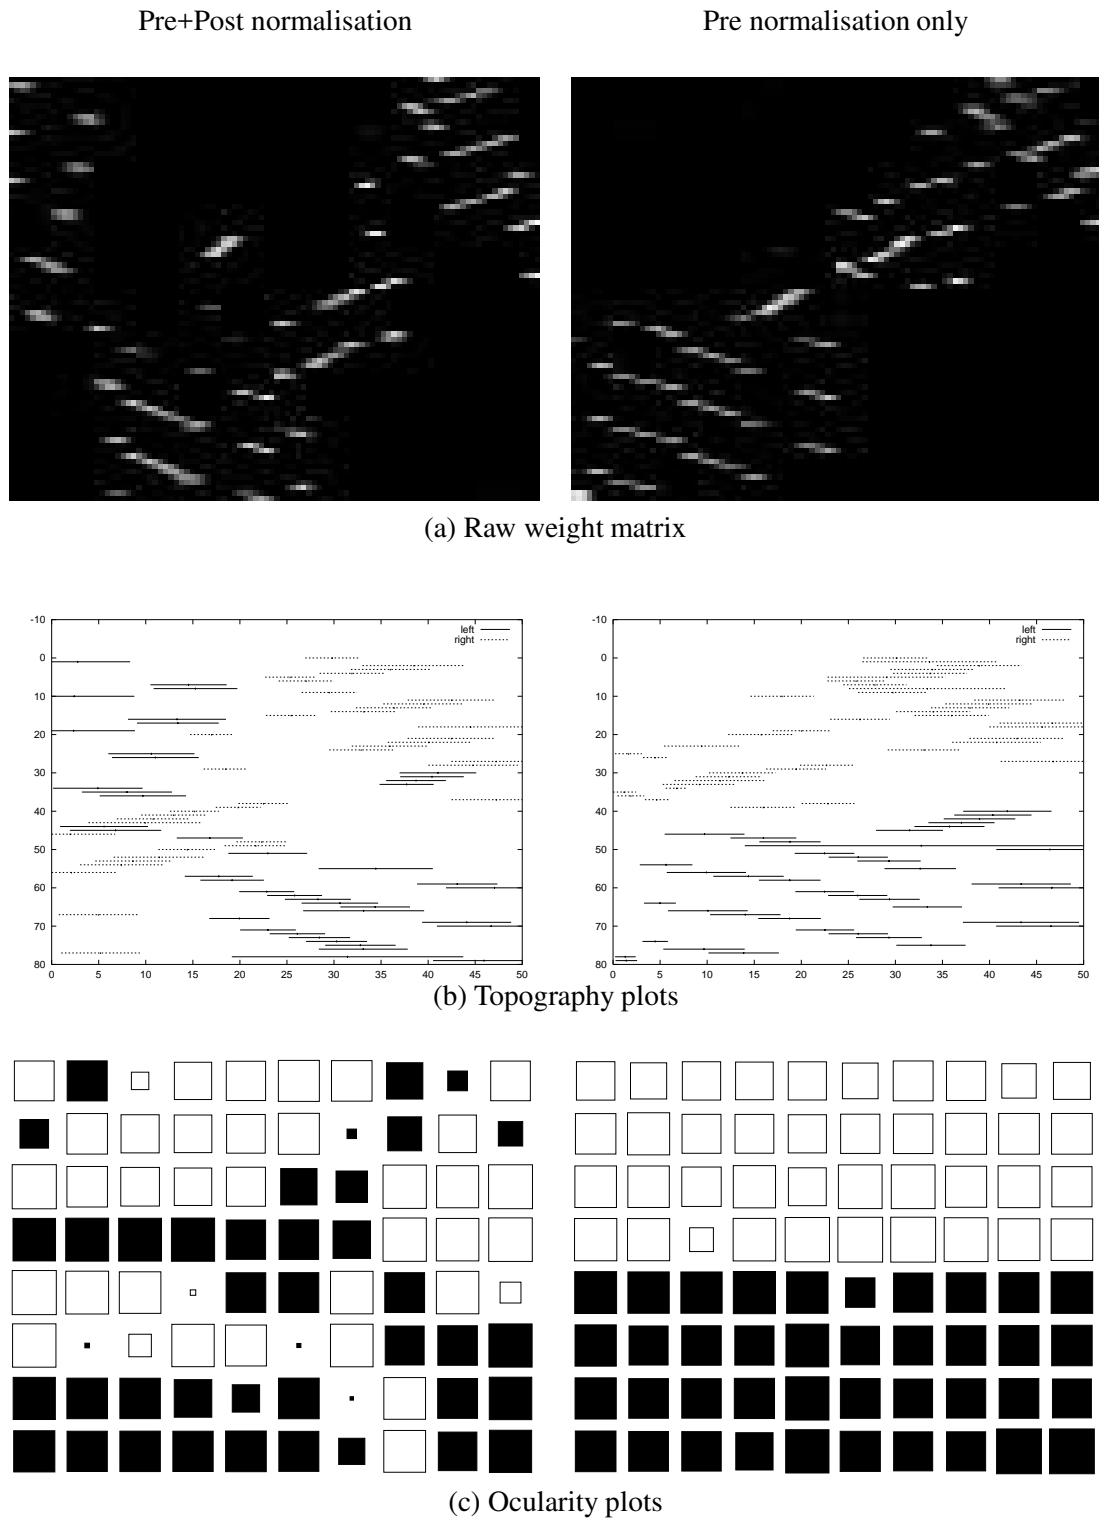


Figure 4.6: The effect of contralateral bias in initial weights – 1: topography and ocularity. The left column shows the results of development using both presynaptic and postsynaptic normalisation; the right column shows the effect of using just presynaptic normalisation.

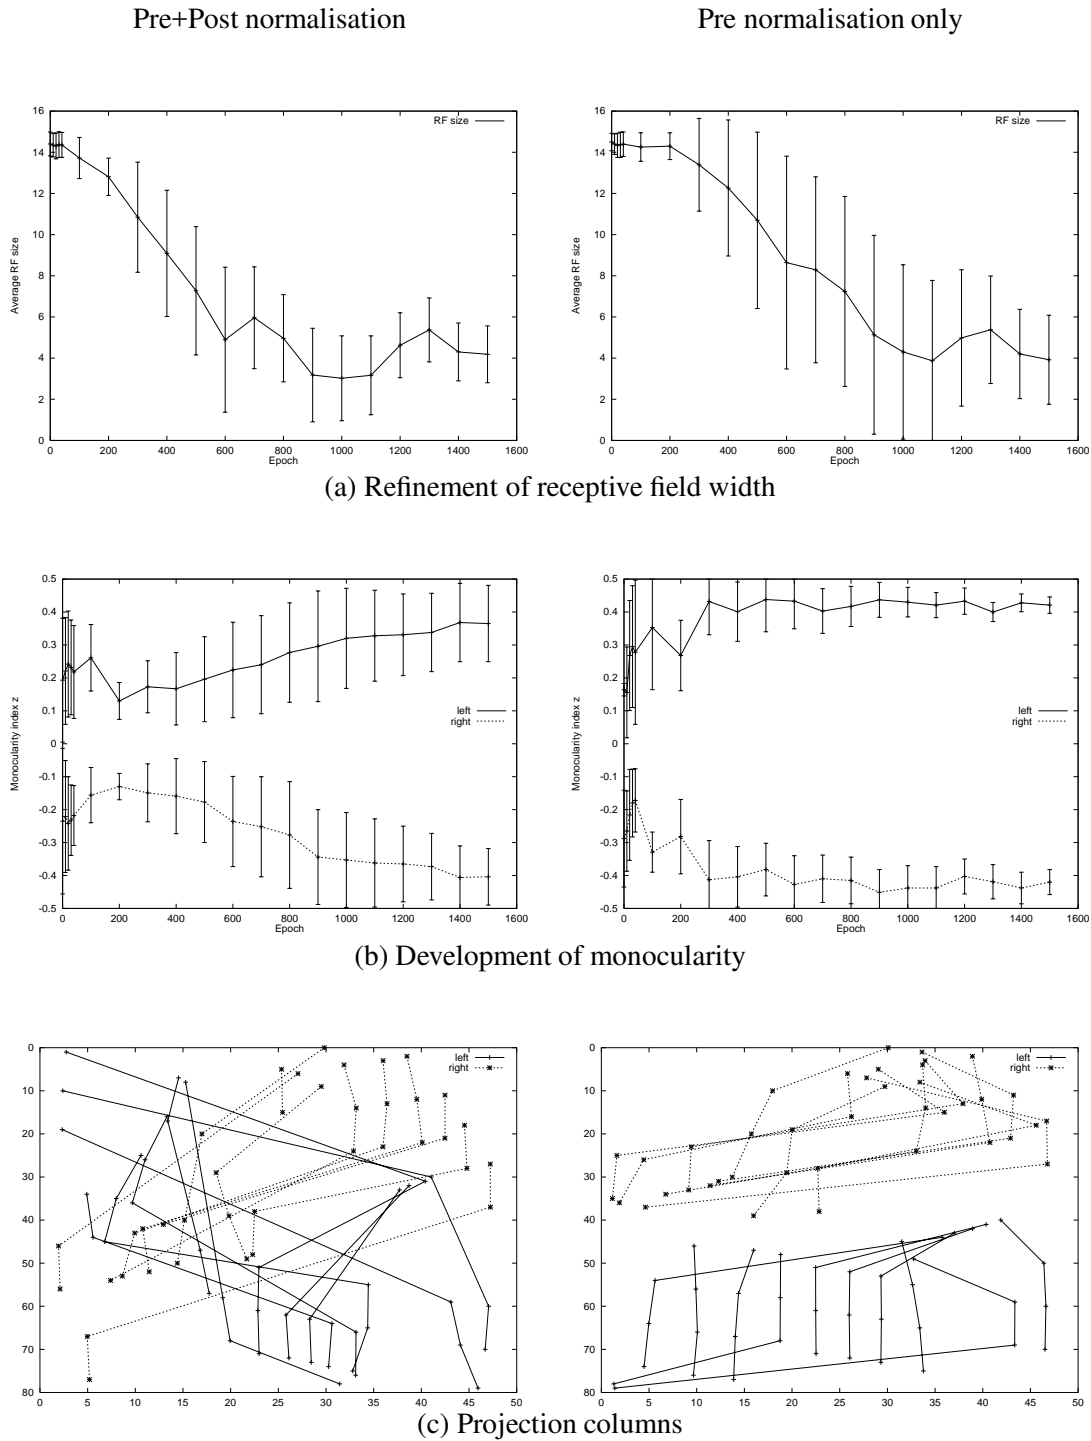


Figure 4.7: The effect of contralateral bias in initial weights – 2: refinement of receptive fields, monocularity and final projection columns. The left column shows the results of development using both presynaptic and postsynaptic normalisation; the right column shows the effect of using just presynaptic normalisation.

4.5.3 Earlier arrival of contralateral axons and topographic bias

As well as imposing a bias in the weights to favour the contralateral inputs going to the top rows of the LGN, we can also impose a topographic bias to ensure that the left end of each retina connects to the left end of the geniculate. This bias can be implemented by assuming that postsynaptic units initially receive input from some subset of all possible presynaptic units. This topographic bias is justified on two grounds. First, some form of topographic bias must be imposed in the retinogeniculate pathway since the mapping of visual space into the geniculate is consistent across animals. If there were no extra mechanism controlling topographic map formation, the mapping of elevation and azimuth into the LGN could be swapped in different ways (such as the superior visual field mapping into the anterior region of the LGN rather than the usual posterior region as shown in Figure 2.3). The second justification for topographic bias is that it is highly unlikely that an immature retinal axon could extend to branch throughout the entire LGN (or for the geniculate dendrite to extend enough to initially receive input from all of the incoming retinal axons). When retinal ganglion axons first innervate the geniculate, they are fairly simple in form and confined to a small region of the LGN, indicating that the axons only contact a few, rather than most of the LGN units (Sretavan & Shatz, 1986). Willshaw and von der Malsburg (1976) also note that rather than allowing each presynaptic unit initially to contact all postsynaptic units, in reality there is probably a more restricted arbor width.

Form of the initial retinotopic bias

Topographic bias is introduced into the initial weights for a set of inputs from one eye to a row of postsynaptic units by setting several consecutive weights to zero. The weight elements set to zero varies systematically across the units within one row of the LGN to provide an overall topographic bias on the initial weights, as shown in Figure 4.8. The degree of initial topographic bias is controlled by a parameter b . This parameter in turn defines two values, o and d (shown in Figure 4.8), which are used to calculate the weight bias:

$$o = \frac{X_{\text{pre}} - b}{2}$$

$$d = \frac{X_{\text{pre}}}{X_{\text{post}}}$$

$$l_x = (x \times d) + o$$

Using these equations, elements $[l_x], [l_{x+1}] \dots [l_{x+b-1}]$ of the weight vector for the postsynaptic unit in column x (and some fixed row) are set to zero. The notation $[l] = l \text{ modulo } X_{\text{pre}}$ is used to provide the periodic boundary conditions. All other weights are set to random values.

Results using retinotopic bias in initial weights

Figure 4.9 shows the development of topography and ocularity in a network whose initial weights were biased for both ocularity and topography. The initial weights are shown in Figure 4.9(a). The ocularity bias was imposed by the left eye innervating the bottom two rows of the LGN and the right eye innervating the bottom four rows (This experiment will be called L2R4, after the ocularity bias). The topographic bias was created using a value of $b = 20$ for weights from the left eye to units in row seven of the LGN and for weights from the right eye to units in row five of the LGN. Weights for all other rows were not topographically biased. Due to the topographic bias for




Figure 4.8: The nature of the topographic bias in the initial weights of one row of LGN units. The top figure shows how the initial bias is calculated. Row x of the diagram corresponds to the set of weights from one eye to the postsynaptic unit in column x (and some fixed row) of the LGN. The portion of the weights coloured grey are the elements of the weight vector which are set to zero; all other elements are set to random values. The bottom figure shows an example set of weights biased using this procedure (with $b = 10$, $X_{pre} = 50$, $X_{post} = 10$).

the left eye, units in row seven of the LGN are slightly dominant for the right eye rather than being equally responsive to both eyes.

Early in development, weights from the right eye tend to move up into the top rows of the LGN, which in turn gives the left eye an advantage in innervating the lower half of the LGN. Hence the layered patterning is in evidence by around 40 epochs, although many units are still binocular, as shown in Figure 4.9(b). The topographic bias has also spread from row seven of the left eye and row five of the right eye into all of the other rows of the LGN, as can be seen most clearly in the weight matrix.

After 100 epochs, the network has settled into two eye-specific layers, each with the same number of postsynaptic units. The only binocular units left are found in the boundary between these two regions (Figure 4.9(c)). At this stage, the network has arranged its global layout, and for the rest of the simulation it then refines the monocularly and receptive field width. Only after the growth rule has finished (after 600 epochs when the radius of the growth rule reaches zero) does the network begin to convert its binocular units at the layer boundaries into monocular units. The receptive fields also gradually refine to receive input from only a small region of the retina.

At the end of development, the network has (almost) produced the desired patterns of ocular dominance and topography. The network has settled into two equal-sized layers of units responding to the two eyes (although at the border between the two layers, four units are wrongly positioned). Each row of the LGN has a smooth retinotopic mapping of visual space. These maps of visual space are aligned throughout the depth of the LGN, as shown by the projection columns of Figure 4.11. (The errant unit in the rightmost projection column for the right eye of Figure 4.11 is due to the periodic boundary conditions in the X dimension of the LGN.)

A similar experiment was performed with slightly different initial weights. This time, the left eye innervated the bottom four rows of the LGN, and the right eye innervated all eight rows of the LGN (This experiment will be called L4R8). To compensate for the increased ocularity

bias, the amount of topographic bias was reduced to $b = 10$. The initial weights, along with subsequent development of the network, are shown in Figure 4.10. Due to the increased ocularity bias, the network has already settled into the two eye-specific layers quite early in development, although the topography is not so well defined in comparison to the previous experiment. The final state of the network is slightly better than in the previous experiment, with a clean segregation of postsynaptic units into two regions of four rows each, smooth topographic mapping in each layer of the LGN, and ordered projection columns.

Figure 4.11 shows the refinement of monocularity and topography over time. Although the L4R8 experiment starts off with more monocular units, the growth rule ensures that in the early stages of development, units are less monocular (explaining the tendency for the average z to initially return towards zero rather than staying at extreme values). This decrease in z may also account for the sudden increase in the average receptive field width for the L4R8 experiment. The overall patterns of development for the two experiments however are pretty similar. The two sources of weight bias can therefore be balanced against each other: if there is less ocularity bias then the topographic bias needs to be increased slightly.

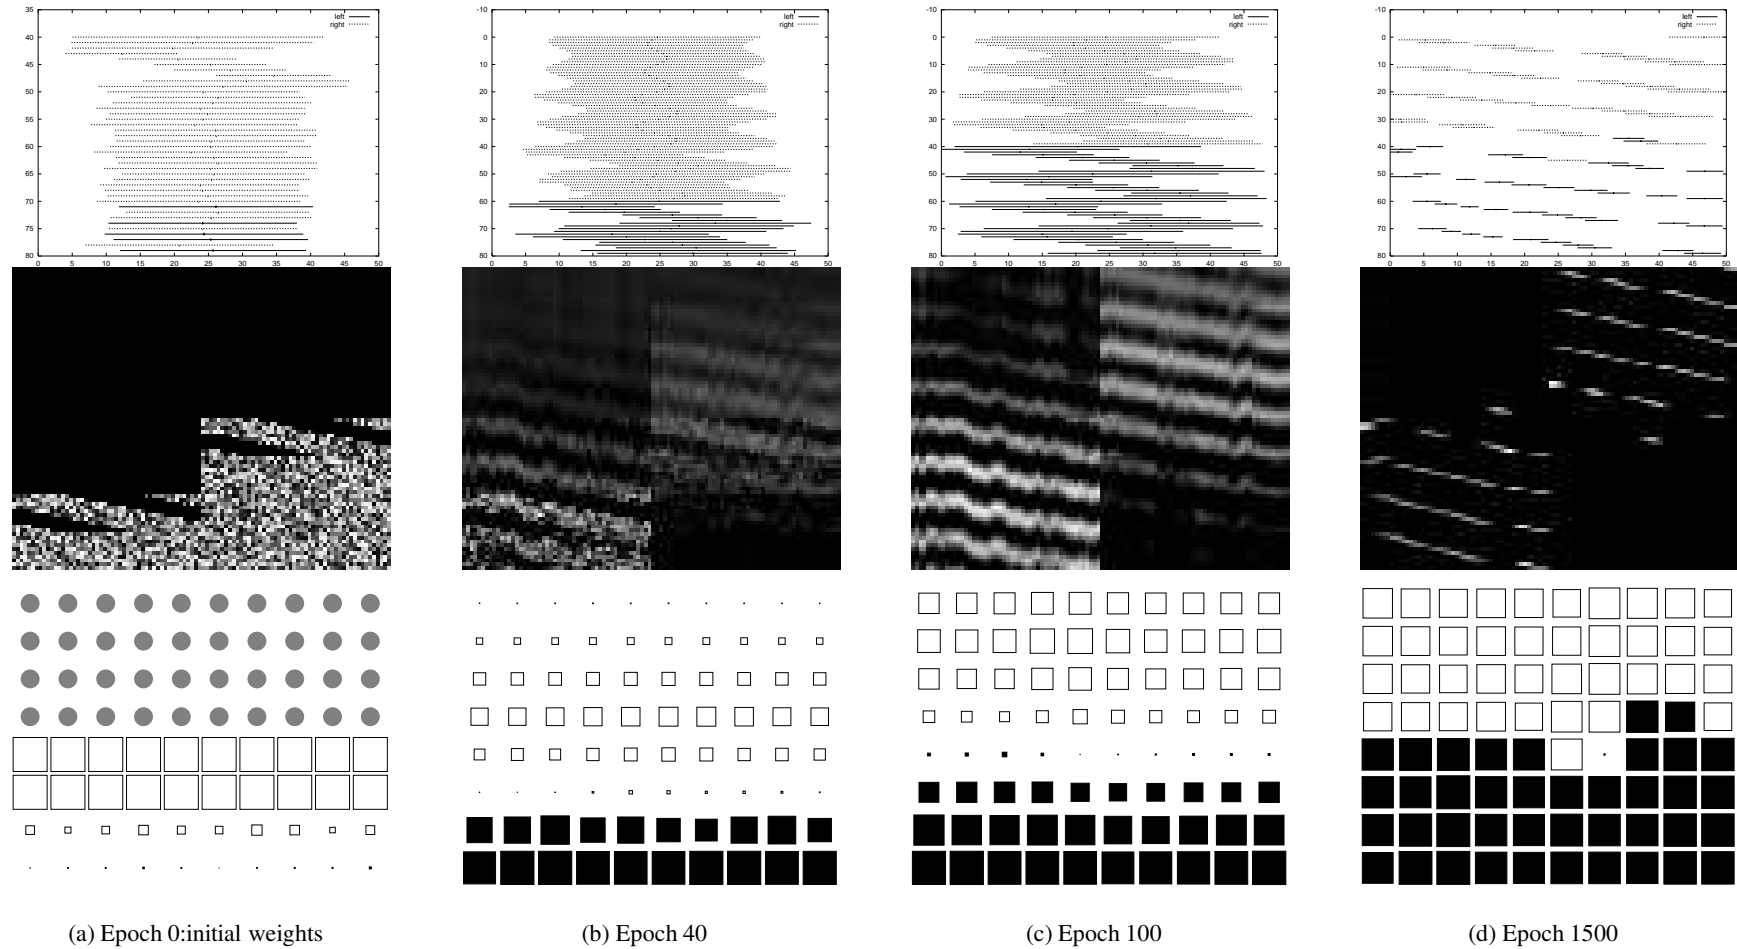


Figure 4.9: The effect of ocular and topographic weight bias on development – 1: left eye innervates two rows, right eye innervates four rows (L2R4).

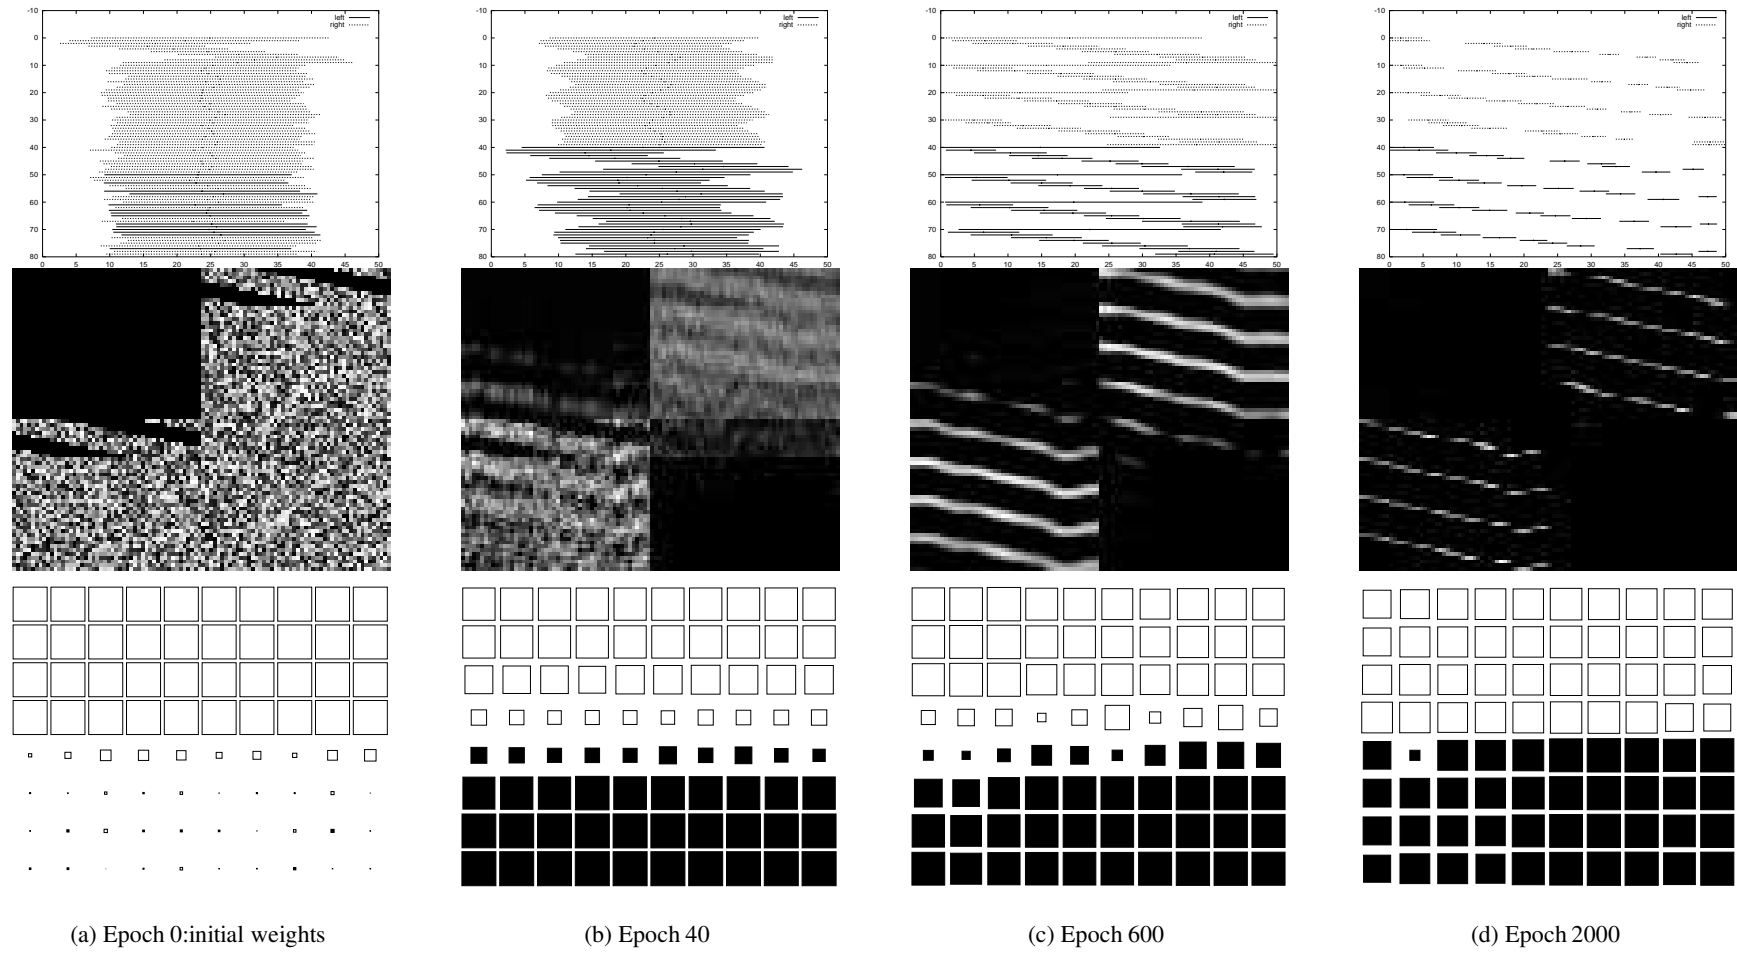


Figure 4.10: The effect of ocular and topographic weight bias on development – 2: left eye innervates four rows, right eye innervates eight rows (L4R8).

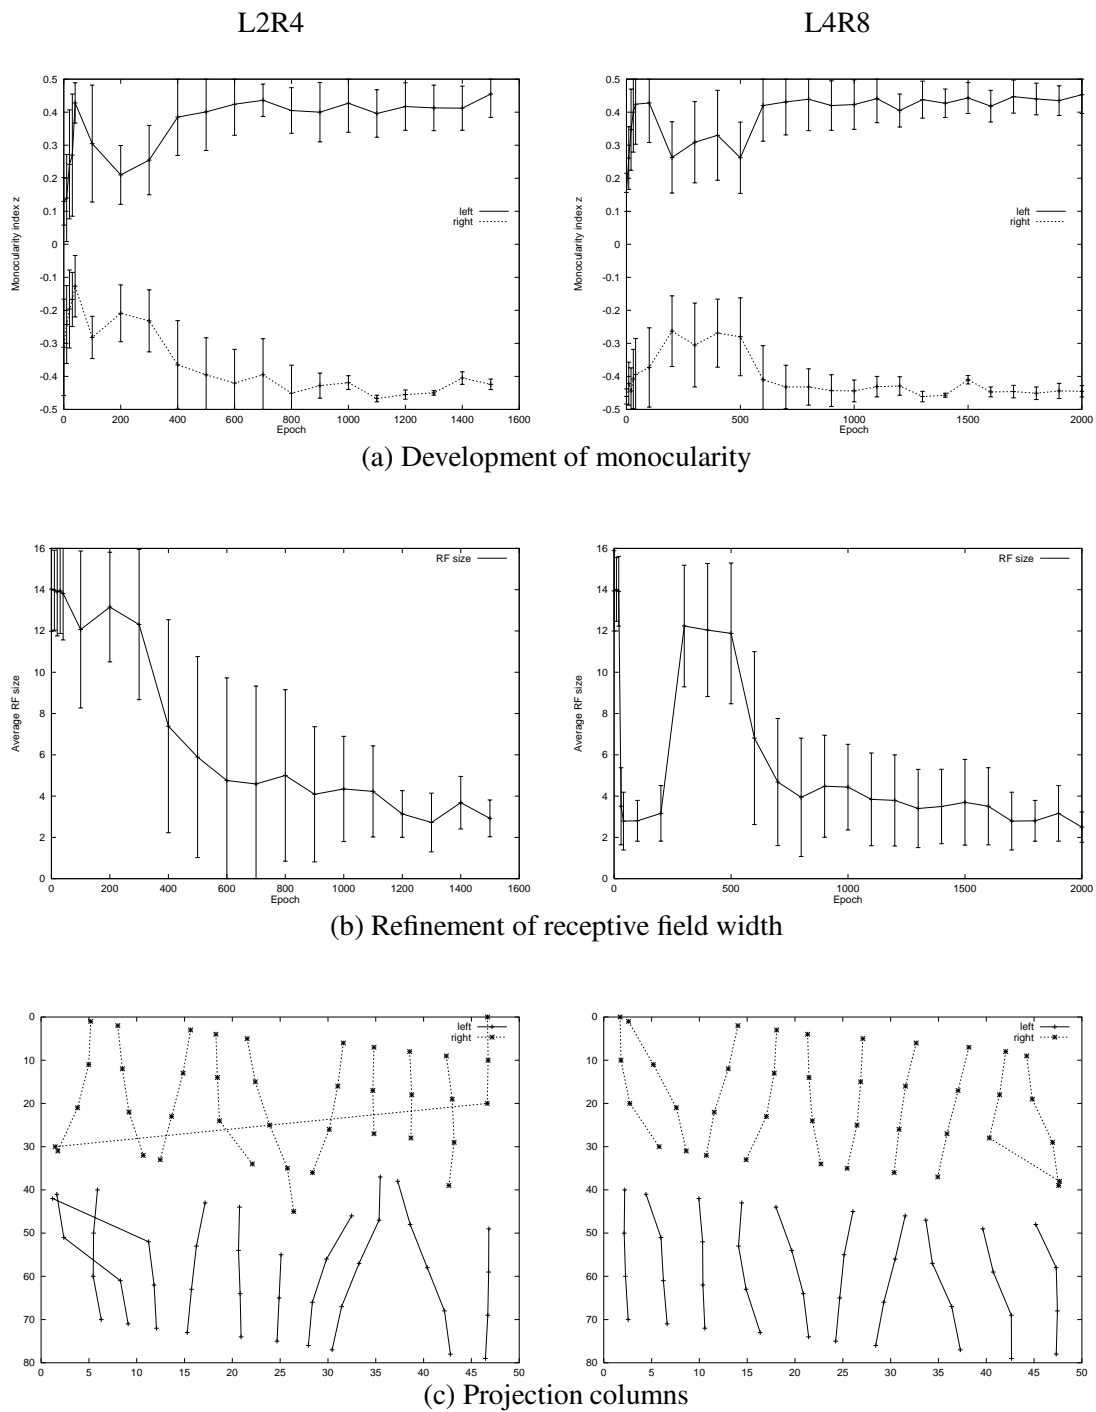


Figure 4.11: Comparison of the L2R4 and L4R8 experiments with respect to the development of ocularity and receptive field width, and the final projection columns. Results for L2R4 are shown on the left and L4R8 on the right.

4.6 The importance of normalisation

As mentioned in Chapter 3, most previous models of ocular dominance require postsynaptic normalisation of weights to allow ocular dominance to develop. However, the results from Figure 4.6 and Figure 4.7 show that ocular dominance develops within this model either with or without any postsynaptic normalisation. These experiments only investigated network development with ocularity bias in the weights. To address the issue of normalisation more thoroughly, a set of experiments were therefore run to examine the effects of all possible combinations of normalisation techniques.

Normalisation can be applied either divisively or subtractively to both the presynaptic and postsynaptic units. In addition, there is the option to ignore normalisation of either pre- or postsynaptic units. This gives us nine combinations of normalisation techniques. The same network with weights initially biased for ocularity and topography (using the same initial weights as shown in Figure 4.10) was run nine times using each combination of normalisation technique. The results from these nine experiments are summarised in Figures 4.12, 4.13, 4.14 and 4.15.

From these figures, several points can be made concerning the role of the different normalisation techniques upon network development.

4.6.1 The effect of normalisation upon topography

First of all, for the normal pattern of topography to develop, the presynaptic normalisation must be divisive. As long as the presynaptic normalisation is divisive, the form of the postsynaptic normalisation is mostly redundant, as shown by Figures 4.12 and 4.14. Although visually the weight matrices look fairly similar for the three experiments with divisive presynaptic normalisation, there are two subtle differences. First, the projection columns for the left eye are marginally more organised when there is no postsynaptic normalisation than for the other two experiments. Second, in the absence of any postsynaptic normalisation, there is no constraint to ensure that all postsynaptic units should receive the same amount of synaptic strength. Figure 4.16 shows the sum of weights to each postsynaptic unit, which indicates that different postsynaptic units receive different amounts of weight strength. There is a systematic trend for postsynaptic units at the extremes of the network (those in rows one and eight of the LGN corresponding to units 0–9 and 70–79) to receive more weight strength than those units at the centre of the LGN (those in rows four and five, corresponding to units 30–39 and 40–49.)

The retinotopic mapping develops best if the presynaptic normalisation is divisive. When the presynaptic normalisation is subtractive, postsynaptic units still refine their receptive fields, but they fail to develop receptive fields similar to their neighbours (as shown in the middle column of plots of Figures 4.12 and 4.14). The subtractive presynaptic normalisation pushes weights to extreme values so that most weights go to zero, with all of the weight strength for the presynaptic unit allocated into one weight. This is in comparison to the divisive presynaptic normalisation, which can distribute its total weight strength amongst many postsynaptic units, as shown in Figure 4.17.

Finally, in the absence of presynaptic normalisation, most of the weight strength for each postsynaptic unit goes either to one or several neighbouring presynaptic units. This is due to the growth mechanism which encourages neighbouring postsynaptic units to develop similar receptive field properties. As shown in Figure 4.18(c), most of the presynaptic units have zero synaptic strength

and therefore the presynaptic units have disconnected from the postsynaptic sheet. (Any activity in these presynaptic units is therefore not propagated to the postsynaptic sheet.) If the growth rule is removed from the model, all presynaptic units can remain connected to the postsynaptic units in the absence of presynaptic normalisation (Figure 4.18(d)), although this prevents the normal development of topography and ocular dominance (Figures 4.18(a) and 4.18(b)). The size of the receptive field for the experiments without presynaptic normalisation is determined by the postsynaptic normalisation: divisive postsynaptic normalisation produces broader receptive fields than the sharper receptive fields produced by subtractive postsynaptic normalisation (as shown in Figure 4.14).

The experiment with neither presynaptic nor postsynaptic normalisation is included here for completeness, although it is difficult to compare with the other experiments. In this experiment, weights increased without bounds causing most weights to reach very large values. Some form of normalisation is therefore necessary to keep the total weight strength in the network within bounds.

4.6.2 The effect of normalisation upon ocular dominance

The form of normalisation also affects development of ocular dominance within the network. The results are summarised in Figure 4.13. The best results are achieved when the presynaptic normalisation is divisive, regardless of the form of postsynaptic normalisation. In these three cases, the network segregates perfectly into the two eye-specific layers.

When the presynaptic normalisation is subtractive, the results are similar to using divisive presynaptic normalisation, although some postsynaptic units failed to become totally monocular. The biggest difference was found for the experiment using subtractive presynaptic normalisation and no postsynaptic normalisation. In the absence of postsynaptic normalisation, there is no guarantee that every postsynaptic unit will receive input. Since the receptive fields produced by subtractive normalisation are much sharper than divisive normalisation, some postsynaptic units fail to receive any input. This produces a large number of “dead” postsynaptic units in the final network (28 out of 80 in the network shown in Figure 4.13).

In the absence of presynaptic normalisation, since all of the receptive fields have developed to receive input from the same part of the presynaptic sheet, all of the postsynaptic units respond to the same eye (in this case the right eye), with no units responding to the other eye. For these simulations, the initial weight strength for the right eye was stronger than for the left eye (due to the contralateral ocularity bias) which caused the network to develop a preference for the right eye. Changing the initial weight bias would affect the eye that dominates network development, but would not change the pattern of ocular dominance produced.

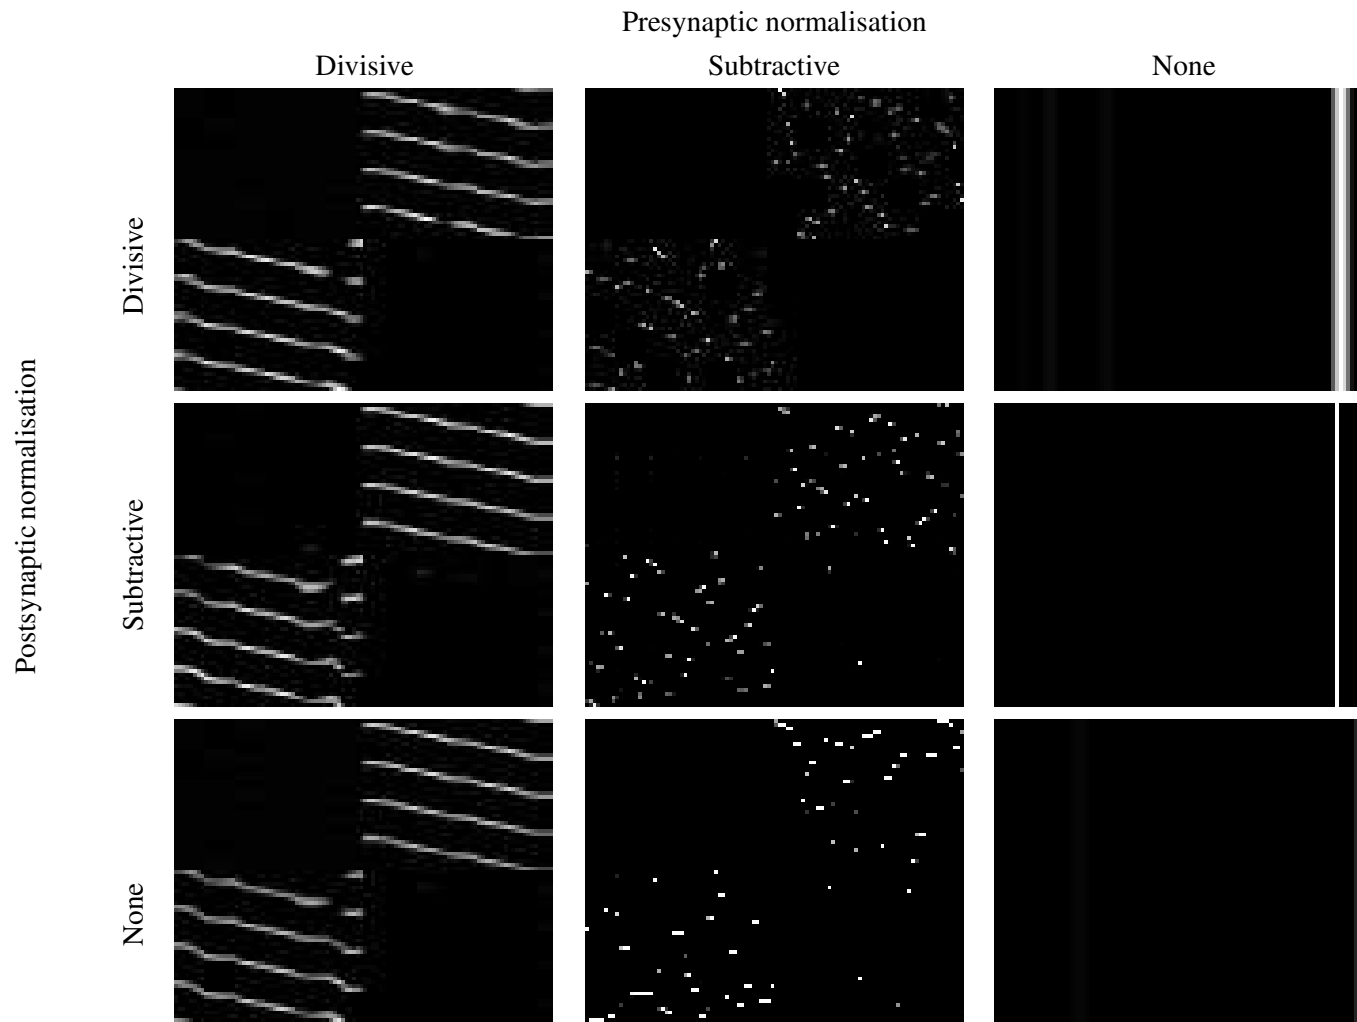


Figure 4.12: The effect of different normalisation methods – 1: weight matrices.

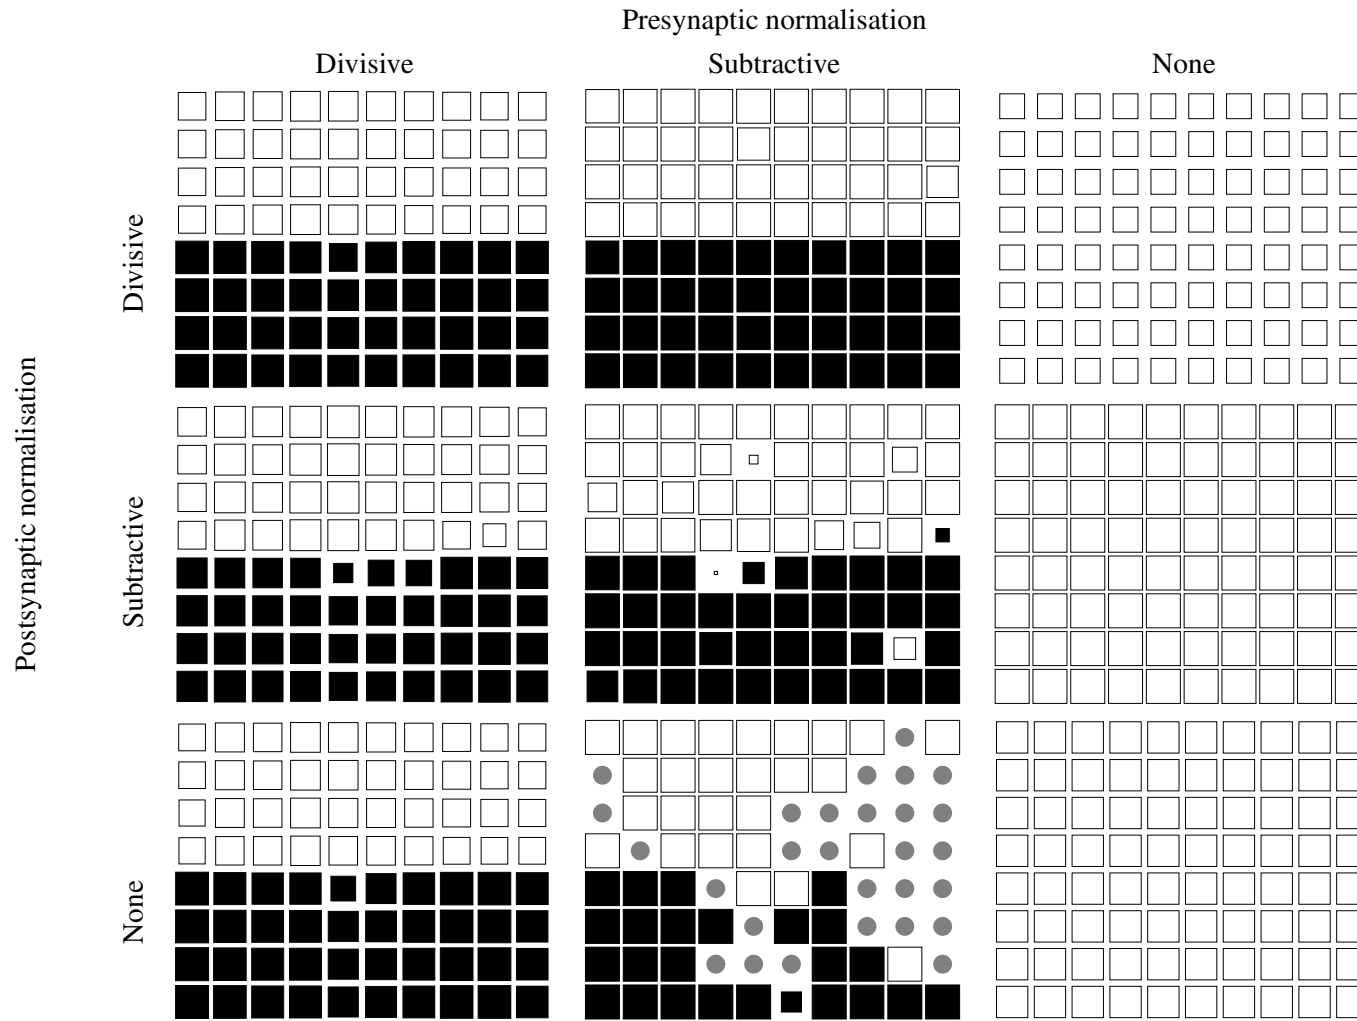


Figure 4.13: The effect of different normalisation methods – 2: ocular dominance plots.

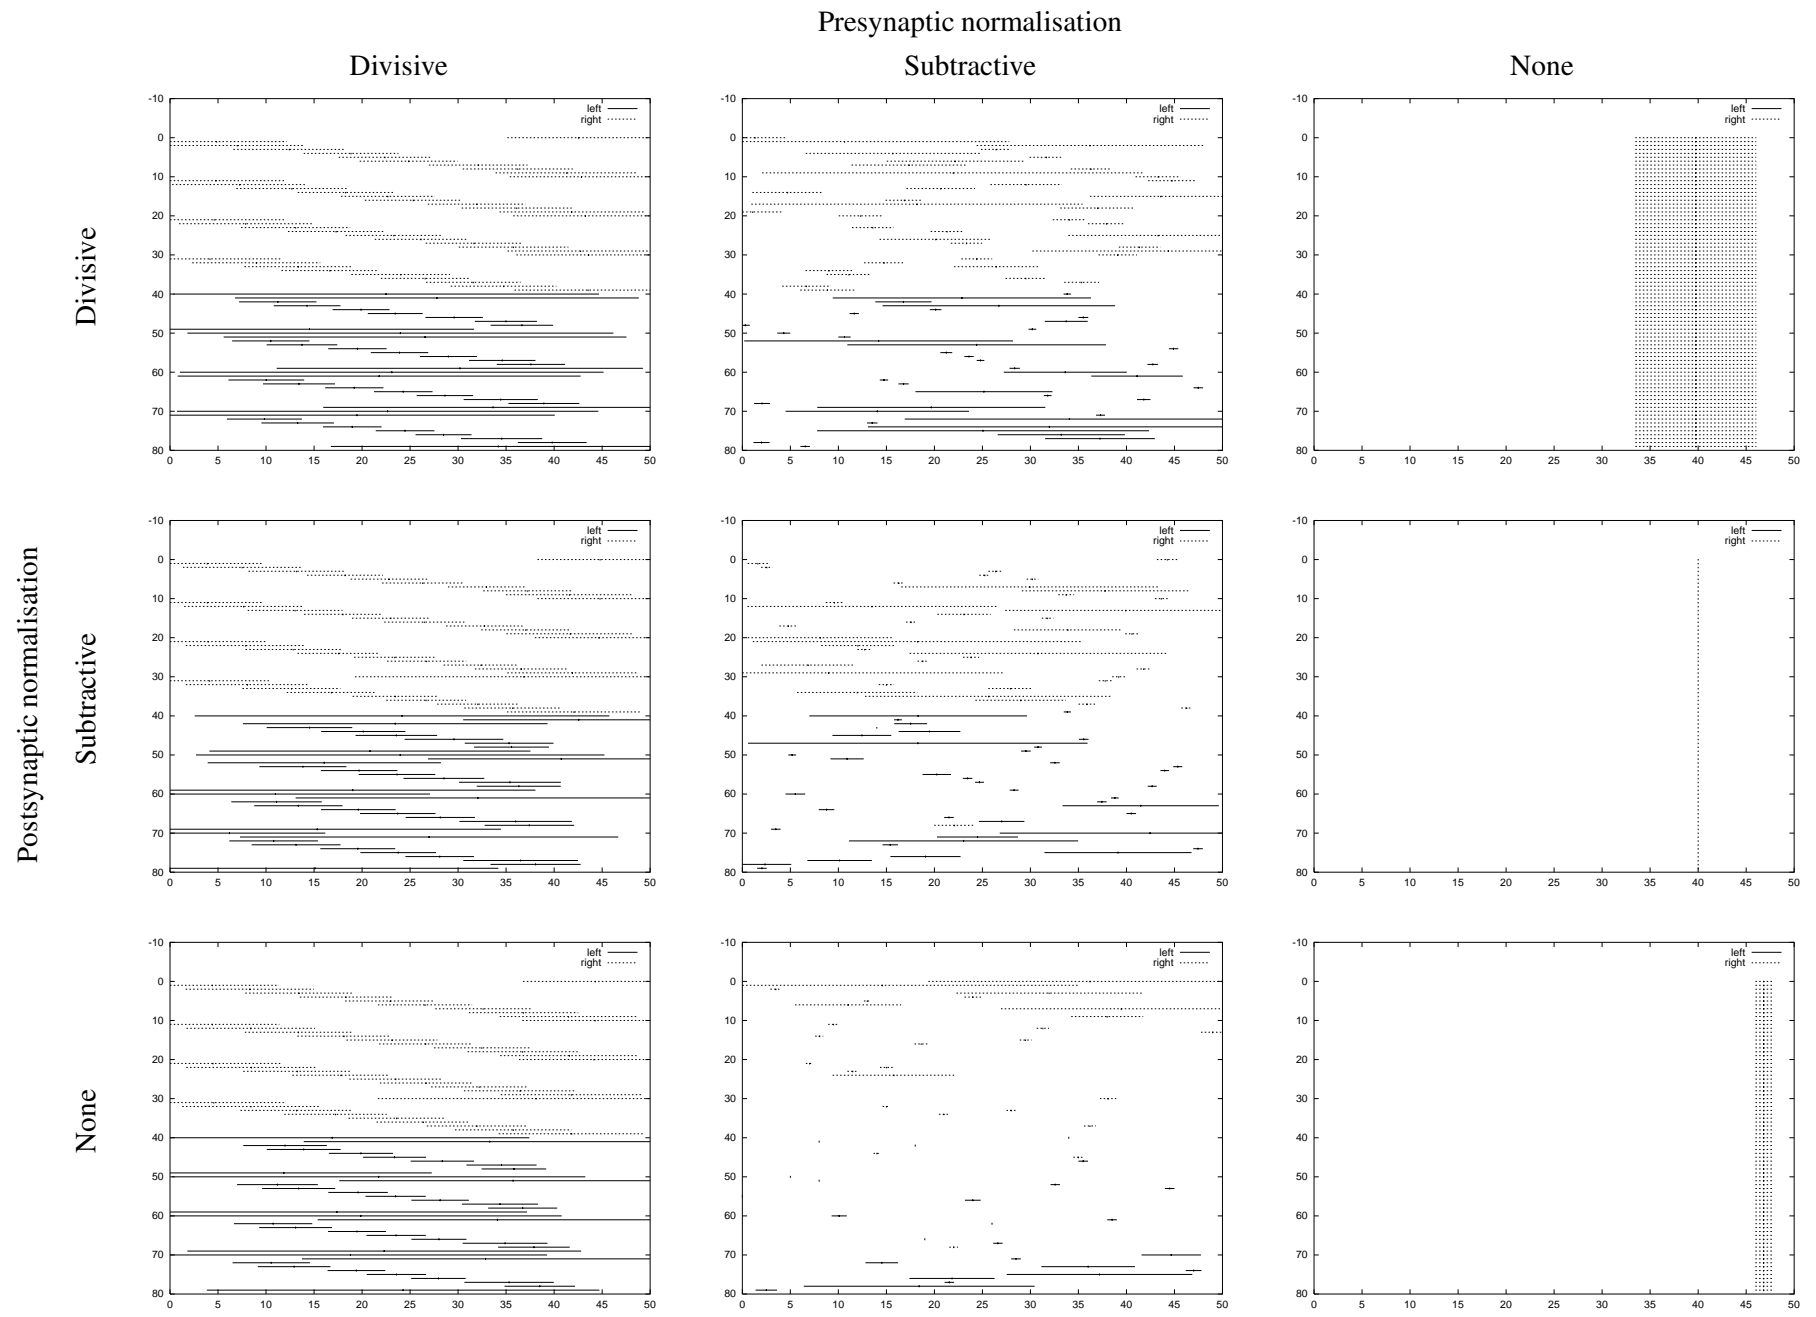


Figure 4.14: The effect of different normalisation methods – 3: topographic plots.

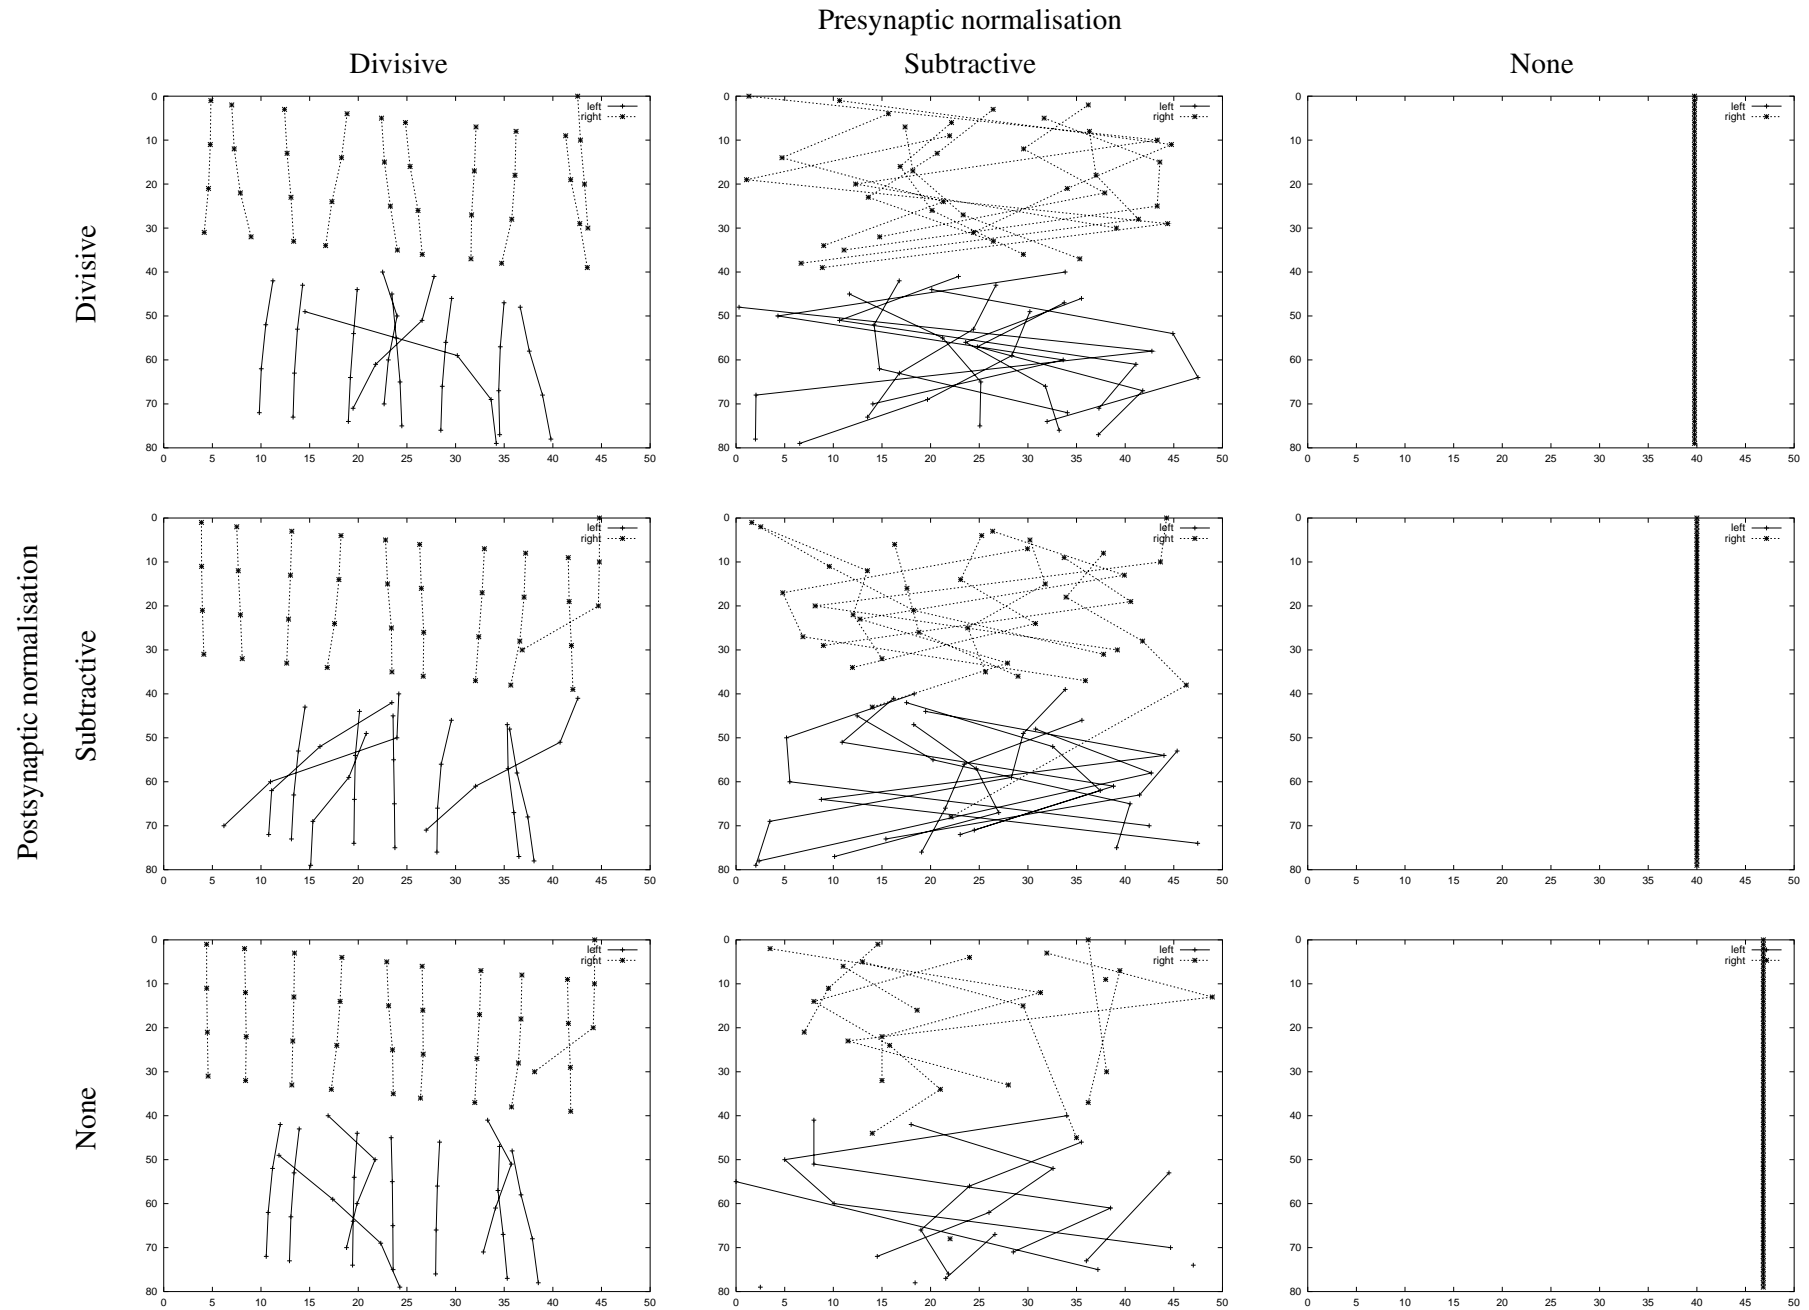


Figure 4.15: The effect of different normalisation methods – 4: projection columns.




Figure 4.16: Sum of weights to each postsynaptic unit in a network with divisive presynaptic normalisation and no postsynaptic normalisation. For each postsynaptic unit j , the sum $\sum_i^{N_{\text{pre}}} w_{ij}$ is shown.




Figure 4.17: Projective field for one presynaptic unit from the right eye using either divisive (solid line) or subtractive (dashed line) presynaptic normalisation. These weights were taken from presynaptic unit 75 in the experiments with no postsynaptic normalisation.

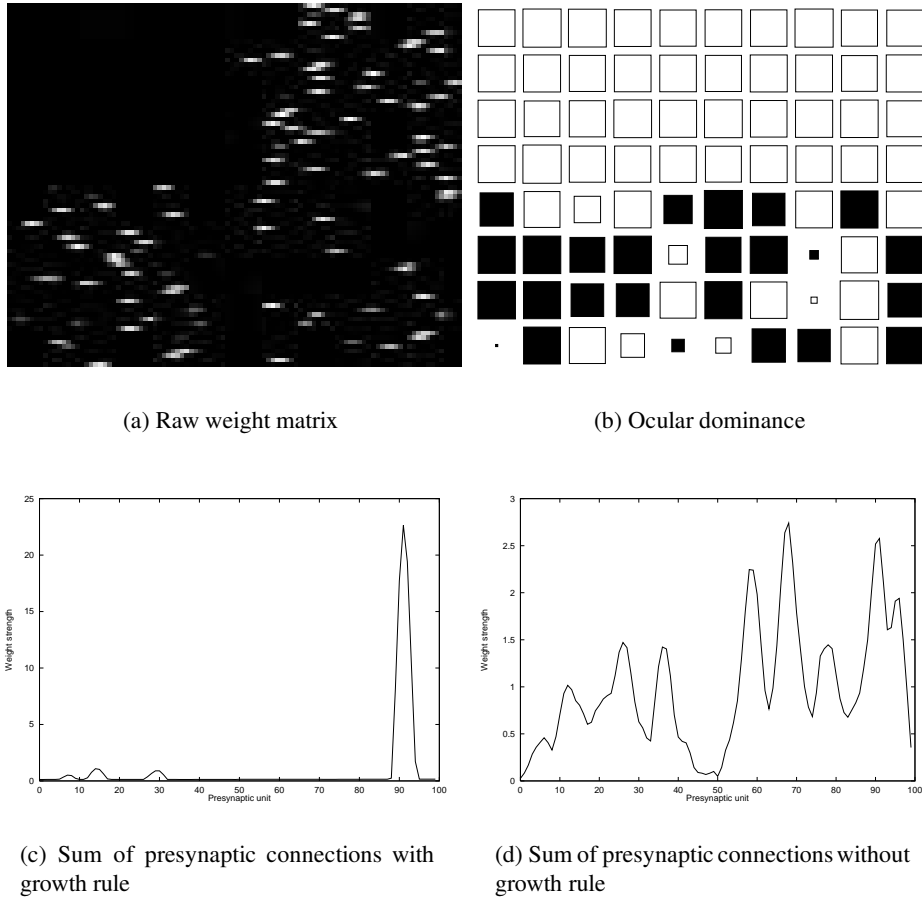


Figure 4.18: The effect of the growth rule upon development in the absence of presynaptic normalisation. (a): Final weight matrix for network with divisive postsynaptic normalisation but no presynaptic normalisation or growth rule. (b): Ocular dominance plot for weights in (a). (c,d): The sum of weights $\sum_j^{N_{\text{post}}} w_{ij}$ for each presynaptic unit. Figure (c) corresponds to the experiment with no presynaptic normalisation and divisive postsynaptic normalisation presented in Figure 4.12. Figure (d) corresponds to the weight matrix shown here in (a).

4.6.3 The importance of the ordering of the normalisation techniques

The previous section has shown that presynaptic normalisation plays a crucial role in development, unlike postsynaptic normalisation. In these experiments, the presynaptic normalisation was always applied before the postsynaptic normalisation. Postsynaptic normalisation may be redundant simply because it is always performed after the presynaptic normalisation. To test for this possibility, a set of experiments were run varying the probability of whether presynaptic or postsynaptic normalisation was applied first.

A new parameter, p_n , was therefore introduced. This controlled the probability of postsynaptic normalisation occurring before presynaptic normalisation. For all of the experiments presented so far in this chapter, this parameter was implicitly set to 0.0 so that divisive presynaptic normalisation was applied first and then subtractive postsynaptic normalisation afterwards. Using the same set of initial conditions as for the previous normalisation experiments, network development was monitored using a range of values for p_n . The results of these experiments are shown in Figure 4.19.

For all values of p_n , the network developed the usual two eye-specific layers, although high values of p_n produced a small number of binocular units and monocular units in the wrong layer. The strongest effect of the p_n parameter was that it affected receptive field size: high values of p_n produced sharper and narrower postsynaptic receptive fields. This sharpening effect of the receptive fields is due to the subtractive normalisation which is sometimes applied first when $p_n > 0.0$. As a consequence of the sharper receptive fields, the topography within each row of the LGN is not so smooth as the value of p_n increases.

4.6.4 Satisfying both forms of normalisation

The previous section has showed that if the subtractive postsynaptic normalisation is applied first, the resulting topographic map is not very smooth, showing discontinuities in each row of the LGN. One possibility for this is simply that if the postsynaptic normalisation is applied first, the presynaptic normalisation can no longer be satisfied afterwards. To quantify how well the two forms of normalisation can be simultaneously satisfied, two error measures were produced:

$$\begin{aligned} \text{presynaptic error} \quad E_{\text{pre}} &= \left[\frac{1}{N_{\text{pre}}} \sum_{i=1}^{N_{\text{pre}}} \left(T_{\text{pre}} - \sum_{j=1}^{N_{\text{post}}} w_{ij} \right)^2 \right]^{\frac{1}{2}} \\ \text{postsynaptic error} \quad E_{\text{post}} &= \left[\frac{1}{N_{\text{post}}} \sum_{j=1}^{N_{\text{post}}} \left(T_{\text{post}} - \sum_{i=1}^{N_{\text{pre}}} w_{ij} \right)^2 \right]^{\frac{1}{2}} \end{aligned}$$

These measures give the root mean squared difference between the target sum and the actual sum of weights. Values smaller than $1\text{e-}2$ for E_{pre} (or smaller than $1.25\text{e-}2$ for E_{post}) indicate that the error is less than 1% of the desired target value.




Table 4.2 shows the error values for weights at different times during development for different values of p_n . From this table it can be seen that the difference between the actual normalised values and the target values are quite small for the normalisation applied first. (By definition, the

normalisation applied second will have a very low error.) The errors also tend to decrease during development, presumably because there are less weight changes made in the network. From these error measures, it is reasonable to conclude that both forms of normalisation can be approximately satisfied at the same time, regardless of whether pre- or postsynaptic normalisation is applied first. This result is in agreement with observations from other models using both pre- and postsynaptic normalisation (Miller & MacKay, 1992; Goodhill, 1992).



Epoch	Normalisation errors							
	$p_n = 0.00$		$p_n = 0.05$		$p_n = 0.20$		$p_n = 1.00$	
	E_{pre}	E_{post}	E_{pre}	E_{post}	E_{pre}	E_{post}	E_{pre}	E_{post}
1	1.09e-2	7.39e-6	1.09e-2	7.39e-6	1.09e-2	7.39e-6	1.02e-6	5.70e-1
350	6.03e-4	8.87e-6	1.61e-6	3.31e-3	1.45e-6	5.45e-3	6.00e-7	9.36e-3
700	6.76e-5	7.85e-6	9.91e-4	6.48e-6	1.74e-3	7.34e-6	4.80e-7	2.75e-2

Table 4.2: Normalisation error measures E_{pre} and E_{post} for different values of p_n at different epochs during development. See Figure 4.19 for the corresponding topography and ocular dominance plots. Error values smaller than $1e-6$ are the errors remaining after the second normalisation has been applied. These errors ideally should be 0.0, but are slightly non-zero due to rounding errors.




Topography plot



Weight matrix



Ocular dominance plot



(a) $p_n = 0.05$

(b) $p_n = 0.2$

(c) $p_n = 1.0$

Figure 4.19: Effect of varying the order of presynaptic and postsynaptic normalisation. The probabilities of postsynaptic normalisation occurring first are: (a) $p_n = 0.05$, (b) $p_n = 0.20$, (c) $p_n = 1.00$.

4.6.5 Capping subtractive presynaptic normalisation

The previous experiments have shown the importance of divisive presynaptic normalisation for the development of topography in the LGN. Networks using subtractive presynaptic normalisation are unable to replicate the topographic projections due to the way that the normalisation pushes individual weights to extreme values. Since there is no maximum weight value imposed on weights in the network, subtractive normalisation forces all of the synaptic weight strength into one of the elements of the weight vector, with all other elements pushed to the minimum weight value of zero, as shown in Figure 4.17. If, however, there is a maximum value imposed on each weight (which will be called w_{\max}), more than one weight for a unit is forced to take a non-zero value. (This assumes $w_{\max} < T_{\text{pre}}$ for presynaptic normalisation and $w_{\max} < T_{\text{post}}$ for postsynaptic normalisation. The case when w_{\max} is equal or greater than the target value for normalisation (T_{pre} or T_{post}) is uninteresting.)

To test the effect of imposing upper limits on individual weights, an extra weight-capping mechanism was introduced to replace the existing weight bound mechanism mentioned in Section 4.2.1. This mechanism constrains each weight to lie in the range $[0, w_{\max}]$. This capping was performed after each of the mechanisms for changing the weights had been applied:

$$w'_{ij} = \begin{cases} 0 & \text{if } w_{ij} < 0 \\ w_{\max} & \text{if } w_{ij} > w_{\max} \\ w_{ij} & \text{otherwise} \end{cases}$$

The effect of imposing a maximum weight value onto individual weights for networks using subtractive presynaptic normalisation (without any postsynaptic normalisation) is shown in Figure 4.21. As can be seen from this Figure (especially when $w_{\max} = 0.2$), the imposition of a maximum weight value on subtractive normalisation does improve the nature of topography within the LGN, although there are several points to note:

1. The value of w_{\max} must be carefully chosen. The larger the value of w_{\max} , the fewer postsynaptic units in different LGN rows are contacted by each presynaptic unit. This is clearly shown in the projective field plots shown in Figure 4.21. Since the target value of presynaptic normalisation is 1.0, values of w_{\max} in the range $[0.2, 0.3]$ produce the best results since they allow presynaptic units to project to four rows of the LGN.
2. The subtractive normalisation must be applied gradually, using the parameter r_s which controls the rate of enforcement of subtractive normalisation. The experiments shown in Figure 4.21 were created with r_s set to 0.05; higher values of r_s causes the normalisation to be stricter, which produces results similar to the earlier results using just subtractive presynaptic normalisation (Figure 4.12) where r_s was set to 1.0. The effect of r_s on enforcing the normalisation constraint can be seen in Figure 4.20: the weights are still bounded, although the mean sum of weights from each presynaptic unit (1.560) is considerably above the desired target value of $T_{\text{pre}} = 1.0$.
3. The capping of weight values does not interfere with subtractive normalisation pushing weights to extreme values: at the end of the simulations presented in Figure 4.21, most

weights have strength 0 or w_{\max} . In contrast, simulations using divisive presynaptic normalisation produce weights distributed over a wide range of values.










Figure 4.20: Sum of weights from each presynaptic unit under relaxed normalisation conditions. This figure shows the final sum of weights from each presynaptic unit for the network shown in Figure 4.21(a). The desired target value for the sum of weights from each units was $T_{\text{pre}} = 1.0$ with $r_s = 0.05$. (Average sum for all presynaptic units = 1.560; standard deviation = 0.532).




Projective field



Weight matrix



Ocular dominance plot



(a) Maximum weight = 0.2

(b) Maximum weight = 0.5

(c) Maximum weight = 1.0

Figure 4.21: The effect of imposing maximum values on individual weights upon development with subtractive presynaptic normalisation (and no postsynaptic normalisation). For these experiments, r_s was set to 0.05. The projective field plots show the final weight vector in each simulation for presynaptic unit 75 (located in the centre of the right eye).

4.7 Summary

This chapter has introduced the Keesing et al. (1992) model of retinogeniculate development. In the first half of this chapter, we have replicated the results of the Keesing model, showing in particular how the form of the initial weight bias influences network development. This replication stage was necessary since the only published information on the model omitted a number of crucial details of the model, including the form of the inputs and the nature of the initial weights (Keesing et al., 1992). We have been able to replicate these initial results, and also expand on the nature of the topographic map across all rows of the LGN, another aspect of the model that was not described in the original publication.

The second half of the chapter has analysed the role of the normalisation mechanisms in network development. In contrast to the original presentation of the model (Keesing et al., 1992), we have shown here that the form of the postsynaptic normalisation is redundant, as long as the presynaptic normalisation is implemented divisively, rather than subtractively. This is a new result, in contrast to previous models of ocular dominance which have used some form of postsynaptic normalisation (reviewed in Chapter 3) to ensure postsynaptic units become monocular. The only other model that relies on presynaptic normalisation of weights is the model by (von der Malsburg, 1979), although it also used a normalisation of the ocular marker concentrations induced into each postsynaptic unit. Postsynaptic normalisation (of either weights or markers) is very likely to produce monocular units, since any increase in weight strengths for some units from one eye is accompanied by a uniform decrease of all weights. Using subtractive enforcement of postsynaptic normalisation increases the tendency for monocular units due to the way that subtractive normalisation forces individual weights to extreme values. We have also shown that in some cases, the presynaptic normalisation can be implemented subtractively, as long as the normalisation is applied slowly and individual weights are constrained to lie between minimum and maximum values. Postsynaptic normalisation may not be required simply because no positive between-eye correlations are present in this model.

Normalisation techniques are often criticised for being too rigid, forcing the weight vectors to always lie exactly on a particular constraint surface, rather than within a wider region of weight space (Swindale, 1996). In this chapter, we have introduced an extra normalisation parameter, r_s , to address this criticism of normalisation techniques. If subtractive normalisation is gradually applied using the r_s parameter as shown in Section 4.6.5, the weights remain finite, but weight vectors no longer lie directly on the constraint surface (as shown in Figure 4.20), producing better topographic maps than when the vectors are forced to lie on the surface.

Chapter 5

The influence of the spatial and temporal wave properties upon development

5.1 Introduction

In Chapter 4 we introduced the Keesing model of retinogeniculate pathway development, investigating in detail the effect of the initial weight bias upon development and the role of normalisation schemes in the network. In this chapter we use the same model to explore the influence of spatial and temporal wave properties upon the development of the retinogeniculate pathway. The chapter is divided into three main experimental sections. First, we examine the temporal properties of the waves by seeing how the rate of wave generation affects development. Second, we vary the width of waves to investigate the spatial properties of the waves. Third, we look at the correlations required between on- and off-centre retinal units to drive development of polarity segregation within units. The results from the model are then compared with the current biological data about retinal waves.

5.2 Exploring the temporal properties of retinal inputs

One of the useful features of this model is that it allows us to vary the temporal properties of the retinal waves to see how development is affected. The model results will be compared with the biological data concerning temporal information later in Section 5.6 after the model results have been presented. Intuitively, the temporal properties of the waves should affect development, especially the monocular status of units: the more often waves are present in each eye, the more likely that units in both eyes are jointly active. This in turn can cause postsynaptic units to remain binocular since presynaptic units in different eyes will be correlated to some degree. In this section, we investigate the temporal properties of the retinal waves by measuring how the rate of wave generation affects development of ocular dominance and topography in the geniculate. Before describing these experiments, the next section analytically describes how the rate of wave generation determines the probability of joint activity in the two eyes.

5.2.1 The relationship between the rate of wave generation and the probability of activity in both eyes

When the waves are independently generated in each eye, three different states describe the activity across the two eyes: either both eyes are active, one of the two eyes is active, or neither eye is active. To calculate the overall probability of both eyes being active at the same time (O_2), one eye being active on its own (O_1), or both eyes being quiet at the same time (O_0), it is first necessary to determine the average delay between one wave finishing and the next wave starting. This value, which shall be called S , depends on p_w , the rate at which new waves are generated.

After an eye has recovered from its refractory period, the probability that a new wave starts in an eye after 0 time steps is p_w . If the wave does not begin at time step 0, then it will start after time step 1 with probability $(1 - p_w)p_w$. Similarly, the probability of a new wave starting at time step 2 is $(1 - p_w)(1 - p_w)p_w$. In general, the probability of a new wave starting at time step n is $(1 - p_w)^n p_w$. The average time when the eye is silent between waves, S , is therefore given by the infinite sum:

$$\begin{aligned}
 S &= 0 \times p_w \\
 &+ 1 \times (1 - p_w)p_w \\
 &+ 2 \times (1 - p_w)(1 - p_w)p_w \\
 &+ \dots \\
 &+ n \times (1 - p_w)^n p_w \\
 &+ \dots \\
 S &= \sum_{k=0}^{\infty} k(1 - p_w)^k p_w
 \end{aligned} \tag{5.1}$$

If we consider the geometric progression $y = \sum_{k=0}^{\infty} q^k$ and take its derivative with respect to q :

$$\begin{aligned}
 y &= \sum_{k=0}^{\infty} q^k = \frac{1}{1 - q} \\
 \frac{dy}{dq} &= \sum_{k=0}^{\infty} kq^{k-1} = \frac{1}{(1 - q)^2}
 \end{aligned} \tag{5.2}$$

This can be substituted into the definition of S :

$$\begin{aligned}
 S &= \sum_{k=0}^{\infty} k(1 - p_w)^k p_w \\
 &= p_w \sum_{k=0}^{\infty} k(1 - p_w)^{k-1} (1 - p_w) \\
 &= p_w(1 - p_w) \sum_{k=0}^{\infty} k(1 - p_w)^{k-1} \\
 &= p_w(1 - p_w) \left(\frac{1}{1 - (1 - p_w)} \right)^2 && \text{(using 5.2 with } q = 1 - p_w) \\
 S &= \frac{p_w(1 - p_w)}{p_w^2} = \frac{1 - p_w}{p_w}
 \end{aligned} \tag{5.3}$$

Once S is known, the overall probability that an eye is active at any time step, O is given by:

$$O = \frac{T}{S + R + T} \quad (5.4)$$

where R is the refractory period, and T the number of time steps that the wave is present on the retina. Once O is calculated, the probabilities O_2 , O_1 and O_0 are given by:

$$O_2 = O \times O = O^2 \quad (5.5)$$

$$O_1 = O(1 - O) + (1 - O)O = 2O(1 - O) \quad (5.6)$$

$$O_0 = (1 - O) \times (1 - O) = (1 - O)^2 \quad (5.7)$$

This analysis can be generalised to n eyes, all independently generating waves: in this case, the probability that i out of n eyes are simultaneously active, O_i , is given as:

$$O_i = (1 - O)^{n-i} \times O^i \times C_i^n, \text{ where } C_i^n = \frac{n!}{i!(n-i)!} \quad (5.8)$$

Figure 5.1 shows how these probabilities vary with the probability that new waves are generated. The overall probabilities for some specific values of p_w are additionally listed in Table 5.1.




Figure 5.1: Overall probabilities of joint retinal activity as a function of the probability of wave generation. This graph shows the overall probabilities for values of p_w only in the range $[0, 0.2]$. Beyond this range, O_0 continues to decay to zero, with O_2 steadily climbing to a value of 0.96 as p_w tends to 1.0.

5.2.2 The effect of wave generation rate upon ocular dominance and topography

To assess the importance of the temporal properties of the waves upon development, a set of experiments were created using identical conditions to those used before (in Figure 4.9), but using different values for the probability of wave generation in both eyes. As shown in Figure 5.1, there are three main different ranges for the parameter p_w (actual values for p_w are given here in brackets assuming $T = 50, R = 1$):

p_w	O	O_0	O_1	O_2
0.002	0.091	0.826	0.165	0.008
0.003	0.130	0.756	0.227	0.017
0.005	0.200	0.640	0.320	0.040
0.020	0.500	0.250	0.500	0.250
0.200	0.909	0.008	0.165	0.826
0.500	0.962	0.001	0.074	0.925
0.800	0.976	0.001	0.047	0.952

Table 5.1: Theoretical probability of eye activity as a function of p_w . Measurements of the corresponding probabilities from the simulations of eye activity were always within four percent of the predicted values.

1. $O_0 > \max(O_1, O_2)$ ($p_w \leq 0.01$). Here it is most likely that both eyes are quiet.
2. $O_1 > \max(O_0, O_2)$ ($0.01 < p_w < 0.04$). In this range, the most likely situation is that one eye is active while the other eye is quiet.
3. $O_2 > \max(O_0, O_1)$ ($p_w \geq 0.04$). In this last range (the widest of the three ranges), both eyes are likely to be jointly active most of the time.

A set of values for p_w were therefore chosen to cover each of these three ranges. The values of p_w used in the experiments are listed in Table 5.1 along with the theoretical probability that both eyes are jointly active. (For comparison, all of the experiments described in Chapter 4 used values of $p_w = 0.02$ for each eye. This corresponds to a value of $O = 0.5$, meaning that, overall, each eye is active for fifty percent of the simulation.)

The correlations between presynaptic units for two extreme values of p_w are compared in Figure 5.2. Despite the large difference in the value of p_w , the only noticeable difference in the correlations are that the within-eye anticorrelations are slightly stronger for the larger value of p_w . Figure 5.3 shows the effect of varying p_w upon development. From this figure it can be seen that normal development depends upon the value of p_w . First, if p_w is too small, most of the units remain binocular and there is no topographic refinement of receptive fields. On the other hand, if p_w is too big, although most units become strongly monocular, the topography begins to break up, producing holes within the representation of the visual space for each row of postsynaptic units (for example, when $p_w = 0.2$).

5.2.3 Introducing the active-covariance rule

The case when p_w is too small is interesting, since it represents the situation when most of the time both eyes are quiet. This should produce ideal conditions for the development of monocularity because units in the two eyes are not likely to be active at the same time. Instead of the weights becoming monocular though, all of the weights in the network developed to the same value, producing completely binocular units with no retinotopic refinement.

An analysis of the covariance rule revealed the reason behind this lack of refinement of receptive fields. In this model, postsynaptic activity depends linearly on presynaptic activity. Given that most of the time the presynaptic units are quiet, the postsynaptic units are also usually quiet.




Figure 5.2: Correlation between presynaptic units as a function of the rate of wave generation. The figure shows the correlation between presynaptic unit 30 and all other presynaptic units for both $p_w = 0.8$ (solid line) and $p_w = 0.005$ (dotted line). The correlation, $Corr(i, j)$, between two presynaptic units i and j is defined as $Corr(i, j) = Cov(i, j) / \sqrt{Cov(i, i) \times Cov(j, j)}$ where $Cov(i, j) = \langle x_i x_j \rangle - \langle x_i \rangle \langle x_j \rangle$ (the covariance between units i and j) and $\langle \rangle$ denotes averaging over all inputs. Presynaptic units 1–50 come from the left eye, and units 51–100 come from the right eye.

When both pre- and postsynaptic activity are below threshold, case four of the covariance rule (defined in Table 3.2 on page 38) increases the strength of connections between the two units. For networks with low overall probabilities of retinal waves being present, non-specific increase in all of the weights in the network, combined with the averaging effects of the growth rule, may cause all of the weights in the network to converge to the same value. This possibility is confirmed by the results presented in Table 5.3, which shows the amount of weight change specified by each case of the covariance rule. For low values of p_w , case four of the rule is responsible for most of the weight change in the network. As p_w increases, although the amount of weight change caused by case four stays roughly the same, the amount of weight change specified by the other cases increases considerably, so that case four is no longer dominant.

To test the hypothesis that case four of the covariance rule is responsible for averaging out weight values when p_w is small, the experiments were repeated with a slightly modified covariance rule. This rule was identical to the covariance rule except that when both pre- and postsynaptic activities were below their threshold levels (case four), no weight adaptation was performed. This rule is referred to here as the “active-covariance rule” to indicate that weight adaptation takes place only when either (or both) pre- or postsynaptic activity is above threshold. (This rule is similar to the rule presented by (Montague et al., 1991), although it was independently discovered during this thesis (Eglen, 1996).)

The effect of varying p_w upon development in a network using the active-covariance rule is shown in Figure 5.4. For all values of p_w chosen, network development with the active-covariance rule was either the same or better (by visual comparison) than with the covariance rule. Of particular interest, when both eyes are normally inactive ($p_w = 0.002, 0.003, 0.005$), the network develops

the normal patterns of ocular dominance and topography. A common feature of both rules however is that the average receptive field width increases inversely with p_w (see Table 5.2). These results must be interpreted with some caution however, due to the rather high standard deviation of receptive field widths in comparison to their mean values.

p_w	Mean \pm standard deviation	
	Covariance	Active-cov.
0.002	N/A	11.384 \pm 4.718
0.003	N/A	5.793 \pm 5.770
0.005	N/A	3.762 \pm 4.965
0.020	4.305 \pm 4.956	1.591 \pm 2.780
0.200	2.199 \pm 0.869	0.889 \pm 0.046
0.500	2.184 \pm 1.001	0.887 \pm 0.054
0.800	2.109 \pm 0.679	0.896 \pm 0.111

Table 5.2: Mean receptive field width for different values of p_w using either the covariance or active-covariance rule. For the entries listed N/A, all postsynaptic units were binocular, which precluded taking any receptive field measurements.

p_w	Frequency of use / %				Weight change per epoch			
	1	2	3	4	1	2	3	4
0.002	0.495	0.011	15.468	84.027	95.5920	-1.1990	-666.186	5830.838
0.003	0.763	0.015	23.032	76.191	164.7320	-1.4550	-1062.795	5279.894
0.005	1.103	0.018	32.381	66.498	257.7480	-1.7380	-1571.903	4600.636
0.020	1.062	1.732	23.810	73.396	1344.4520	-393.7110	-7218.146	4541.466
0.200	1.324	3.427	23.559	71.690	3108.4650	-819.3380	-13704.282	4365.681
0.500	1.459	3.540	25.339	69.661	3339.5660	-852.4340	-14409.363	4250.804
0.800	1.609	3.453	27.526	67.412	3448.7770	-846.1370	-14552.396	4185.087

(a) Covariance rule

p_w	Frequency of use / %				Weight change per epoch			
	1	2	3	4	1	2	3	4
0.002	0.382	0.124	11.815	87.680	112.7830	-24.6050	-806.365	6017.041
0.003	0.409	0.368	12.171	87.052	242.7040	-86.3470	-1652.527	5920.608
0.005	0.403	0.718	11.412	87.467	432.6690	-183.7100	-2872.521	5948.982
0.020	0.671	2.123	14.894	82.312	1513.0623	-562.6000	-8218.576	5538.130
0.200	1.063	3.689	18.945	76.303	3235.3990	-957.5940	-14278.994	4969.223
0.500	1.137	3.863	19.737	75.264	3473.8570	-999.0490	-15003.074	4875.800
0.800	1.249	3.813	21.323	73.615	3577.4920	-987.0899	-15108.254	4773.361

(b) Active-covariance rule

Table 5.3: The frequency of use and the amount of weight change per epoch for each case of the covariance and active-covariance rule. Cases refer to the cases of the covariance rule as shown in Table 3.2. For the active-covariance rule, although values of weight change for case four are given here, these weight changes were ignored.

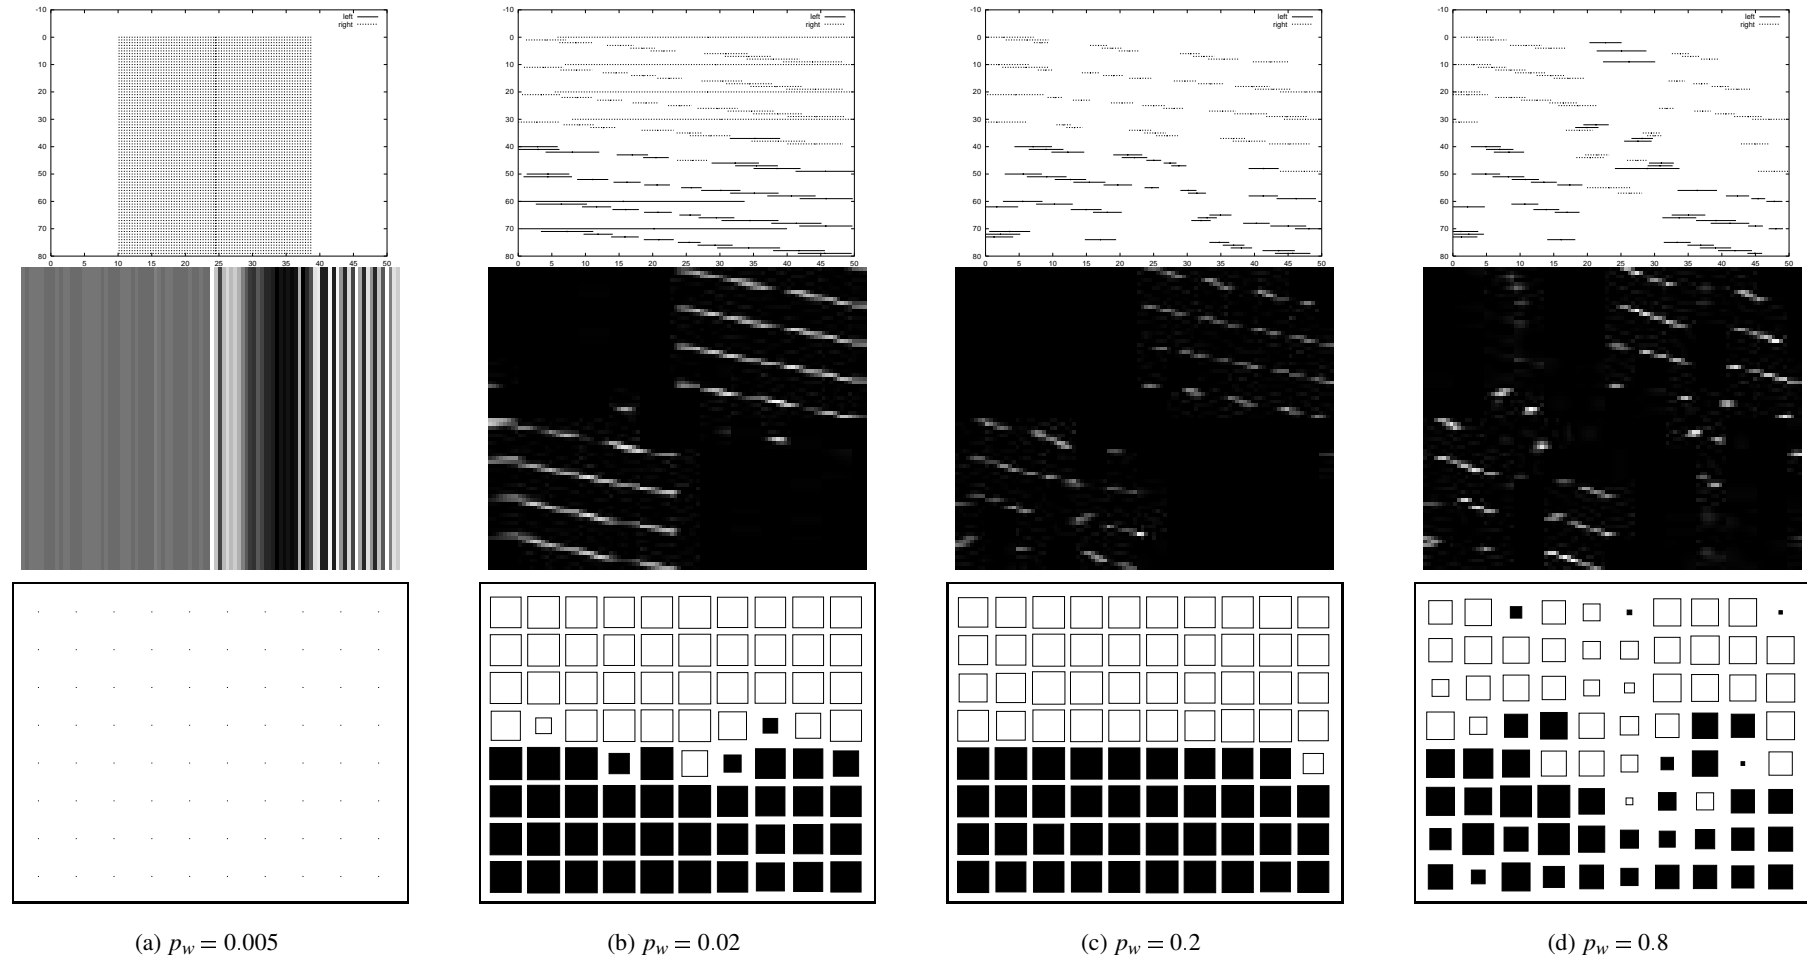


Figure 5.3: Effect of varying p_w upon topography and ocular dominance – 1: development with the covariance rule. For each value of p_w we show (from top to bottom) the final topography plot, weight matrix and ocular dominance plot. When $p_w = 0.005$, all weights adapt to (almost) the same value, producing completely binocular units (hence the empty ocular dominance plot). The plot of the weight matrix for this experiment has exaggerated the small differences between weights, and should not be compared with the other weight matrices.

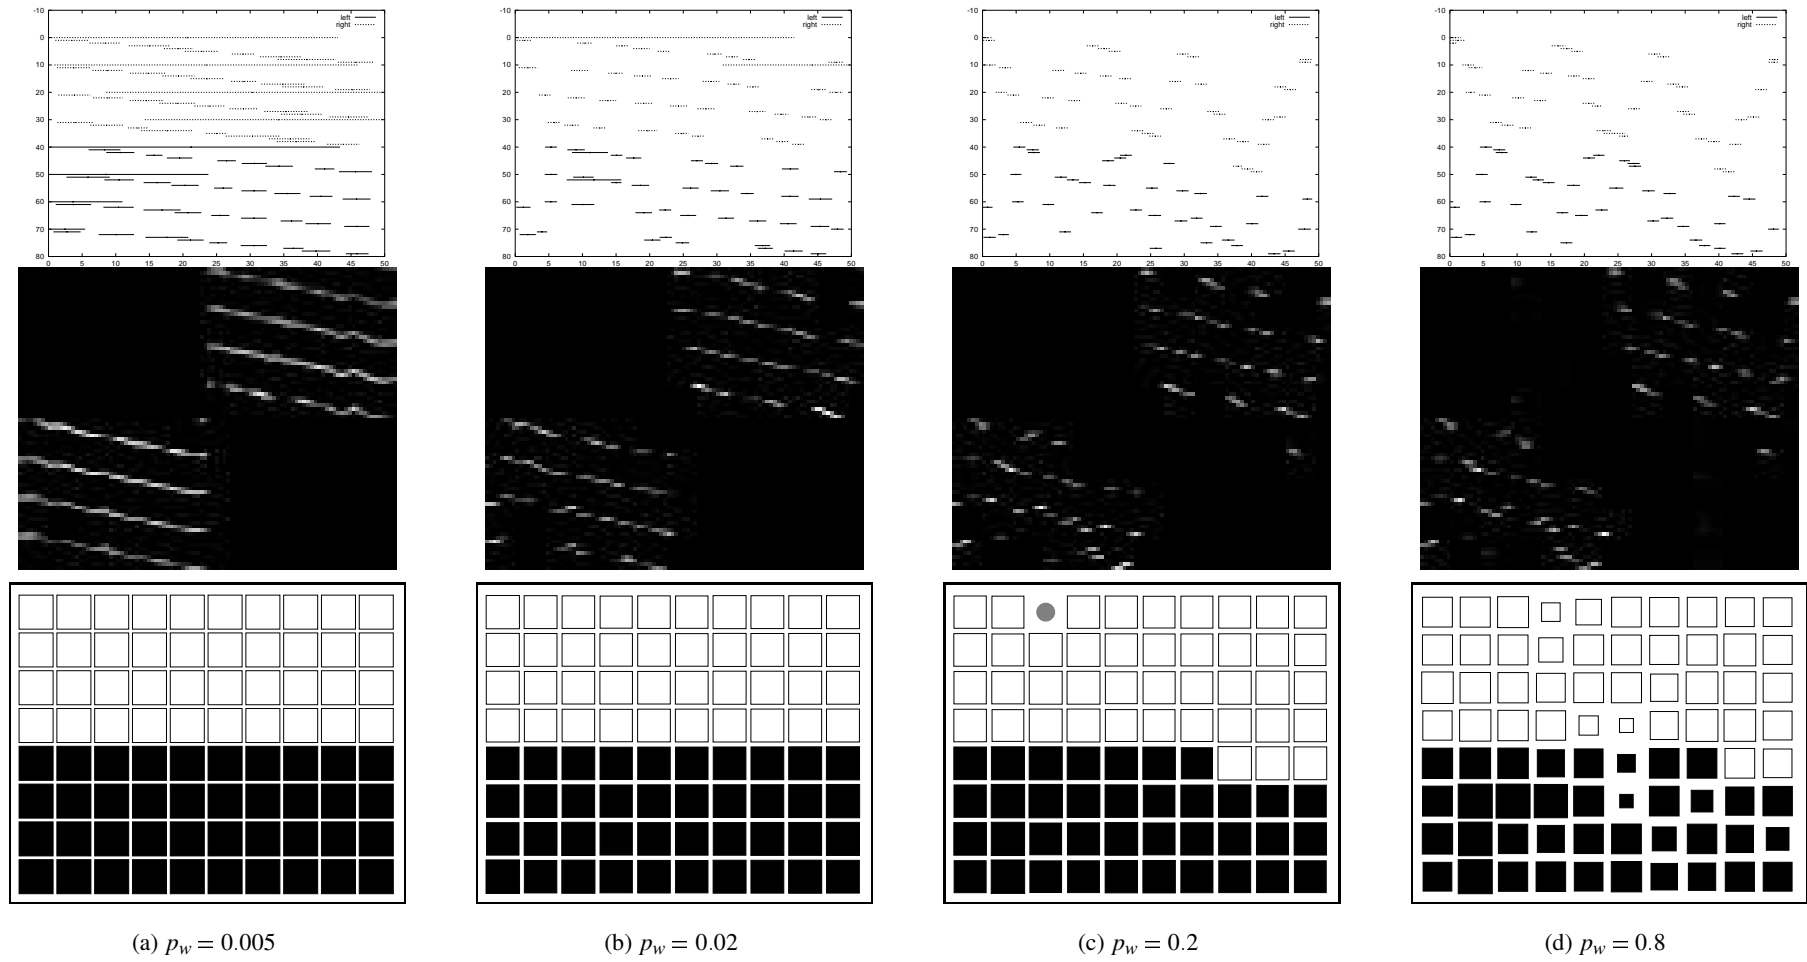


Figure 5.4: Effect of varying p_w upon topography and ocular dominance – 2: development with the active-covariance rule. For each value of p_w we show (from top to bottom) the topography plot, the weight matrix and the ocular dominance plot at the end of development.

5.3 Development of the LGN under conditions of monocular deprivation

The previous section investigated the effect of wave generation rate upon development. In these experiments, the probability of wave generation was the same for both eyes. Although to date there are no measurements comparing the properties of waves in the two eyes, it is reasonable to assume that in normal development the waves in the two eyes have similar spatio-temporal properties. However, experimental manipulations to the visual pathway have reduced the level of activity in one eye, revealing more details about the mechanisms underlying development of the visual system.

A common technique for reducing activity levels from one eye is simply to remove the eye early in development. The effect of enucleation upon the development of the visual cortex is fairly well established (Shatz & Stryker, 1978; Kalil, 1980): in such cases, the non-deprived eye occupies more of the postsynaptic sheet than normal at the expense of the deprived eye. Self-organising models of cortical development can reproduce the data from these experiments due to the competitive nature of the models: the more active presynaptic units tend to connect to more postsynaptic units (von der Malsburg, 1979; Goodhill & Willshaw, 1994).

Monocular deprivation also affects LGN development. Removal of one eye at day E44 (in the cat) allowed the remaining non-deprived eye to innervate the whole of the LGN. The only segregation found within the LGN was a novel lamination pattern consisting of a dorsal magnocellular division and a ventral parvocellular division (Garraghty et al., 1988).

To test whether this model can reproduce the results of monocular deprivation in the LGN, the probability of wave generation was lowered for the left eye in comparison to the right eye. Five different values of p_w (0.005, 0.003, 0.002, 0.001 and 0.000) were chosen for the left eye whilst p_w for the right eye was fixed at 0.020 for all five experiments. Given the low values of p_w and the results from Section 5.2.2, the active-covariance rule was used to modify weights. Initial weights were the same as shown in Figure 4.9, namely the left eye initially innervated the bottom two rows of the LGN and the right eye innervated the bottom four rows.

Figure 5.5 shows the final ocular dominance and topography plots for these experiments. As expected, the lower the value of p_w for the left eye, the fewer postsynaptic units connected to that eye. Some of the units in the bottom half of the LGN that are normally responsive to the left eye instead are innervated by the right eye. The size of layer A increases at the expense of layer A1, as layer A grows downwards into parts of the LGN that usually belong to layer A1. Table 5.4 quantifies the number of postsynaptic units that are responsive either to the left, right or neither eye. (All postsynaptic units developed to become either strongly monocular or dead.) This table shows that as p_w for the deprived eye decreases, the number of postsynaptic units responsive to the non-deprived eye increases.

Given the nature of the active-covariance rule, the basic result that the non-deprived eye occupies more of the LGN was expected. For units in the deprived eye, most of the time the active-covariance rule is specifying weight decrements (especially with case three of the rule), since the presynaptic activity is low but the postsynaptic activity is high due to the input activity from the non-deprived eye. However, there is one slight complication in the model as it stands. For the extreme case when the left eye never produces any waves ($p_w = 0.000$), there are still six units responsive to the left eye. This occurs simply because the presynaptic normalisation forces each

presynaptic unit to keep connections to the postsynaptic sheet, regardless of the activity level of the presynaptic units.

5.3.1 A weaker form of divisive normalisation

For the monocular deprivation simulations, divisive normalisation was used to fix the sum of weights from each presynaptic unit. Considering the divisive normalisation from Section 3.6.1 in more detail, there are two cases to consider when the weights are to be modified so that the sum of weights equals some target value:

- Multiplicative normalisation: when the sum is *less* than the target value, each weight is *multiplied* by some constant greater than one.
- Divisive normalisation: when the sum is *greater* than the target value, each weight is *divided* by some constant greater than one.

Both of these cases are called divisive normalisation, although strictly speaking only the second case divides each weight. Divisive normalisation therefore reduces the amount of weight strength back to some constant, whereas multiplicative normalisation increases the amount of weight strength. Arguments for normalisation methods in computational models are usually based on providing some form of resource limitation to prevent weights from increasing without bounds. However, by using both divisive and multiplicative normalisation, the constraint is stronger: not only does it prevent weights increasing without bounds, but it also prevents weights decreasing without bounds. This constraint is too strong, since most models also employ other techniques for ensuring weights do not decrease without bounds (such as ensuring each weight is non-negative).

A weaker form of normalisation was therefore used in the model by ignoring the multiplicative element of presynaptic normalisation. This effectively changes the normalisation for presynaptic unit i from that given in Section 4.2.1 to:

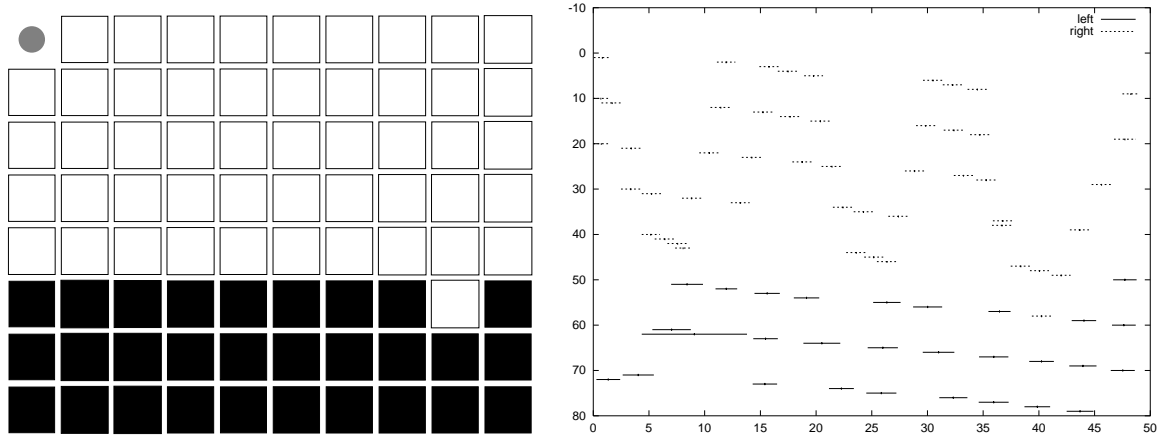
$$\sum_{j=1}^{N_{\text{post}}} w_{ij} \leq T_{\text{pre}}$$

The same set of monocular deprivation experiments were repeated using this new form of divisive normalisation. The results are shown in Figure 5.6, with counts of the number of postsynaptic units responding to each eye shown in the right-hand side of Table 5.4. The overall results are broadly similar to the experiments with both divisive and multiplicative normalisation: the less active the eye, the fewer postsynaptic units receive input solely from that eye. However, for very small values of p_w for the left eye (such as 0.000 or 0.002), no postsynaptic units are responsive to the left eye: all units are either responsive to the non-deprived right eye or are dead. The number of dead postsynaptic units also increased when multiplicative normalisation was ignored. The weaker form of divisive normalisation therefore produces better results than the divisive and multiplicative normalisation for replicating the outcome of monocular deprivation.

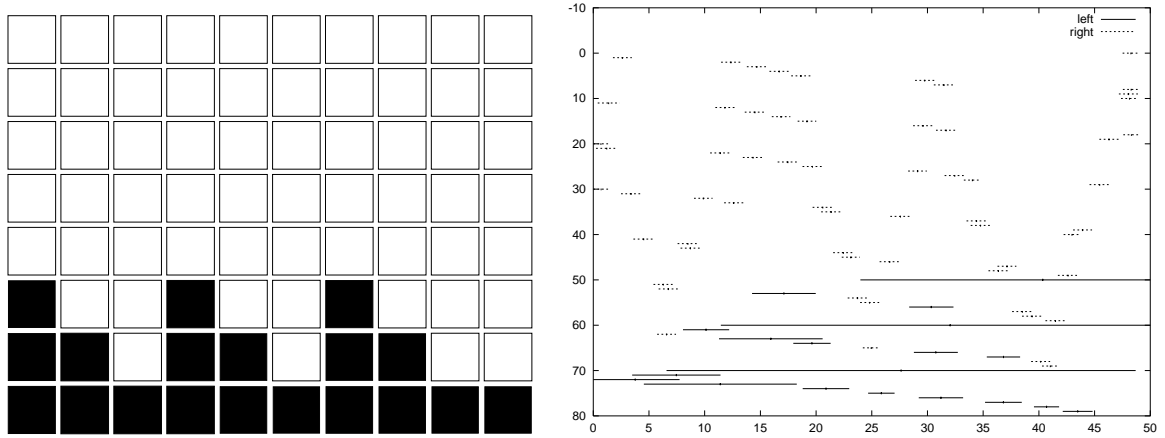
5.3.2 The nature of the retinotopic mapping

Regardless of the form of the divisive normalisation, the nature of the retinotopic mapping for the non-deprived eye is different to the normal mapping. As shown in the topographic plots of

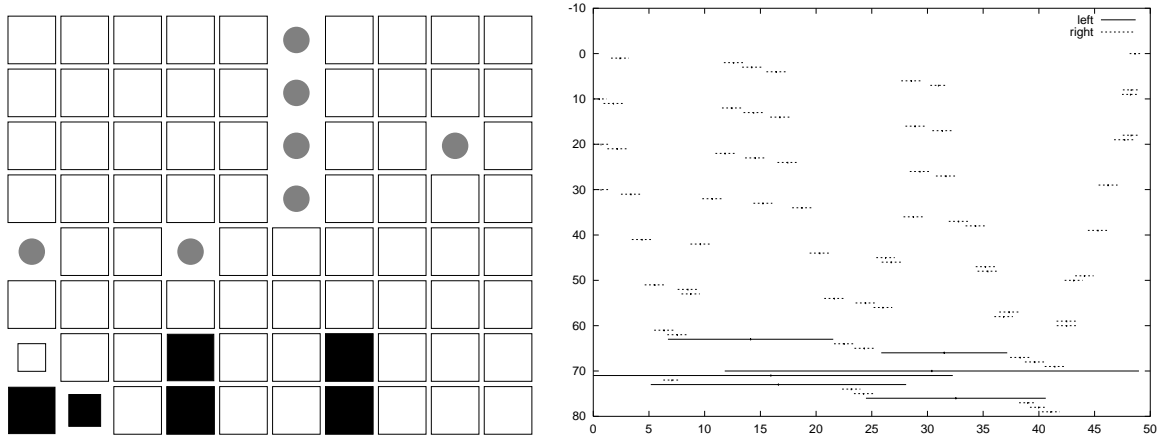
Multiplicative and divisive normalisation



(a) $p_w = 0.005$



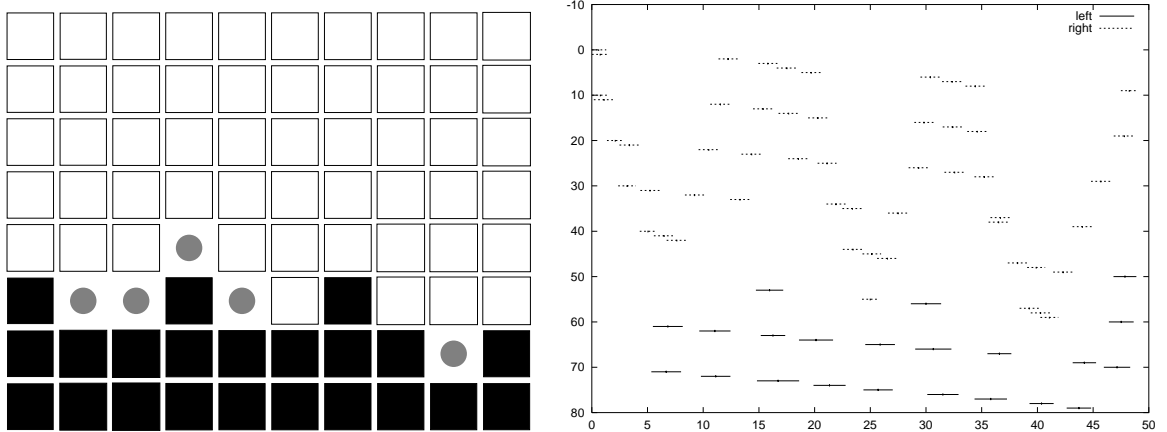
(b) $p_w = 0.002$



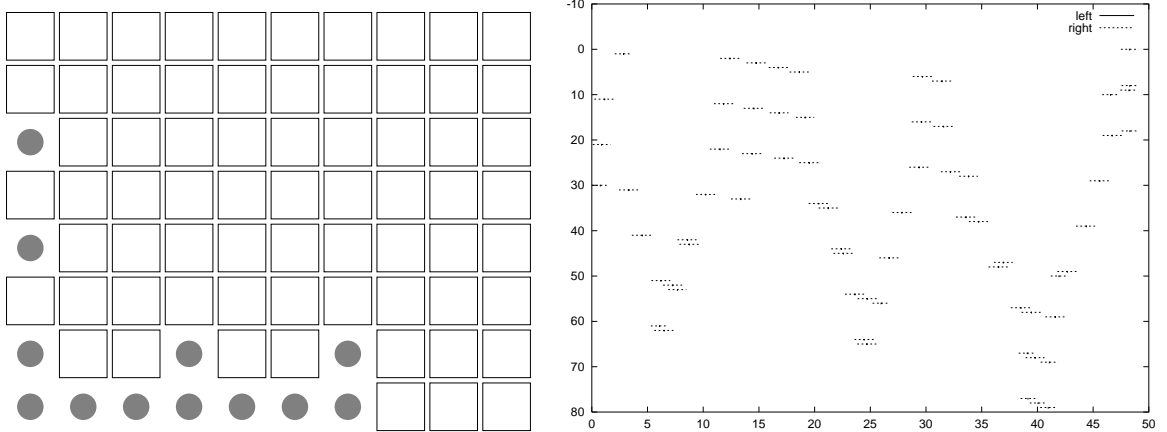
(c) $p_w = 0.000$

Figure 5.5: Network development with reduced activity levels for the left eye – 1: development with multiplicative and divisive presynaptic normalisation.

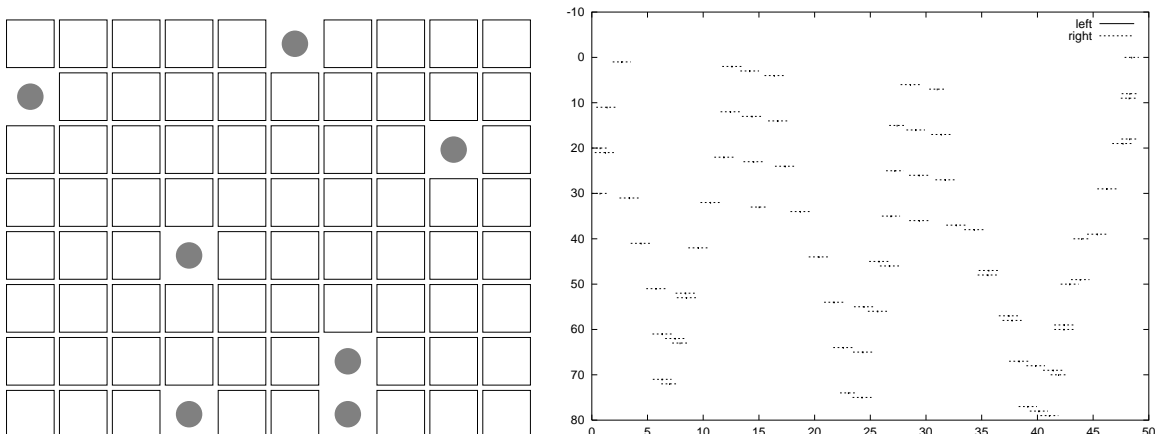
Divisive normalisation



(a) $p_w = 0.005$



(b) $p_w = 0.002$



(c) $p_w = 0.000$

Figure 5.6: Network development with reduced activity levels for the left eye – 2: development with only divisive presynaptic normalisation, ignoring multiplicative normalisation.

Left p_w	(a) D and M			(b) D only		
	Left	Right	Dead	Left	Right	Dead
0.005	29	50	1	22	53	5
0.003	22	58	0	7	64	9
0.002	19	61	0	0	68	12
0.001	15	62	3	0	71	9
0.000	6	67	7	0	73	7

Table 5.4: Number of postsynaptic units responsive to the left, right or neither eye at the end of development under reduced activity conditions for the left eye. The value of p_w for the right eye was fixed at 0.020, and the value of p_w for the left eye varied as given in the table. (a): Results using both divisive and multiplicative presynaptic normalisation. (b): Results using only divisive presynaptic normalisation.

Figures 5.5 and 5.6 (especially when $p_w = 0.000$), each row of LGN units no longer completely covers the visual space — gaps are present within each row. The projection columns for these maps show that units in different rows respond to neighbouring parts, rather than the same parts, of the presynaptic sheet (Figure 5.7). This mapping is produced because the non-deprived eye innervates more rows of the LGN, allowing more freedom for neighbouring vertical units to represent neighbouring, rather than identical, parts of visual space.




Figure 5.7: Final projection columns for a network with no waves generated within the left eye. Projection plots taken from network shown in Figure 5.6(c). One of the projection columns jumps from the far right of the visual field to the far left and then back to the right due to the periodic conditions in the LGN.

5.4 The effect of wave width upon development

Having considered the temporal properties of the waves, we now turn our attention to the spatial properties of retinal waves. The width of the wave controls the degree to which neighbouring presynaptic units within an eye are correlated. To investigate the effect of wave width σ_w upon the development of topography and ocularity, the LGN was reduced to a one-dimensional array of twenty postsynaptic units. This reduction in the postsynaptic sheet from a two-dimensional to a one-dimensional array allowed us to ignore the issue of projection columns and focus on the topography and ocularity of individual postsynaptic units within just one row of the LGN. It also shows that the set of weight modification rules can produce topographic and ocular refinement within networks of different sizes to the usual one-dimensional retina and two-dimensional geniculate.

For this network, apart from the changes in size of the postsynaptic sheet ($X_{\text{post}} = 20, Y_{\text{post}} = 1, N_{\text{post}} = 20$), all of the other parameters were the same as listed in Table 4.1. Initial weights were created with a topographic bias of $b = 20$. Figure 5.9 shows the final maps created using a range of values for σ_w . For values of σ_w smaller than 2.5, units have small receptive fields and are mostly monocular. As the wave width increases, so does the size of the receptive field. This is because the extent of the within-eye correlations widens with σ_w , as shown in a plot of the presynaptic correlations in Figure 5.8.

The dependence of wave width upon σ_w is summarised in Figure 5.10. Increasing σ_w also affects the overall topographic arrangement: when σ_w is small, the neighbouring postsynaptic units are organised into a smooth topographic mapping. As the wave width increases, although neighbouring units still tend to have similar receptive fields, the global topographic mapping is lost.

As shown by the ocularity plots in Figure 5.9, the wave width also affects the final ocularity of each postsynaptic unit. For small wave widths, most postsynaptic units become monocular. For wider waves, units tend to remain binocular (although the results when $\sigma_w = 4$ in Figure 5.9 are an exception). Although the waves are still uncorrelated in their activity patterns, the wider the waves, the more chance there is that waves will be present in both eyes at the same time. Given that the covariance rule will tend to connect active presynaptic units to the same postsynaptic unit, postsynaptic units will tend to receive inputs from both eyes.




Figure 5.8: Correlations between presynaptic units as a function of wave width. The figure shows the correlation between presynaptic unit 30 and all of the other presynaptic units for various values of σ_w . As the wave width increases, the correlation function between neighbouring presynaptic units of the same eye gets broader, and distant presynaptic units get more anticorrelated. In contrast, the between-eye correlations remain uncorrelated regardless of the value of p_w . (The correlation between units is defined in Figure 5.2.)

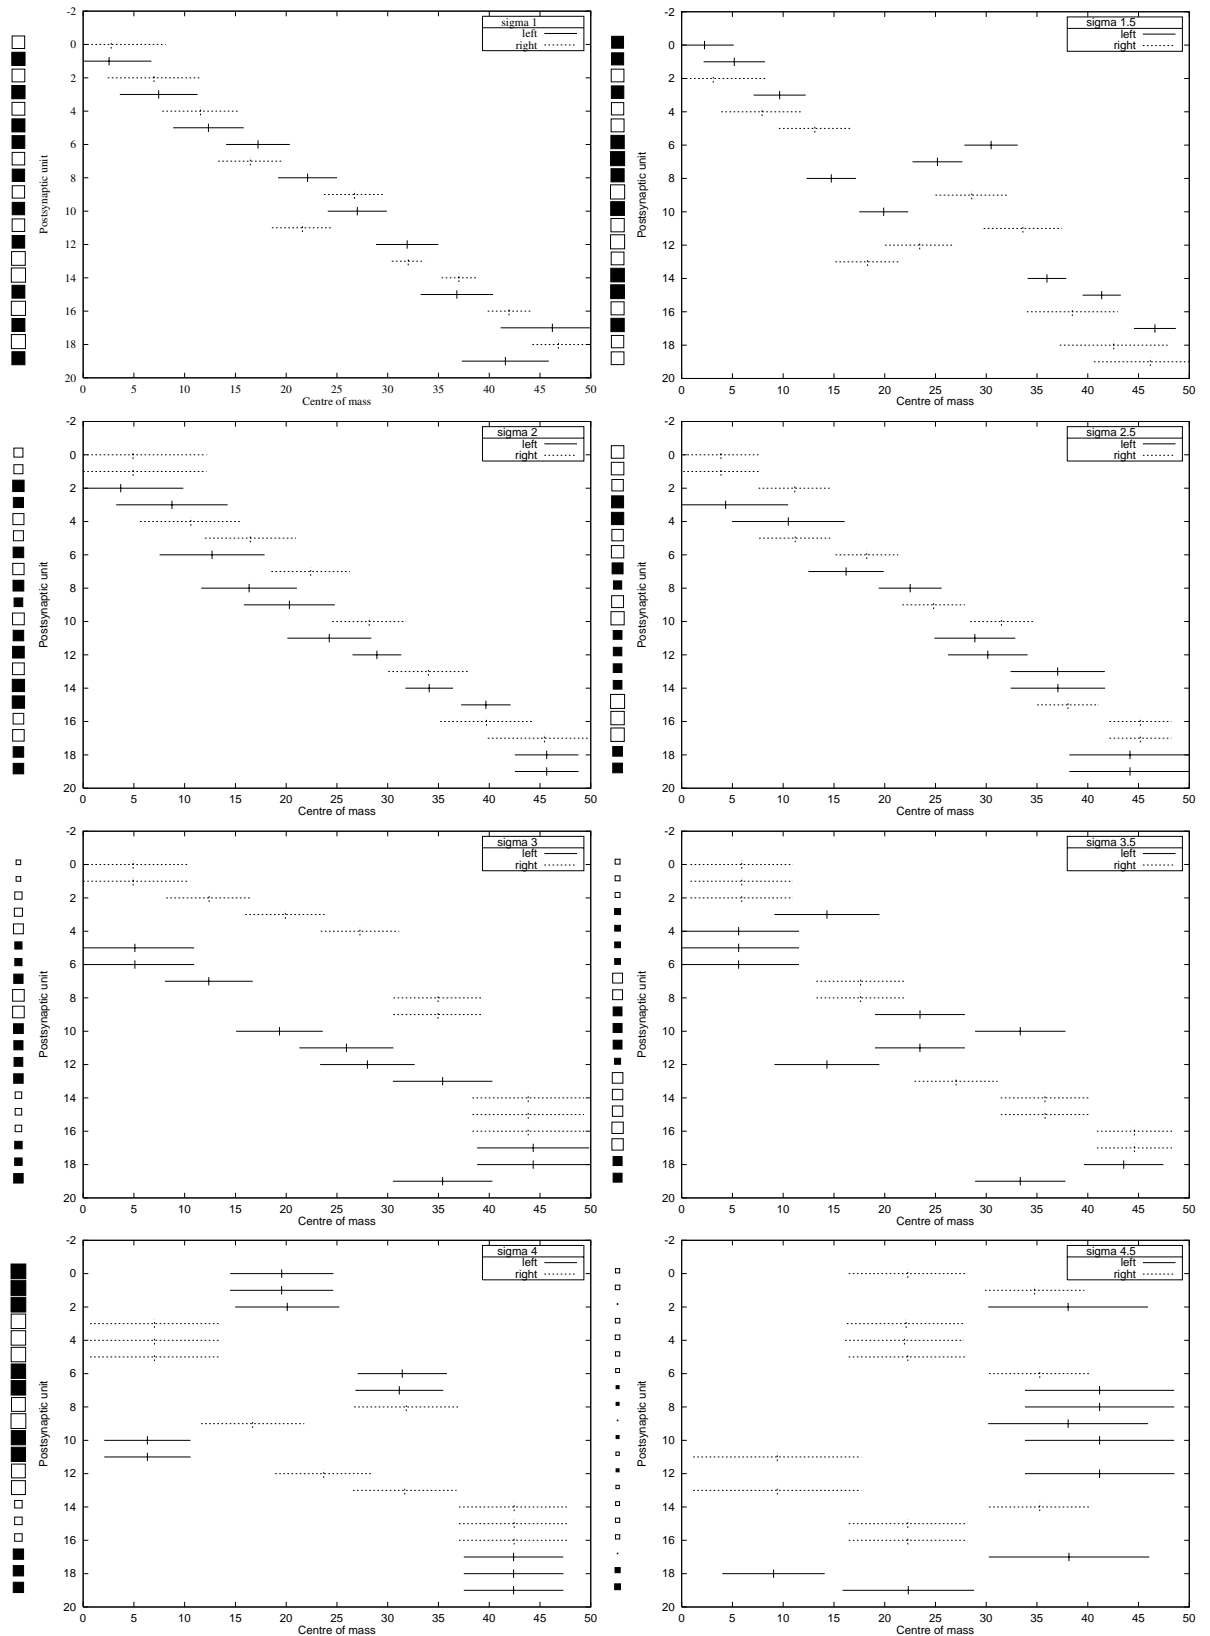


Figure 5.9: The effect of wave width upon topography and ocular dominance. The two retinæ innervate a one-dimensional array of twenty postsynaptic units. Eight experiments were run using the same parameters except for σ_w which increased from 1.0 to 4.5 in steps of 0.5. The value of σ_w is given in the key of each plot. The ocularity plot for each network is given to the left of the topographic plot, and aligned with the postsynaptic unit in the topographic plot.




Figure 5.10: The effect of wave width upon receptive field width. Plots for the left eye are offset slightly to the left from the true value of sigma so that data points for both the left and right eyes can be clearly seen. Error bars indicate ± 1.0 standard deviation from the mean value.

5.5 The effect of polarity-specific wave patterns upon LGN sublamination

In the mature LGN, each relay cell receives input from only on- or off-centre retinal cells, but not both. Relay cells are therefore classified as either ‘on-centre’ or ‘off-centre’ according to the dominant polarity (on- or off-centre) of retinal inputs. Like-centred LGN cells also tend to cluster into sublamina, as described in Section 2.3.5. It has been suggested that development of these sublamina is driven by activity-dependent processes, although until recently no polarity-specific patterns amongst retinal waves had been discovered which could drive such segregation (Morgan & Thompson, 1993). The first signs of polarity-specific firing patterns in the waves were recently reported by Wong and Oakley (1996). At P9 in the ferret retina, both on- and off-centre cells participate in retinal waves, and cannot be distinguished by their activity patterns. By P16 however, an asynchrony in the wave patterns of on- and off-centre cells has developed: on-centre cells tend to fire less often than off-centre cells. This means that when on-centre cells are active, neighbouring off-centre cells are also active, but when off-centre cells are active, neighbouring on-centre cells are usually quiet (Wong & Oakley, 1996).

5.5.1 Representation of on- and off-centre retinal units

To investigate whether these polarity-specific patterns drive LGN sublamination, the model retina was modified to represent both on- and off-centre units. In the mature retina, the number of on- and off-centre retinal ganglion cells is roughly equal, although there are slightly more off cells (Wassle, Boycott, & Illing, 1981a; Wassle, Peichl, & Boycott, 1981b). The relative percentages of these cells in the developing retina though are still unknown because on- and off-centre units cannot be morphologically discriminated until the period of segregation in the inner plexiform layer (P9 onwards; see Table 2.3). Counts of on- and off-centre units in ferret retina during P9–21 suggest a roughly equal number of on- and off-centre units during the period of on- and off-centre

sublamination (Wong & Oakley, 1996). On- and off-centre cells independently tile the ganglion cell layer, with the nearest neighbouring unit nearly always a unit of the opposite sign (Wassle et al., 1981a, 1981b). For the model retina therefore, presynaptic units were alternately labelled on-centre and off-centre (starting with the leftmost unit as on-centre). This arrangement ensured an equal number of on- and off-centre units arranged such that the nearest neighbouring unit was always of the opposite polarity.

To mimic the polarity-specific firing of retinal cells, retinal inputs were generated using the same procedure as listed in Section 4.3 and then slightly modified. At each time step, after a wave has been positioned on a retina, a value p is chosen at random from the uniform distribution $[0,1]$. This value controls whether the activity of on- or off-centre units should be set to zero according to the rules in Table 5.5.

if $p \leq p_{\text{on}}$	set $x_i = 0 \forall i = 2, 4, \dots, X_{\text{pre}}$	(keep on-centre activity)
else if $p \leq p_{\text{off}}$	set $x_i = 0 \forall i = 1, 3, \dots, X_{\text{pre}} - 1$	(keep off-centre activity)
else	change nothing	(keep both on- and off-centre activity)

Table 5.5: Rules for generating polarity-specific waves.

This formulation allows us to vary the correlations between on- and off-centre units by varying the values of p_{on} and p_{off} . To mimic the late onset of polarity-specific waves, we only introduce this extra mechanism after a fixed number of epochs have elapsed (1200 epochs for the experiments reported here). We also allow the width of the waves to increase during the simulation so that neighbouring on- or off-centre units are correlated. This increase in wave width during development is consistent with wave recordings from the ferret retina (Wong et al., 1993, see also Section 5.6). (Narrow waves are not wide enough to produce sufficiently high correlations between neighbouring on- or off-centre units.) Table 5.6 lists the parameters used in the following simulations for generating polarity-specific waves along with details of the other simulation parameters.

Epoch	Action
0	Wave width set to $\sigma_w = 0.85$
400	Wave width increased to $\sigma_w = 2.00$
1200	Polarity-specific wave patterns introduced

Table 5.6: Parameters for the polarity-specific wave experiments. To allow the network more time to develop, t_r was set to 700 epochs. Values for p_{on} and p_{off} are given in the text. All remaining simulation parameters were the same as in Table 4.1, with initial weights the same as in Figure 4.10.

5.5.2 Measuring the degree of polarity dominance

The segregation of on- and off-centre inputs into postsynaptic units is assessed by taking a measurement of the fraction of input that comes from on-centre units. This measure is analogous to the measure of ocular dominance, and so will be called ‘polarity dominance’. For a postsynaptic

unit j , the polarity dominance is denoted f_j , and is defined as:

$$f_j = \frac{f_j^{\text{on}}}{f_j^{\text{on}} + f_j^{\text{off}}} - \frac{1}{2}$$

$$f_j^{\text{on}} = \sum_{i=1}^{N_{\text{pre}}/2} w_{(2i-1)j} \quad f_j^{\text{off}} = \sum_{i=1}^{N_{\text{pre}}/2} w_{(2i)j}$$

The sign of f_j shows whether the dominant polarity of input unit is on-centre (+) or off-centre (−), regardless of the eye of origin. The magnitude of this value indicates the degree to which the unit receives input from presynaptic units of the dominant polarity. For each postsynaptic unit, the value of f_j varies between $[-0.5 : +0.5]$. A value of $+0.5$ indicates that the unit receives all of its input from on-centred presynaptic units, and -0.5 indicates all the input comes from off-centred units. A value of 0.0 indicates that the unit receives equal input from on- and off-centre presynaptic units. To monitor the overall development of polarity dominance, we also take the mean and standard deviation of f_j , separately for all positive and negative values of f_j , at regular intervals during the simulation.

The value of f_j for each postsynaptic unit is visualised in the ocular dominance plot by drawing either a + or − sign centred over the top of the square indicating the ocularity status of the postsynaptic unit. The size of the + or − sign is scaled in proportion to the magnitude of f_j such that the maximum possible sign corresponds to a magnitude of 0.5 .

5.5.3 Anticorrelated polarity-specific waves


The first polarity segregation experiment was performed with $p_{\text{on}} = 0.5$ and $p_{\text{off}} = 1.0$. Choosing these values ensured that at each time step after polarity-specific waves were introduced, only on- or off-centre units (but not both at the same time) would be active, ensuring that units of opposite centre signs were anticorrelated.

The correlation matrices during the three key stages of wave development are shown in Figure 5.11. First, up until 400 epochs the presynaptic correlations are the same as seen in previous experiments. Second, after 400 epochs the waves are widened, which broadens the correlations amongst more distant presynaptic units. Third, after 1200 epochs the polarity-specific waves ensure that neighbouring on-centre units (separated by an off-centre unit) are more correlated than the nearest neighbouring off-centre units (and the same for neighbouring off-centre units).


Figure 5.12 shows the development of topography, ocularity and polarity segregation at different stages during development. Up until 1200 epochs, the network develops in a typical manner, although the ipsilateral eye occupies slightly more of layer A than normal due to the higher value of σ_w used after 400 epochs. At this point, most postsynaptic units still receive an equal mixture of on- and off-centre inputs, as shown in the ocular dominance plots and the mean values of f_j in Figure 5.13. After 1200 epochs, segregation of on- and off-centre inputs gradually begins. The segregation is boosted once the growth rule finishes after 1400 epochs. By around 2500 epochs, most postsynaptic units receive input from presynaptic units of only one polarity. However, although individual postsynaptic units are responsive to either on- or off-centre units, there is no sublamination present: neighbouring postsynaptic units within the same row of the LGN do not always have the same polarity. Instead, each row of the LGN has a mixture of on- and off-centre

postsynaptic units.


The introduction of polarity-specific waves did not radically affect the nature of the topographic mapping, as shown by the topographic plots in Figure 5.12. The final projection columns for the network, shown in Figure 5.14, are well organised, with only a few units in the wrong position. The average wave width at the end of development was slightly bigger than for the previous simulations (mean = 3.795, s.d. = 0.695, compared with mean = 2.496, s.d. = 0.737 for L4R8 of Figure 4.11 and mean = 2.922, s.d. = 0.894 for L2R4 of Figure 4.11). This increase in receptive field width was expected due to the increase in the width of the retinal waves after 400 epochs.



(a) Epoch 400



(b) Epoch 1200



(c) Epoch 3000

Figure 5.11: Slices of the input correlation matrix before and after the onset of polarity-specific wave patterns. Each correlation plot shows the correlation between presynaptic unit 25 (an on-centre unit) and all other presynaptic units. Conventions for correlations are the same as in Figure 5.2. ($p_{\text{on}} = 0.5, p_{\text{off}} = 1.0$.) (a): Initial correlations from the start of the simulation up to 400 epochs, just before waves are widened. (b): Correlations from epoch 1 up to epoch 1200, just before polarity-specific waves are introduced. (c): Correlations from epoch 1 to the end of the simulation at 3000 epochs.

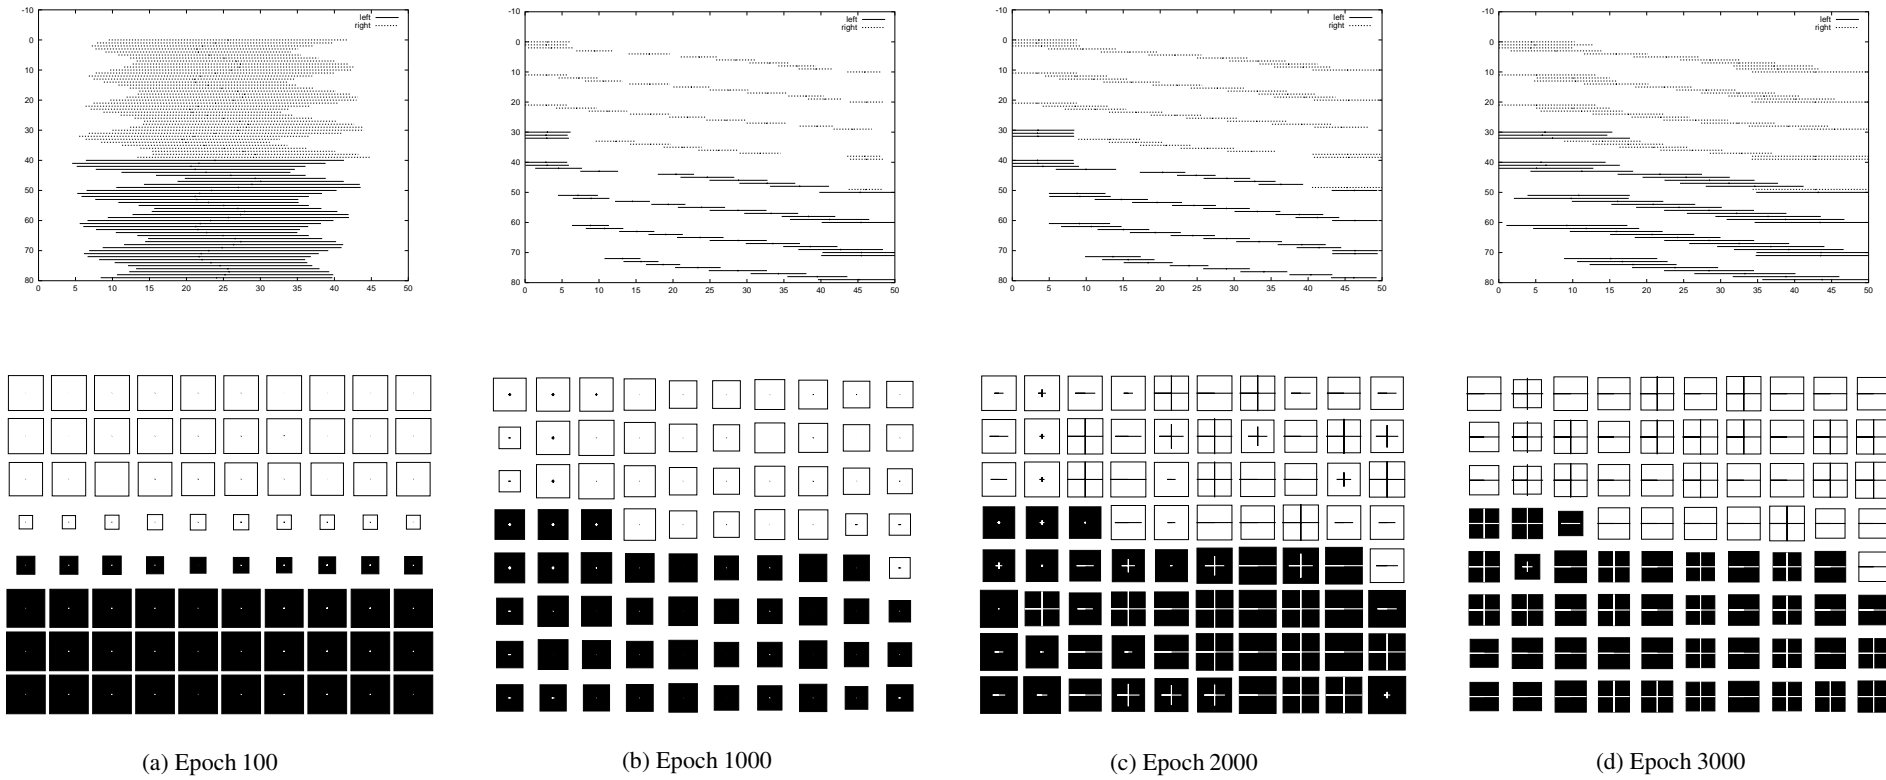


Figure 5.12: Time course of polarity dominance using anticorrelated on- and off-centre waves. The topography plot (above) and ocular dominance plot (below) is shown at four different points during development. ($p_{\text{on}}=0.5, p_{\text{off}}=1.0$).




Figure 5.13: Development of polarity dominance in postsynaptic units. The mean value of f_j over all postsynaptic units is shown at regular intervals during development. Error bars indicate ± 1.0 standard deviation from the mean value. This experiment was continued after 3000 epochs for another 2000 epochs to ensure that the network had converged. It is clear that by around 3000 epochs, polarity segregation is almost complete.






Figure 5.14: Final projection columns for the network with polarity-specific waves as input. The corresponding topography and ocular dominance plots for this experiment are shown in Figure 5.12.

5.5.4 Weakening the anticorrelations between on- and off-centre inputs



The parameters p_{on} and p_{off} control the nature of the correlations between the on- and off-centre units. In the previous section, the values $p_{\text{on}} = 0.5, p_{\text{off}} = 1.0$ were used. This ensured that once the polarity-specific waves had been introduced, an on-centre unit and off-centre unit would never be jointly active, maximising the chance that units of different polarity would not jointly innervate the same postsynaptic unit. However, such anticorrelations have not been found in the developing retina. Instead, as mentioned before, when on-centre cells are active, neighbouring off-centre cells also tend to be active, but not vice-versa (Wong & Oakley, 1996). These weaker anticorrelations can be modelled by reducing the value of p_{off} below 1.0 to allow some retinal waves to have both on- and off-centre activity (using the default rule from Table 5.5).

Since off-centre cells are active more often than on-centre cells during the period of polarity-specific waves, a set of experiments were performed with p_{on} reduced below 0.5, and p_{off} set to $p_{\text{on}} + 0.5$. This means that off-centre waves are present half of the time, with the rest of the time divided between on-centre waves and mixed on- and off-centre waves. As p_{on} is reduced to 0.0, only off-centre and mixed on- and off-centre waves are generated. This is the situation found in the developing retina (Wong & Oakley, 1996).



Figure 5.15 shows the results of development using $p_{\text{on}} = 0.48, 0.46, 0.44, 0.42$ with $p_{\text{off}} = p_{\text{on}} + 0.5$ in each case. The plots of polarity segregation for each experiment clearly show that as p_{on} decreases, the overall degree of polarity segregation is greatly reduced. The other network features, ocular dominance and topography, were unaffected. The existence of mixed on- and off-centre waves therefore specifically inhibits polarity segregation in this model.





(a) $p_{\text{on}} = 0.48; p_{\text{off}} = 0.98$



(b) $p_{\text{on}} = 0.46; p_{\text{off}} = 0.96$



(c) $p_{\text{on}} = 0.44; p_{\text{off}} = 0.94$



(d) $p_{\text{on}} = 0.42; p_{\text{off}} = 0.92$

Figure 5.15: Polarity segregation for different correlations between on- and off-centre presynaptic units. For each value of p_{on} and p_{off} the mean value of f_j throughout development and the final ocular dominance plot are shown.

5.6 Comparison with biological data

To date, little quantitative data has been published on the spatio-temporal properties of the retinal waves. Instead, more attention has been paid to examining the mechanisms underlying the generation and propagation of the waves (Wong et al., 1995; Feller, Wellis, Stellwagen, Weblin, & Shatz, 1996). Despite the relative paucity of data, in this section we compare the limited biological data of the wave properties with those used in the model.

5.6.1 Temporal characteristics of retinal waves

Wong et al. (1993) provide average values for the burst duration and mean inter-burst interval for ferret retina. These values allow a broad comparison to be made between the temporal properties of the waves used in the simulations and the waves found in the natural retina. Table 5.7 summarises the results of this comparison.

Even when we take the maximum values for the overall probability of eye activity (O_{\max}), we see that the values of O are quite small. Table 5.1 shows how O varies with the parameter p_w from the computer simulations. From this, we can see that even when the waves are present most often on the retina at P21, this corresponds to a value for p_w much less than 0.020. It is important to note however that these calculations are only order of magnitude figures: for example, we do not compare the relative sizes of the retina and the speed at which the wave travels across the retina.

If the estimates are roughly correct, it would indicate that the activity-driven developmental mechanisms need to be robust enough to account for the long periods of silence between waves. We have shown here that it is possible for the covariance rule to be adapted to prevent weight adaptation when input and output are below threshold, to account for the long periods of silence between waves. Alternative mechanisms may also allow the network to cope with long periods of silence.

5.6.2 Spatial characteristics of retinal waves

The spatial characteristics of the waves have been analysed in some detail by examining the correlation index between pairs of ferret retinal cells (Wong et al., 1993). These correlation indexes reveal two features about the waves:

1. The correlation between pairs of retinal cells decreases with the distance between cells. This fits in with the model assumption that the retinal waves can be modelled as Gaussian functions of position from the centre of the wave. Providing an estimate of the width of Gaussians from the biological data is complicated, however, due to the arbitrary correspondence between the size of the natural and model retina.
2. During the first few postnatal weeks the waves tend to get wider, producing weaker correlations over a wider distance. This continues until around P30 when the broad waves eventually become constant maintained levels of activity, similar to that seen in the mature retina.

During development, the waves therefore get progressively wider until they eventually change into maintained background levels of activity. More details about the spatial width of retinal

Day	A	$I \pm sd$	O_{\min}	O_{\max}
P0	1–4	75 ± 6	0.012	0.055
P5	2–4	101 ± 20	0.016	0.047
P12	2–7	71 ± 20	0.022	0.121
P15	2–7	60 ± 9	0.028	0.121
P21	2–9	28 ± 3	0.061	0.265

The overall probability of eye activity, O , is given using:

$$O = \frac{A}{A + I}$$

A = Mean wave duration in seconds (active)

I = Mean interval between waves in seconds (inactive)

Minimum and maximum values for O were calculated by taking the lowest and highest values for O out of the four calculations using the two extreme values of A and the two values $I \pm sd$.

Table 5.7: Burst and inter-burst intervals at various postnatal stages in the ferret. The range of values for the overall probabilities [O_{\min} , O_{\max}] can be compared to the values used in the simulations, shown in Table 5.1. Data taken from (Wong et al., 1993, p926).

waves and more detailed modelling of the retina are required before we can use the biological data to guide selection of model parameters such as wave width. As is shown in the work with the on- and off-centre units in Section 5.5, wide waves do not inhibit topographic refinement, as long as narrower waves are initially used to lay out the overall topography.

5.6.3 The nature of correlations between on- and off-centre retinal units

In Section 5.5 it was assumed that polarity-specific waves were either on-centre, off-centre or mixed. The results from the model showed that as the frequency of on-centre waves decreased, so did the degree of polarity segregation. The biological data does not support this result because segregation occurs in the absence of on-centre waves. Wong and Oakley (1996) note that during the period of polarity segregation, the on-centre cells fire for only around 25–35% of the time that off-centre cells fire, and that when on-centre cells fire, neighbouring off-centre cells nearly always also fire. Assuming that no on-centre waves are generated, this corresponds to simulation parameters $p_{\text{on}} = 0.0$ and $p_{\text{off}} = 0.66$. Given the failure of the network to develop polarity-specific units when $p_{\text{on}} = 0.42$ and $p_{\text{off}} = 0.92$ (Figure 5.15), it was not expected to see the network segregate using these more realistic values. An extra simulation was run using $p_{\text{on}} = 0.0$, $p_{\text{off}} = 0.66$, which produced similar results to those for $p_{\text{on}} = 0.42$, $p_{\text{off}} = 0.92$ (mean value of f_j for on-centre units = 0.018, s.d. = 0.016, $n = 36$; mean value of f_j for off-centre units = -0.014 , s.d. = 0.011, $n = 44$).

Here the results from the network are therefore in disagreement with the biological data. The current model can only segregate on- and off-centre inputs when the inputs are anticorrelated. These anticorrelations are only present much later in life, when the retina is performing its normal function of converting light into electrical signals (Mastrorarde, 1989). However, by this time, units in the LGN are already polarity-specific and are organised into sublamina. This rules out any

light-evoked anticorrelations between on- and off-centre cells driving polarity segregation.

5.7 Discussion

The formulation of the waves used in this thesis is relatively simple, but has allowed us to investigate the effect of the spatio-temporal wave properties upon development. This chapter has highlighted the following issues:

1. The probability of wave generation affects development. When the waves are frequently generated, the global retinotopic mapping breaks down although there is still some degree of local order. In contrast, the monocular status of postsynaptic units are hardly affected. At the other end of the scale, when waves are not generated very often, the covariance rule causes most weights to converge to the same value, preventing both receptive field refinement and ocular segregation. A modified weight adaptation rule, the active-covariance rule, is presented which does not increase weights when both input and output activities are below threshold. This restores normal development in the model under conditions of sparse input. A simple comparison with biological data shows that such sparse inputs are likely to be present in the natural retina during the period of retinogeniculate development.
2. The model can account for monocular deprivation results if the presynaptic normalisation is weakened. This is achieved by allowing the sum of weights from a presynaptic unit to be less than, rather than strictly equal to, a target value. Without this weaker form of normalisation, presynaptic units remain connected to the postsynaptic sheet despite their lack of activity. It would be interesting to see how the projections columns generated by the network (shown in Figure 5.7) compare with the projection columns taken from a monocularly deprived animal.
3. Wave width affects overall topography and ocularity segregation. As the width of waves is increased, although units still refine their receptive fields, the overall size of receptive fields increases and the overall topographic mapping breaks down. Network development is robust to wide waves however, as long as narrow waves are initially used to set up the overall topographic mapping. A similar result has recently been found for retinal ganglion cells in the developing turtle retina (Sernagor & Grzywacz, 1996).
4. The model predicts segregation of on- and off-centre inputs, but only in the presence of anticorrelated on- and off-centre activity. The current model cannot segregate into polarity-specific units using input correlations similar to those recently discovered in the ferret retina (Wong & Oakley, 1996). In contrast, recent modelling work using a BCM rule has shown that a single postsynaptic unit can become responsive to only one polarity using these correlations (Lee & Wong, 1996).

5.7.1 New mechanisms introduced

This chapter has introduced two new mechanisms to allow the model to account for different aspects of development. Both of these are now discussed.

The active-covariance rule

The active-covariance rule specifies no weight change when both pre- and postsynaptic activities are below threshold (case four of the covariance rule). This rule allows the network to develop as normal when waves are rarely present on the retina. The effect of ignoring case four is not quite the same as preventing weight adaptation when both eyes are quiet: case four is normally invoked when both eyes are quiet, but is also invoked when at least one eye is active. Hence the fraction of time case four is invoked should always be greater than the value for O_0 — the fraction of time for which both eyes are quiet. (The remaining time case four is used is when there is some input to the network, but not enough to drive all postsynaptic units above threshold). Comparing the values of O_0 and the usage of case four in Tables 5.1 and 5.3, the case when both eyes are quiet accounts for almost all of the usage of case four of the covariance rule. The active-covariance rule therefore effectively prevents weight update in the absence of any presynaptic activity.

Case four of the covariance rule has attracted attention in other related research. First, Willshaw and Dayan (1990) noted that the covariance rule has “the biologically undesirable property that changes can occur when neither pre- nor postsynaptic neuron is active”, and mention that this case just changes weights by a constant amount in the absence of any activity. This is a positive term however, and so it does not represent a decay term as might be expected. Second, Montague et al. (1991) have shown that the case is not required for development of ocular dominance columns, although in their case the network’s inputs were anticorrelated and they did not compare the results of using the covariance rule with and without this case. (The inclusion of case four may inhibit ocular dominance in their model.)

Weakening presynaptic normalisation

In Section 4.2.1, the parameter r_s was used to weaken subtractive normalisation so that the network could develop using subtractive presynaptic normalisation. In this chapter, we have presented another way of weakening normalisation such that $\sum_{j=1}^{N_{\text{post}}} w_{ij} \leq T_{\text{pre}}$. (Here we have only weakened divisive normalisation, but subtractive normalisation can be similarly weakened by ignoring the case when normalisation should add a constant value to each unit.) With this constraint, the weights are allowed to be positioned within the constraint space, rather than strictly on the constraint surface. This weaker constraint allows the model to replicate monocular deprivation results showing that if an eye is not active, it will eventually disconnect from the postsynaptic sheet (all of the weights associated with the deprived eye will go to zero). In comparison, the standard form of normalisation prevents presynaptic units from disconnecting from the postsynaptic sheet, regardless of the activity levels of presynaptic units.

5.7.2 The polarity experiments

Unfortunately, the results from the polarity-specific wave experiments are at odds with the current biological data on two counts. First, postsynaptic units only become responsive to inputs of one polarity when the on- and off-centre inputs are anticorrelated. As these anticorrelations are weakened in correspondence with the biological data, polarity segregation is inhibited. Recent success at polarity segregation of one postsynaptic unit in the presence of weaker anticorrelations using the BCM rule (Lee & Wong, 1996) suggests that the covariance rule may not be suitable for modelling polarity development. Second, even when we assume that on- and off-centre in-

puts are anticorrelated, although postsynaptic units become responsive to one polarity, there is no consistent sublamination effect. This lack of sublayering is due to two reasons:

1. No systematic bias. In the ferret, the sublayering is consistent across animals, so that the on-centre units always cluster in the rostral regions of a layer, and the off-centre units into the caudal regions (see Figure 2.4). Similar clustering tendencies, although slightly weaker, are found in the cat. This would seem to indicate some form of systematic bias in the natural system, either in the initial arrangement of connections or through chemical gradients. However, no such bias has been suggested in the biological literature. It is possible to introduce a polarity bias into the model, although there are difficulties with putting the bias either in the initial weights or using chemical gradients.

Weights could be initialised so that certain postsynaptic units received more initial weight strength from either on- or off-centre presynaptic units. This can be objected to on both theoretical and practical grounds. First, assuming development begins before the period of polarity segregation (around P16 in the ferret), some retinal ganglion cells are still classified as mixed on-off cells since their dendritic trees have not segregated into one sublayer of the IPL. This segregation is activity-dependent, complicating prediction of a retinal cell's polarity (Bodnarenko & Chalupa, 1993; Bodnarenko, Jeyarasasingam, & Chalupa, 1995). It is therefore difficult to know beforehand the polarity of each retinal ganglion cell. Without knowing the polarity of a presynaptic unit it is difficult to know how to bias its weights in advance to help the layering. Second, even if we know the polarity of each presynaptic unit, any initial weight bias may be lost during the early stages of development before the onset of polarity-specific waves.

Chemical gradients could also be used to bias on-centre units to connect to the top of each eye-specific layer. Figure 5.16 shows two possible biases that could be used. The first should ensure that eye-specific layers are preserved, but this entails that the chemical bias knows where the division between layer A and layer A1 occurs. The second bias is much simpler, but may have the side-effect of destroying the eye-specific layers if the polarity bias is too strong. Both weight bias and chemical gradients therefore have different drawbacks.

2. Nature of growth rule. The growth rule used here does not discriminate between on- and off-centre units. Any weight increases favouring an on-centre unit will be distributed to both neighbouring on- and off-centre units by the growth rule, and hence there is no significant segregation of polarity until the growth rule terminates. Once the growth rule stops, postsynaptic units can sharpen their responses to units of one polarity. There is no longer any direct interaction between neighbouring postsynaptic units however, and so neighbouring units cannot influence each other with respect to the polarity to which it becomes dominant.

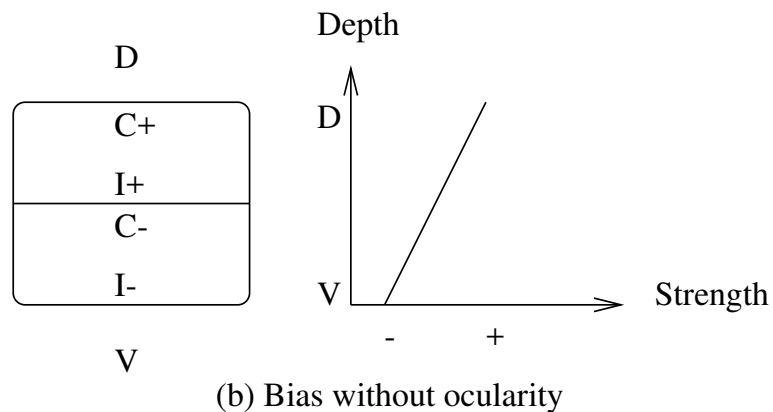



Figure 5.16: Chemical gradients to bias the formation of on- and off-centre sublamina in the LGN. On the left a coronal slice of the cat LGN is shown, and on the right a possible chemical gradient, showing the strength of some chemical as a function of LGN depth. The strength of chemical indicates an affinity for on-centre (+) cells. (a): Chemical bias with ocularity information to prevent disruption of eye-specific layering. (b): A simpler chemical bias without ocularity information. This second bias will probably affect eye-specific segregation as a side-effect of polarity segregation, causing the LGN to be ordered according to polarity rather than ocularity. Key: C – contralateral eye input. I – ipsilateral eye input. D – dorsal. V – ventral.

Chapter 6

Factors affecting ocular dominance stripe formation

6.1 Introduction

This chapter considers the mechanisms driving the development of ocular dominance stripes in some of the main models reviewed in Chapter 3. The first half of the chapter considers why stripes, rather than segregated layers, are formed in these models. Simulations of the model by Obermayer et al. (1991) are replicated, which show how ocularity variance controls stripe width in a two-dimensional postsynaptic sheet. In the second half of this chapter, the Obermayer model is extended with a three-dimensional postsynaptic block to see how topography and ocularity map into a three-dimensional output space representing the LGN.

6.2 Reasons for ocular dominance stripe formation in models

Most of the models of ocular dominance reviewed in Chapter 3 form stripes of alternating eye input rather than completely segregated regions for each eye. This is in contrast to the model investigated in this thesis which generates completely segregated eye-specific regions. There are three different arguments for why stripes form in models:

1. Lateral interactions
2. Competition between topography and ocularity
3. Primary/secondary feature argument

Each of these arguments are now considered in turn.

6.2.1 Lateral interactions

Some models of ocular dominance do not consider the development of topography (Swindale, 1980; Miller et al., 1989). In these models, the major determiner of stripe width is the extent of the lateral interactions amongst postsynaptic units. The local excitatory connections encourage nearby pairs of units to be correlated, and therefore adapt to similar stimuli. The longer-range inhibitory connections perform the opposite task of ensuring that pairs of distant units are anticorrelated,

preventing them from adapting in the same direction to the input stimuli. Miller's model is slightly more complicated: in the absence of inhibitory lateral connections, as long as there is presynaptic normalisation then stripes still develop (Miller et al., 1989).

6.2.2 Competition between topography and ocularity

For those models considering the development of both ocular dominance and topography, stripes arise through competition between these two features as well as the extent of the lateral interactions (von der Malsburg, 1979; Goodhill, 1992; Andrade & Moràn, 1996). Neighbouring presynaptic units from the same eye ("same-eye units") tend to connect to the same region of the postsynaptic sheet. Similarly, units at the same retinotopic position but from different eyes ("same-position units") tend to connect to the same postsynaptic region. If these two tendencies are roughly equal, a compromise is made whereby neighbouring postsynaptic units receive inputs from either units from the same eye but different positions, or from the same position but different eyes. Stripes form as a consequence of these two competing tendencies.

What are the mechanisms underlying the tendency for presynaptic units to connect to the same region of the postsynaptic sheet? The correlated activity of same-eye units ensures that they connect to similar parts of the postsynaptic sheet. For same-position units however, there are two different mechanisms at work. First, if there are positive between-eye correlations, pairs of same-position units are likely to be positively correlated. This correlated activity will ensure that they connect to the same region of the postsynaptic sheet. Second, the initial weights of the network are usually biased so that same-position units initially connect to roughly the same region of the postsynaptic sheet. So, even in the absence of any positive between-eye correlations, same-position units will connect to similar regions of the postsynaptic sheet as long as the initial weights are appropriately biased. (This is seen for example in Goodhill's model (Goodhill, 1992, p90) when the correlation parameter $h = 0$ which produced no between-eye correlations.)

Under this framework, stripe width is affected by the relative strengths of the correlations of same-eye units and same-position units. When pairs of same-eye units are more correlated than same-position units, the stripes tend to be quite wide — many neighbouring postsynaptic units respond to the same eye. As the strength of the same-position correlations increase relative to the same-eye correlations, stripe width decreases as some neighbouring postsynaptic units tend to connect to same-position units in preference to same-eye units.

6.2.3 Primary/secondary feature argument

The work by Obermayer et al. (1991) put a different perspective on why stripes arise and the mechanisms controlling stripe width. They considered a five-dimensional feature vector, (x, y, ϕ, q, z) , to represent a pattern of activity distributed across two retina. Features x and y represented the centre of mass of the stimuli, ϕ and q the orientation of the stimuli, and z the ocularity of the stimuli. Feature vectors were input to a Kohonen network with a two-dimensional arrangement of postsynaptic units. In this model, the Kohonen network was required to perform dimension reduction: to map from a five-dimensional input space to a two-dimensional output space. This had previously been investigated as an abstract mapping problem (Kohonen, 1988; Ritter & Schulten, 1988). In such situations, the network automatically discovers the elements of the input vector

with highest variance, and represents them in the network. Other elements of the input vector with variance below a critical value do not get represented in the map (Ritter & Schulten, 1988). This is referred to as the “automatic selection of feature dimensions” (Kohonen, 1988). As is shown in more detail in the next section, the variance of the ocularity feature controls stripe formation: the higher the ocularity variance (compared with the other input variances), the wider the ocular dominance stripes.

6.3 Obermayer’s model for stripe formation

All of the models mentioned in the last section provide different reasons for how stripe width varies under different conditions. Out of all of the models, the model by Obermayer et al. (1991) is arguably the simplest because it uses feature vectors to represent the neural activity distributed across two retinæ. This model has therefore been chosen to investigate the nature of stripe formation in the cortical simulations. We also extend the postsynaptic sheet into a three-dimensional block so that the model can be applied to the problem of retinogeniculate development.

6.3.1 Implementation details of Obermayer experiments

The network consists of a presynaptic sheet with three units fully connected to a set of postsynaptic units. The postsynaptic units are arranged into a three-dimensional block for the purposes of neighbourhood weight updating. (For the two-dimensional simulations, the Z dimension of the postsynaptic block was set to one.)

The initial weights are set at random, with no topographic or ocular bias. One iteration of the model consists of generating and presenting a feature vector, calculating postsynaptic unit activations and updating the weights of the winning and neighbouring units. One epoch of the model corresponds to 100 iterations, after which various parameters, such as the weight-update rate and size of neighbourhood, are updated. Table 6.1 describes the equations and parameters used for these experiments.

6.3.2 Visualisation of maps formed in Obermayer experiments

Each postsynaptic unit receives inputs via three weights. The first two weights code for the centre of mass and the third weight codes for the ocularity to which the unit is most responsive. Visualisation of each postsynaptic unit’s preferred stimulus is straightforward. The centre of mass for each postsynaptic unit j is drawn on a graph at the point $(x = w_{1j}, y = w_{2j})$. Lines are drawn between the points of neighbouring postsynaptic units to indicate the location of neighbouring postsynaptic units. The ocularity of each postsynaptic unit is encoded in $z = w_{3j}$. The value of z for all postsynaptic units is visualised using a Hinton diagram — the size of the square is proportional to the magnitude of z (scaled to a maximum size of z_{pat}) and the colour of the square indicates the sign of z (black for negative values representing dominant left-eye input and white for positive values representing dominant right-eye input).

6.3.3 How the variance of ocularity affects stripe width

To illustrate the principle of feature selection within the standard cortical model, several experiments similar to those presented in (Obermayer et al., 1991) were replicated. Three-dimensional

Inputs	$x_1 = \text{rnd}(0, 15), x_2 = \text{rnd}(0, 15), x_3 = \pm z_{pat}$
Activation rule	$y_j = \sum_{i=1}^3 (w_{ij} - x_i)^2$
Winner	Find m : $y_m \leq y_j \forall j$
Weight modification	$\Delta w_{ij} = \epsilon(t) d(j, m) (x_i - w_{ij})$ (t is the epoch number)
Distance measure	$d(j, m) = n(P_x(m) - P_x(j), P_y(m) - P_y(j), P_z(m) - P_z(j))$
Index function	$j = P_x(j) + (X_{\text{post}} \times P_y(j)) + ((X_{\text{post}} \times Y_{\text{post}}) \times P_z(j))^\dagger$
Neighbourhood function	$n(x, y, z) = \exp(-\frac{x^2}{\sigma_x^2} - \frac{y^2}{\sigma_y^2} - \frac{z^2}{\sigma_z^2})$, $ x \leq r_x, y \leq r_y, z \leq r_z$
Rate of weight-update	$\epsilon(t) = \begin{cases} 0.05 & t \geq 1000 \\ \min(0.05, \epsilon_0 - \text{floor}(t/50) \times 0.1) & \text{otherwise} \end{cases}$
<p>† Given the index j for a postsynaptic unit, its position within the three-dimensional postsynaptic block, $(P_x(j), P_y(j), P_z(j))$, can be calculated using this equation.</p>	

Parameter	Value	Epoch	(r_x, r_y, r_z)	$(\sigma_x, \sigma_y, \sigma_z)$
Number of epochs	1000	0	(5,5,5)	(3.0, 3.0, 3.0)
Initial weight-update rate, ϵ_0	0.8	200	(3,3,3)	(2.0, 2.0, 2.0)
Inputs per epoch	100	500	(1,1,1)	(1.0, 1.0, 1.0)
Postsynaptic block size $(X_{\text{post}}, Y_{\text{post}}, Z_{\text{post}})$	(32,32,Z)			

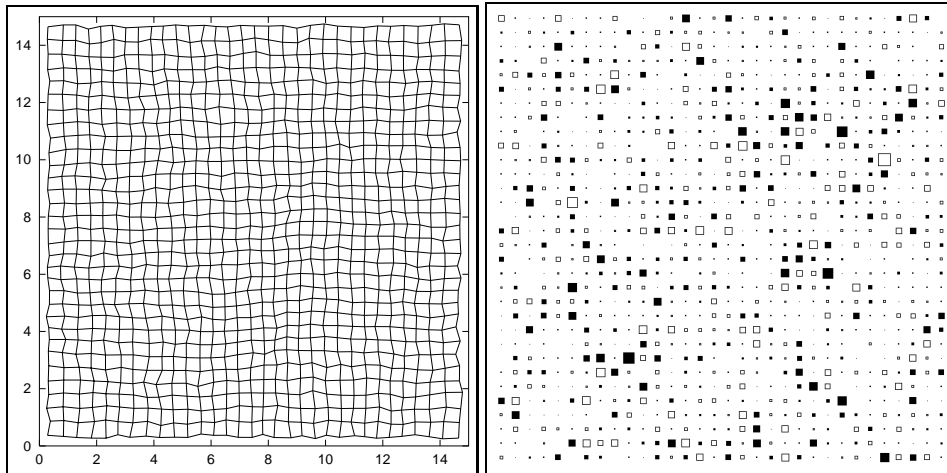
Table 6.1: Network equations and parameters used for Obermayer experiments. The only parameters varied in the experiments presented in this chapter were Z , which controlled the size of the postsynaptic block, and z_{pat} , which determined the degree of ocularity variance in the input vector. Network equations are the same as those used by Obermayer et al. (1991).

input vectors (x, y, z) were used as input to a Kohonen network with a two-dimensional sheet of postsynaptic units of size $X_{\text{post}} = 32, Y_{\text{post}} = 32$ (assuming $Z_{\text{post}} = 1$). Each feature vector represented a pattern of activation across two notional retinas. The features x and y coded for the centre of mass of the activation, whereas z encoded the degree of monocularity of the activation over the two notional retinas. The higher the value of z , the more monocular the input stimuli. Rather than presenting patterns of activity on two retinas and then calculating the feature vector, the feature vectors were directly generated. The x and y elements were independently chosen at random from the uniform distribution $[0, 15]$ and z was set to $\pm z_{pat}$, with the sign of z chosen at random. z_{pat} was a control parameter which was varied in experiments.

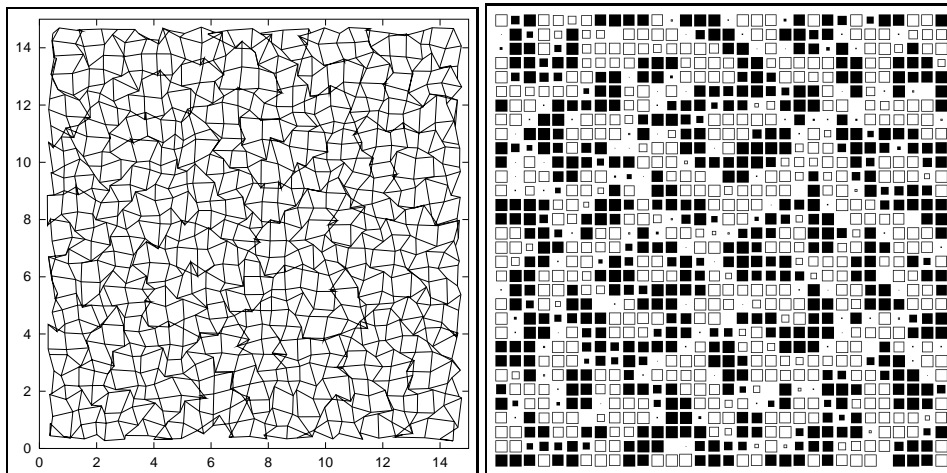
Figure 6.1 shows the effect of varying the z_{pat} parameter on map development. Table 6.2 shows the number of postsynaptic units responsive to each eye along with the mean ocularity strength. When z_{pat} is small, the ocularity component of the input vector is largely ignored, and the map smoothly represents the x, y parameters (Figure 6.1(a)). In this map, visual space is regarded as a primary feature because the map is ordered according to the values of (x, y) . As z_{pat} increases above its first critical threshold ($z_{pat} = 0.6$), the network has to code for all three features of the input vector, producing the slightly bumpy map of visual space. Thin stripes are present in the map's representation of the z parameter (Figure 6.1(b)), but visual space is still the primary map feature. In this map, ocularity is a secondary feature because the map represents the same value of z in more than one region of the postsynaptic sheet. As z_{pat} is further increased, the map of visual space becomes more convoluted as it has to account for the wider ocularity stripes (Figure 6.1(c)).

Once z_{pat} exceeds a second critical value ($z_{pat} = 2.0$), ocularity becomes the primary map feature and visual space is relegated to a secondary feature (Figure 6.2). In this network, each point of visual space is represented twice in the postsynaptic sheet, once for each eye. The map of visual space has folded over on itself, which is more clearly seen in Figure 6.2(b) when the topographic map is plotted separately for postsynaptic units with positive and negative values of z .

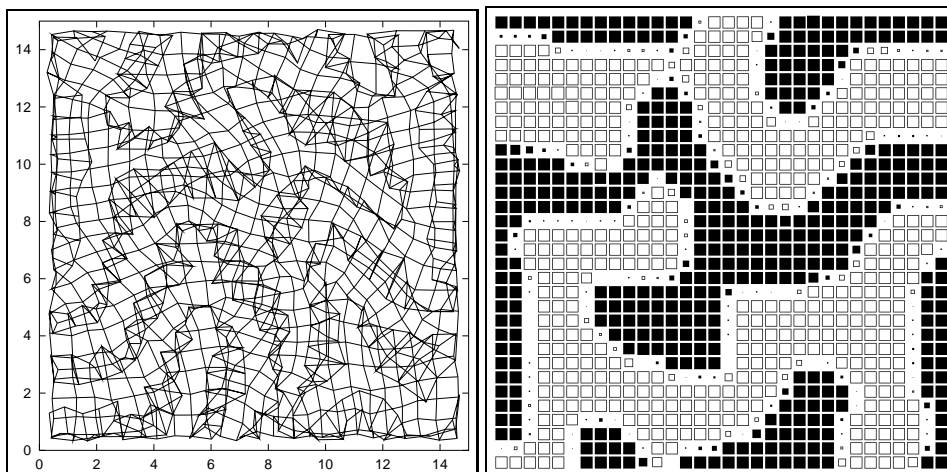
The feature vector representation used here is much more abstract than an explicit representation of neural activity in two retina as used in other models. However, in agreement with the other models, it also predicts that the more monocular the stimuli, the wider the strips. (Due to the way the feature vectors were generated, it is difficult to directly interpret z in terms of between-eye correlations. A simple approximation is that when $z = 0$, the two eyes are completely correlated, and as z increases, the between-eye correlations decreases.) The model also shows, however, that if the variance of the ocularity is sufficiently high, the ocular map segregates into two eye-specific regions.



(a) $z_{pat} = 0.2$

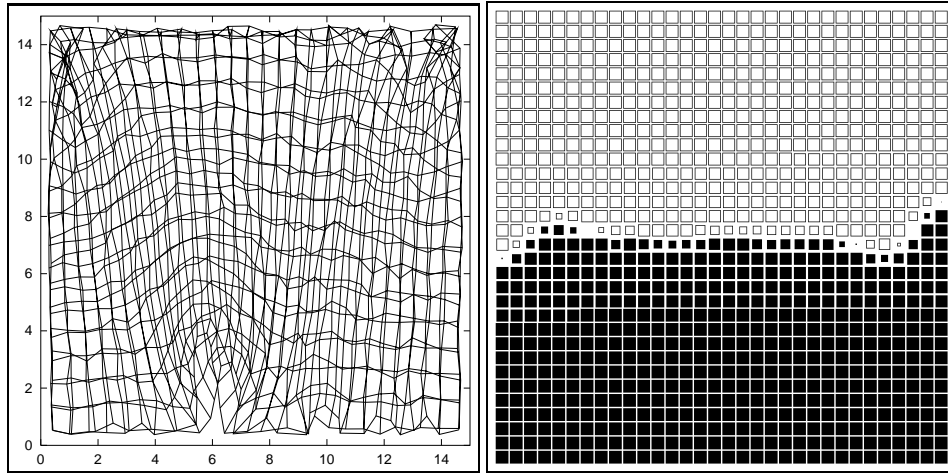


(b) $z_{pat} = 0.6$

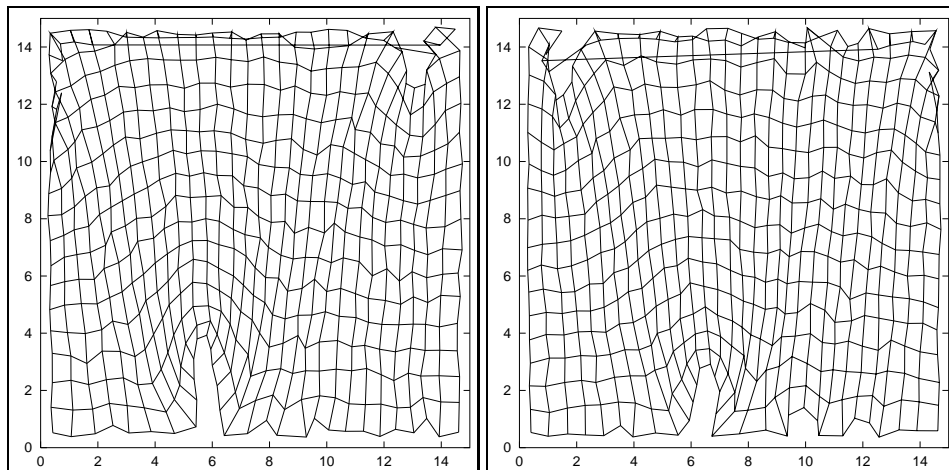


(c) $z_{pat} = 1.0$

Figure 6.1: Final topography and ocularity maps in a two-dimensional Kohonen network with feature vectors $(x, y, \pm z_{pat})$ as input and $z_{pat} = 0.2, 0.6, 1.0$. (a): $z_{pat} = 0.2$. (b) $z_{pat} = 0.6$. (c) $z_{pat} = 1.0$. See also Table 6.2 for mean z values.



(a) Topography and ocularity plots ($z_{pat} = 2.0$)



(b) Representation of visual space for each eye ($z_{pat} = 2.0$)

Figure 6.2: Final topography and ocularity maps in a two-dimensional Kohonen network with feature vectors $(x, y, \pm z_{pat})$ as input and $z_{pat} = 2.0$. (a): Topography and ocularity plots. Conventions for plots same as in Figure 6.1. (b): Topography plots for each eye. In the left (right) plot, the (x, y) values from weight vectors are plotted for units with a negative (positive) z value. Lines are drawn between the points representing nearest vertical and horizontal units with the same sign of z . The 32×32 postsynaptic sheet has divided into two maps of size 32×16 representing visual space independently for each eye. The kinks at the bottom of each map (above the point $x = 6$) arise from a rectangular map (of width 32 and height 16 units) uniformly covering a square input space.

6.4 A new model of retinotopic map formation in the LGN

In the previous experiments, the units in the postsynaptic sheet were arranged in a two-dimensional sheet. In the experiments presented in this section, the postsynaptic sheet is extended into a three-dimensional block to make the model similar to the LGN. The three-dimensional arrangement of the postsynaptic block is shown in Figure 6.3. All other elements of the model remained the same.

When the network is presented with the three-dimensional feature vector, postsynaptic units become monocular even with values of z_{pat} that were too low for ocularity to be represented in the two-dimensional network. Figure 6.4 shows the results of presenting the same retinal inputs as used in the two-dimensional experiments into the three-dimensional postsynaptic block of size ($X_{post} = 32, Y_{post} = 32, Z_{post} = 4$), for $z_{pat} = 0.2$ and $z_{pat} = 2.0$.

To analyse the maps, the three-dimensional postsynaptic block has been divided into multiple two-dimensional planes (as illustrated in Figure 6.3). Z-plane n corresponds to all of the units in the plane $Z = n$ of the postsynaptic block. Figure 6.4 shows the results of network development with $z_{pat} = 0.2$. In this case, the top two Z-planes of the network responded to the right eye and the bottom two Z-planes of the network to the left eye. Within each Z-plane, the representation of visual space was complete and topographic. Figure 6.5 shows the topographic maps from the same network when the postsynaptic sheet is divided into multiple X-planes. (X-plane n corresponds to all of the units in the plane $X = n$ of the postsynaptic block.) Each X-plane covers a small part of the visual space, but all of the X-planes taken together cover the entire visual space. Additionally, neighbouring X-planes cover neighbouring parts of the visual space. (X-planes 1 and 32 cover a slightly larger region of visual space than the other planes due to boundary effects.) The X-planes are similar to the projection columns found in the LGN (Sanderson, 1971a).

However, with a bigger value of z_{pat} , such as $z_{pat} = 2.0$, ocularity is the primary map feature and visual space a secondary feature (similar to the two-dimensional network shown in Figure 6.2). Each Z-plane of the network developed in the same manner, although the border between left- and right-eye regions varied systematically through the planes. In Figure 6.4, the left half of each Z-plane is responsive to the right eye, and the right half of each Z-plane is responsive to the left eye. In this case, the topographic plots in both the X- and Y-planes show no correspondence with the LGN projection columns. A sample of some topographic maps in the X- and Y-planes are shown in Figure 6.6. For the X-plane plots we find that the maps for the planes equidistant from the centre (X-planes i and $33 - i$ for all values of $i = 1 \dots 16$) cover almost the same region of the input space. For each pair of X-planes, X-plane i contains units responsive to the right eye, and X-plane $33 - i$ contains units responsive to the left eye. In contrast, each Y-plane topographic plot has folded over on top of itself, so that each part of visual space is covered by two units (one for each eye) in each Y-plane.

6.5 Correspondence with topography and ocularity in the LGN

The cat LGN has a very distinctive three-dimensional shape, as shown in Figure 2.3. It can be approximated however as a three-dimensional block of postsynaptic units, with roughly equal extent in the lateral-medial (X) and anterior-posterior (Y) dimensions, but a much smaller dorsal-ventral (Z) extent. The maps in Figure 6.4 can be compared with the retinotopic and ocularity




Figure 6.3: Three-dimensional arrangement of postsynaptic units in the Kohonen network. Postsynaptic units are positioned uniformly throughout the three-dimensional block. The X dimension of this block corresponds to the medial—lateral dimension of the LGN. Likewise, the Y dimension corresponds to the anterior—posterior dimension, and the Z dimension corresponds to the dorsal—ventral dimension. The figure shows two different two-dimensional slices through the postsynaptic block. The grey slice is a Z-plane which samples all postsynaptic units at the same dorso-ventral position. The black slice is a X-plane which samples all postsynaptic units at the same medio-lateral position. Key: A – anterior. P – posterior. D – dorsal. V – ventral. L – lateral. M – medial.

mappings within the LGN.

First, when $z_{pat} = 0.2$, the plots show a similar retinotopic organisation to that found in the LGN: units at the same (X, Y) position but at different depths (Z) receive input from the same part of visual space. Furthermore, the map has segregated into eye-specific layers (where a layer is used here to mean the same as a Z-plane) in the same way as the LGN. (In the LGN, the contralateral inputs always innervate the top of the LGN. In these experiments there is no preference for the contralateral inputs to go to the top of the LGN, and so the eye providing input to the top Z-plane varies from simulation to simulation. It should be relatively simple to ensure that the contralateral eye always dominates the top half of the postsynaptic block by placing a suitable bias on the initial weights.) The map has automatically oriented itself to represent the input features with the highest variance (x, y) along the largest dimensions of the postsynaptic block (X, Y) . The remaining input feature with the smallest variance, z , is mapped onto the smallest dimension of the postsynaptic block (Z).

This segregation into eye-specific layers is dependent however on the number of Z-planes in the postsynaptic block. When $Z = 4$, the network can easily divide into two halves, so that Z-planes 1 and 2 can respond to one eye, and Z-planes 3 and 4 can respond to the other eye. When $Z = 3$ however, the network fails to completely segregate into separate layers, as shown in Figure 6.7. Monocular units for each eye are found in each Z-plane, although as shown in Table 6.2, there is a tendency for units from the left eye to settle in Z-planes 1 and 2 and for the right eye to settle in Z-planes 2 and 3. This non-complete segregation into eye-specific layers also affects the topography of units within each layer: units responding to different eyes respond to different parts of the visual space. This failure to segregate into eye-specific layers is most likely to be due to the odd number of Z-planes in the network, making it impossible for an equal number of Z-planes to be responsive to each eye.

However, when $z_{pat} = 2.0$, the map develops a different retinotopic and ocular structure to that found in the LGN. For the map in Figure 6.4, the primary feature, ocularity, is mapped along the X dimension of the postsynaptic block, rather than the Z dimension as before. The remaining

inputs, (x and y), map along the X and Y dimensions of the postsynaptic block. The structure of the map within each Z-plane is almost identical, except for a shift in boundary position between different ocularity values. This does not correspond to the topographic and ocularity maps found in the LGN, since both space and ocularity are primary features in the LGN.

Figure	z_{pat}	Z_{post}	Z-plane	$z > 0$ (right eye)		$z < 0$ (left eye)	
				n	mean \pm s.d.	n	mean \pm s.d.
6.1(a)	0.2	1	1	529	0.039 ± 0.029	495	-0.041 ± 0.030
6.1(b)	0.6	1	1	503	0.475 ± 0.162	521	-0.474 ± 0.154
6.1(c)	1.0	1	1	545	0.854 ± 0.301	479	-0.851 ± 0.305
6.2	2.0	1	1	509	1.964 ± 0.177	515	-1.962 ± 0.191
6.7	1.0	3	1	310	0.741 ± 0.306	714	-0.877 ± 0.230
			2	523	0.673 ± 0.347	501	-0.671 ± 0.346
			3	731	0.880 ± 0.227	293	-0.770 ± 0.305
6.4	0.2	4	1	1016	0.180 ± 0.028	8	-0.044 ± 0.026
			2	967	0.122 ± 0.034	57	-0.086 ± 0.041
			3	72	0.094 ± 0.035	952	-0.122 ± 0.032
			4	24	0.040 ± 0.037	1000	-0.178 ± 0.033
6.4	2.0	4	1	554	1.938 ± 0.284	470	-1.893 ± 0.383
			2	541	1.891 ± 0.351	483	-1.937 ± 0.216
			3	479	1.930 ± 0.244	545	-1.890 ± 0.355
			4	468	1.885 ± 0.390	556	-1.947 ± 0.242

Table 6.2: Mean values of z for the two- and three-dimensional Obermayer experiments. The number of units with values of z greater than (right eye) and less than (left eye) zero are shown for each experiment, along with the mean and standard deviation of z . The left-hand column indicates the figure number where the relevant maps are displayed.

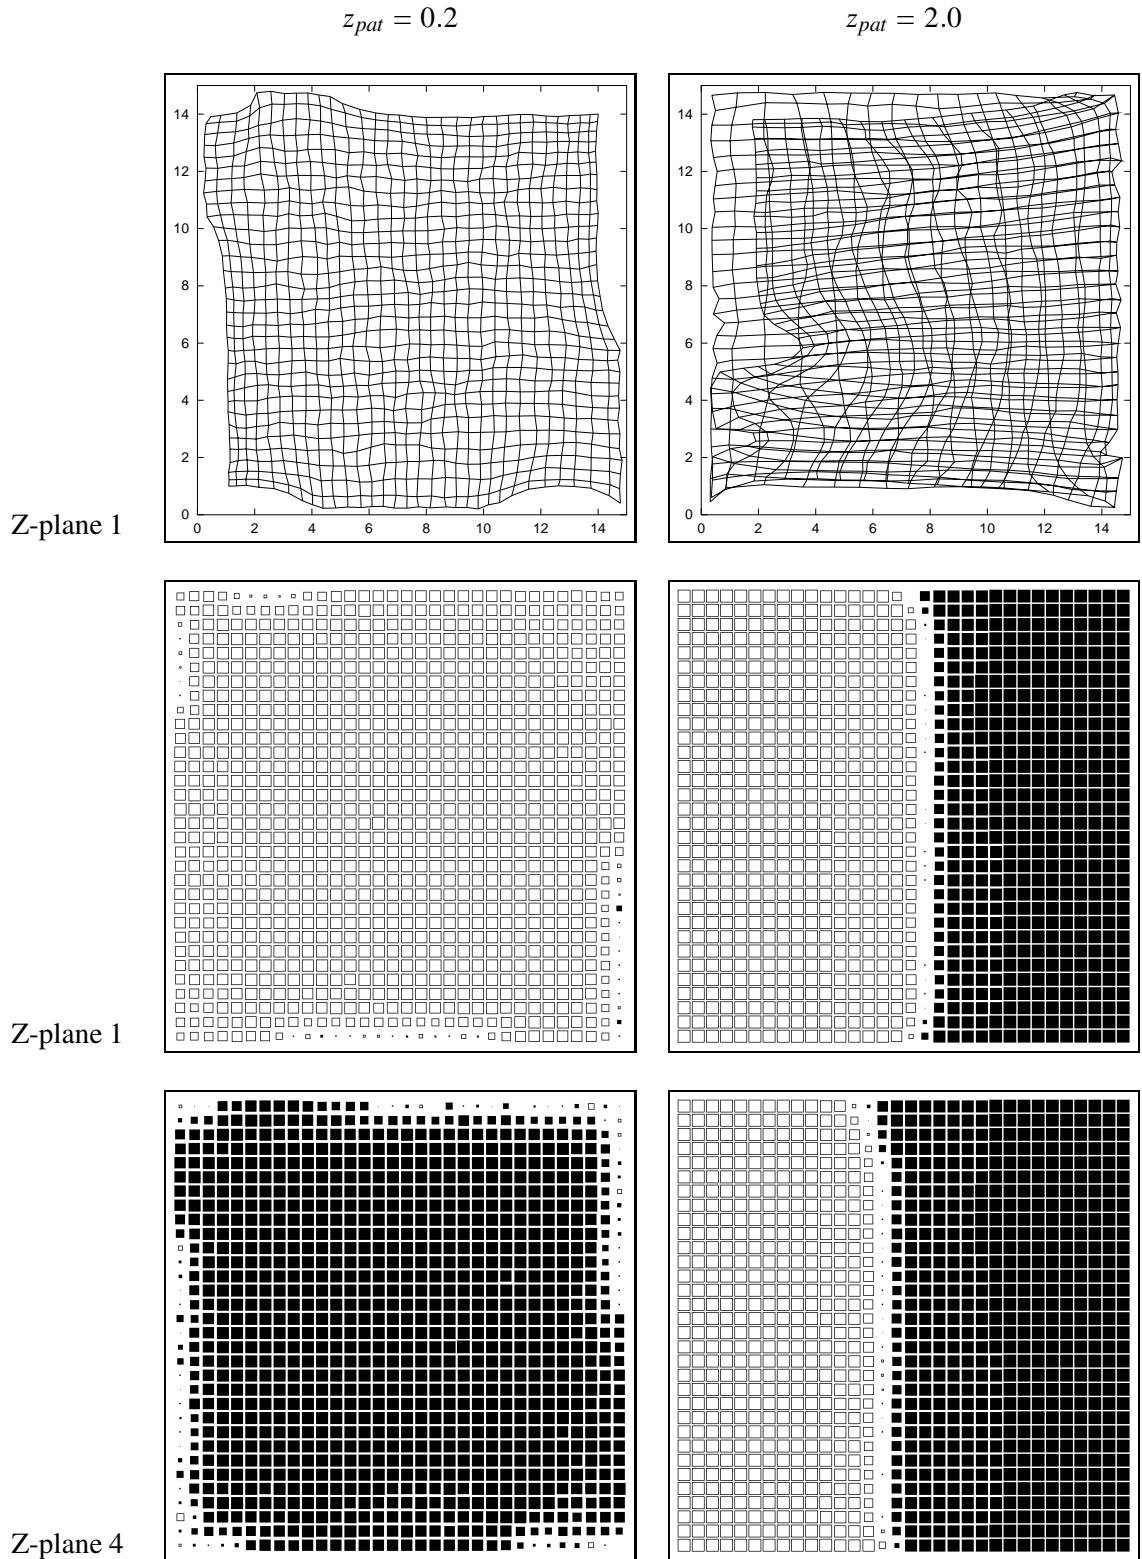




Figure 6.4: Final topography and ocularity plots for three-dimensional feature vector (x, y, z) mapping into a three-dimensional postsynaptic block for $z_{pat} = 0.2$ and 2.0 . The top row shows the topography plots for all of the units in the top Z-plane ($Z = 1$) of the postsynaptic block. The maps for the other Z-planes are similar and so are not shown. The middle and bottom rows show the ocularity plots for units in the top ($Z = 1$) and bottom Z-planes ($Z = 4$) of the network. For $z_{pat} = 0.2$, Z-planes 1 and 2 and Z-planes 3 and 4 were similar, although the central Z-planes (2,3) had slightly fewer monocular units. For $z_{pat} = 2.0$, all Z-planes were similar, although the position of the boundary between left and right eyes gradually moved in the same direction from one Z-plane to the next.




(a) All 32 X-planes

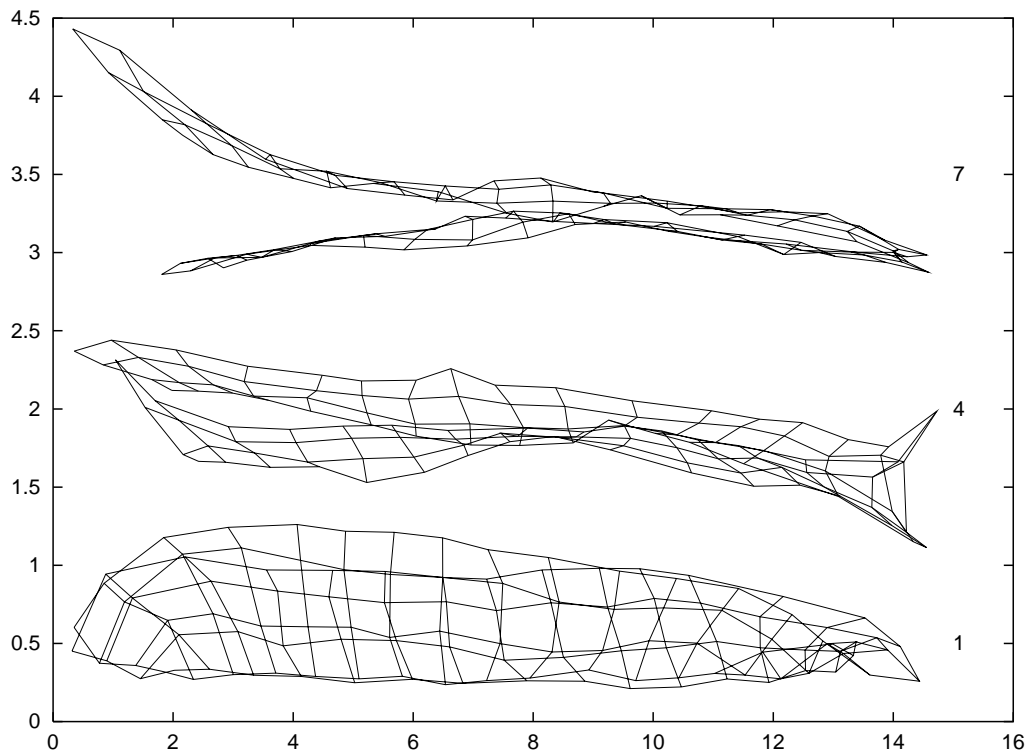


(b) X-planes 1,4 and 7

Figure 6.5: Final topography plots analysed in the X-planes for a three-dimensional network with $z_{pat} = 0.2$ and $Z_{post} = 4$. (a): The 32 topographic maps for each set of 32×4 units in the X-plane $X = i$, for all $i = 1 \dots 32$. The map on the far left corresponds to X-plane $X = 1$, with the remaining maps following in order. (b): Three of the topographic maps in more detail, corresponding to (from left to right) X-planes 1,4 and 7. Each map on its own covers only a small region of the visual space, but together all of the maps completely cover the visual space.

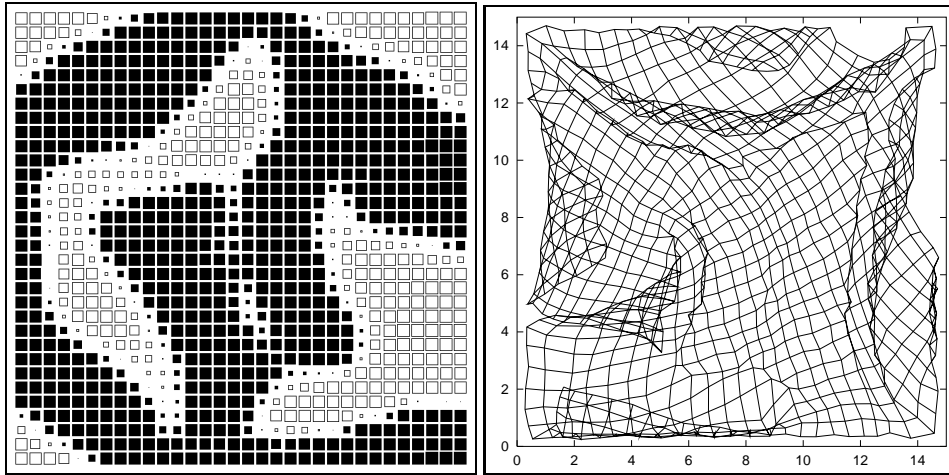
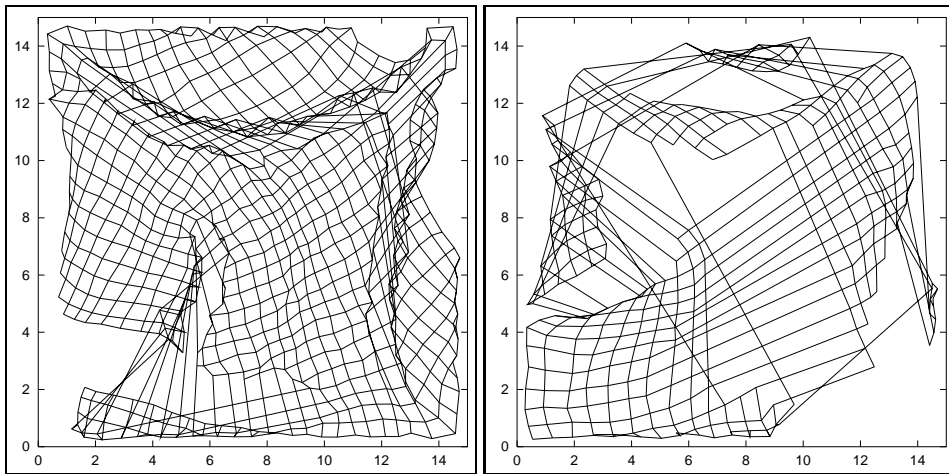


(a) X-planes 1 and 32



(b) Y-planes 1, 4 and 7

Figure 6.6: Final topography plots analysed in X-planes and Y-planes for a three-dimensional network with $z_{pat} = 2.0$ and $Z_{post} = 4$. (a): The topographic maps for X-planes 1 and 32. These are the most medial (solid line) and lateral (dotted line) X-planes in the postsynaptic block. The two maps almost completely cover the same region of the input space. (b): The topographic maps for Y-planes 1, 4 and 7. Each map has folded over on top of itself to represent the same part of visual space twice. For each map, the fold has occurred in the region of visual space corresponding to $X \approx 14$.

(a) $z_{pat} = 1.0$, Z-plane 1

(b) Representation of visual space for each eye in Z-plane 1

Figure 6.7: Final topography and ocularity maps in a three-dimensional network with $z_{pat} = 1.0$ and $Z_{post} = 3$. (a): Topography and ocularity plots for Z-plane 1. (b): Topography plots for Z-plane 1, separated for each eye (left: $z < 0$, right: $z > 0$). At first sight, the topography plot in (a) looks similar to a normal convoluted map when ocularity is present as a secondary feature. Separating the topographic plot according to the ocularity of the postsynaptic unit shows however that the units responsive to each eye represent different parts of visual space.

6.6 Discussion

In this chapter, we have presented three main arguments for the development of ocular dominance stripes. Although these arguments were developed with segregation into ocular stripes in mind, two of them have some connection with ocular segregation in the LGN. In relation to the Keesing model, the first argument, the range of lateral interactions, does not apply since the model does not use lateral connectivity. The model does use a growth rule, but this is only used periodically during development, is excitatory only, and is activity-independent.

The second argument presented for stripe development, the competition between topography and ocularity, is relevant to retinogeniculate development. This argument suggests that stripes form when same-position units and same-eye units compete for the same region of postsynaptic space. The force driving this competition is the strength of the correlation between the presynaptic units. In the case of retinogeniculate development, retinal waves produce no between-eye correlations, and so we would expect same-eye units to cluster together rather than same-position units. In the absence of any bias in the initial weights, it is therefore expected that same-eye units should cluster together. However, as seen in the first set of retinogeniculate simulations (in Figure 4.4 on page 65), the postsynaptic sheet does not divide into two eye-specific regions unless there is some form of initial weight bias to help segregation.

The initial weight bias used in Chapter 4 is justified by the evidence that contralateral axons arrive at the LGN a few days earlier than ipsilateral axons (Shatz, 1983). However, the simulations presented here in Section 6.3.3 show that it is possible for ocularity maps to segregate into two eye-specific regions without any form of initial weight bias. This is probably because the Obermayer model uses the Kohonen rule for weight updates, and this rule adapts the weights of a group of neighbouring postsynaptic units at the same time. In contrast, the Keesing model has weaker interactions between neighbouring postsynaptic units, making global reorganisation of topography and ocularity more difficult.

The third argument for ocular dominance stripe formation is also relevant to the issue of retinogeniculate development. Ocular dominance stripes can form in networks when the dimensionality of features in the input space is larger than the dimensionality of the output space. In the case of cortical development of topography and ocularity, the input space has three features and the output space is only two-dimensional. As long as the variance of the ocularity feature is above a critical value, ocular dominance stripes form. In this chapter, ocularity variance is assumed to increase as the input stimuli become more monocular. In Section 6.3.3 we have replicated the results from (Obermayer et al., 1991) which show that as the ocularity variance increases, stripe width increases. For sufficiently high values of ocularity variance, ocular dominance stripes within a two-dimensional postsynaptic sheet coalesce into two eye-specific regions.

When the Obermayer model is extended so that postsynaptic units are arranged into a three-dimensional block, ocularity stripes are no longer present, as shown in Section 6.4. In this case, the three input features can independently map onto the three dimensions of the postsynaptic block, although the relative variance of each of the features still controls the layout of the map. When the variance of the ocularity feature is small in comparison with the centre of mass features ($z_{pat} = 0.2$), ocularity maps onto the smallest dimension of the postsynaptic block (Z). The other features representing centre of mass are mapped onto the X and Y dimensions of the postsynaptic

block. So, even when the network does not need to perform dimension reduction, the variance of each feature controls how they map into the postsynaptic block.

As the variance of ocularity is increased ($z_{pat} = 2.0$), ocularity maps onto the X dimension of the postsynaptic block, one of the two big dimensions. The centre of mass features are then forced to map onto the remaining dimensions of the postsynaptic block. However, because the value of z for each postsynaptic unit is always $\pm z_{pat}$, the x centre of mass value can also map onto the X dimension of the postsynaptic block, rather than the Z dimension which has eight times fewer units (32 versus 4). In this case, the Z dimension is ignored. Hence, we find that most units (excluding those on the boundary between left and right eye regions) at the same (X,Y) position but at different depths respond to the same eye and the same region of visual space.

In the case of retinal waves driving retinogeniculate development, since the between-eye correlations are very small, the corresponding variance of ocularity should also be quite small (favouring $z_{pat} = 0.2$ rather than $z_{pat} = 2.0$). The Obermayer model therefore predicts that the ocularity feature maps onto the dorso-ventral dimension of the LGN because this is the smallest dimension of the LGN and ocularity is the feature with the smallest variance. In principle this means that the retinal axons could self-organise into eye-specific layers given completely random initial conditions. However, in practice this probably does not happen because of the limited retinal axonal branching. Furthermore, there is no real need for global reorganisation to occur if other mechanisms, such as the earlier arrival of contralateral axons, already impose a global map layout.

Chapter 7

Conclusions

The aim of this thesis has been to investigate the hypothesis that spontaneous waves of activity can guide development of the retinogeniculate pathway. In this chapter we summarise the main results from this thesis. We then give directions for future work and outstanding biological issues that will further help understand the processes underlying retinogeniculate development.

7.1 Discussion of the main results

All of the experimental results mentioned in this section refer to the work presented in Chapters 4, 5 and 6.

7.1.1 Replication of Keesing model results

The first experimental results presented in the thesis replicate the findings of the initial model of retinogeniculate development (Keesing et al., 1992). The original presentation of the model omitted values for most critical parameters as well as the method for generating retinal waves. Replication of their results therefore required finding suitable parameters and making appropriate assumptions about the retinal waves. Once the model was built, it also allowed us to investigate issues that were not previously addressed, such as the formation of projection columns in the LGN.

7.1.2 Normalisation

In contrast to most previous models of ocular dominance, we have shown here that presynaptic normalisation is crucial for ocular dominance development. Postsynaptic normalisation, used in many other models to provide competition between inputs from the two eyes, is redundant in the current model. The presynaptic normalisation should be implemented divisively for retinotopic projection columns to form in the model. Subtractive normalisation can also be used, but only if normalisation is slowly enforced and by placing appropriate upper bounds on individual weights. Postsynaptic normalisation may not be required simply because no positive between-eye correlations are generated by the retinal waves.

7.1.3 Spatio-temporal retinal wave properties

The spatial and temporal properties of spontaneous retinal waves affect development. First, we have shown that retinotopic map formation is reliant on a certain rate of wave generation. Ocular dominance, on the other hand, is more robust to changes in wave generation rate. Examination of the current biological data on temporal wave properties suggests that waves are not generated very often, and so the network needs to cope with long periods of silence between waves. This was achieved in the network by introducing a new weight modification rule, the active-covariance rule, which does not adapt weights when both pre- and postsynaptic activity are below threshold.

Second, the width of retinal waves also affects development. As the waves get wider, we find that the final receptive field size of postsynaptic units also increases, as more neighbouring presynaptic units are correlated. However, this also has the effect of inhibiting ocular segregation since units in different eyes are more likely to be jointly active. Network development is robust to wider waves though as long as narrower waves are used initially to establish the overall map layout. This increase in wave width during development is consistent with current biological data (Wong et al., 1993).

7.1.4 Monocular deprivation

The retinogeniculate model can also account for the results of monocular deprivation which shows that as an eye becomes less active, the number of postsynaptic cells responsive to that eye decreases. At the limit, when an eye never generates any spontaneous activity, no postsynaptic cells respond to that eye (Garraghty et al., 1988). These results are dependent on weakening presynaptic normalisation so that the weights are constrained to lie within, rather than strictly on, the constraint surface. This allows presynaptic units to disconnect from the postsynaptic sheet when their weights are never reinforced.

7.1.5 Segregation of on- and off-centre inputs

Postsynaptic units in the model can refine from initially receiving mixed on- and off-centre inputs to a mature state of receiving inputs from units of only one centre-sign. However, neighbouring postsynaptic units did not consistently organise into sublayers, in contrast with the natural system (Stryker & Zehs, 1983). The model also required on- and off-centre units to be anticorrelated for polarity segregation to occur. Although these correlations are present after the onset of visually-driven activity, they do not appear to be present in the developing retina (Mastrorarde, 1989; Wong & Oakley, 1996). As the anticorrelations between on- and off-centre units in the model were weakened, in line with recordings from the developing retina, polarity segregation was inhibited.

7.1.6 Models of ocular dominance stripe formation

Most computer models of ocular dominance produce stripes of alternating eye input, similar to those found in the visual cortex. These stripes usually form as a result of positive between-eye correlations or some retinotopic bias in the unit weights. We have taken one of these models and extended the model into a three-dimensional output space (Obermayer et al., 1991). When the variance of the ocularity is low in comparison with the variance of the retinotopy features, we find that the ocularity feature maps onto the smallest axis of the postsynaptic block, the dorso-ventral

axis, in a similar fashion to the mapping of ocularity in the LGN. This self organisation can, in principle, account for the overall layout of visual space and ocularity in the LGN. In practice however, factors such as limited axonal branching may prevent such global reorganisation.

7.2 Future work

In this section we consider several directions in which this work could be extended.

7.2.1 Modelling of the retina

The model retina used here is very simplistic and could be improved in several ways. First, we have considered the model retina to represent only a small patch of the natural retina. A larger-scale retina could include non-uniform sampling to investigate if the gradient of cell density in the retina is preserved in the mapping of visual space in the geniculate (Sanderson, 1971b). Second, if the contralateral retina is made larger than the ipsilateral retina, would the geniculate develop a monocular segment for the region of contralateral retina that is missing from the ipsilateral retina? Third, introducing further classification of presynaptic units as either X or Y ganglion cells would also allow us to investigate the issue of clustering of X and Y cells within geniculate layers.

In the work presented here, the retinal waves have been modelled simply as Gaussian functions of distance from a wavefront. We have not considered various other wave properties, such as their velocity and how they interact with other waves travelling across the retina (Feller et al., 1996). Including these other wave properties may introduce finer spatio-temporal correlations that can be exploited by the network.

7.2.2 Scaling up the network

The size of the retinogeniculate model was relatively small to keep the computational demands of the simulations fairly low. Given the encouraging results from the current model, it may be worth scaling up the size of the model on two counts. First, the two models presented in this thesis should be merged so that we use retinal waves (rather than feature vectors) to generate activity for two two-dimensional retinæ innervating a three-dimensional LGN. This new model would be much closer to the natural system than either of the two models investigated in this thesis. Second, the number of units in each layer could be increased, respecting also the relative numbers found in the retina and geniculate during early development. This is not straightforward however, given the ongoing cell death during development and variation in population sizes between animals (Williams et al., 1986; Wong & Hughes, 1987; Spear, Kim, Ahmad, & Tom, 1996). The relative dimensions of the geniculate could also be taken into account when designing the new model (Elgeti, Elgeti, & Fleischhauer, 1976). Making a model which respects the relative numbers of cells in the retina and LGN would allow for comparisons of convergence and divergence with the natural system.

7.2.3 Polarity segregation

The failure to segregate on- and off-centre inputs using correlations similar to those present in the developing retina is likely to be due to the choice of weight modification rule and neighbourhood

update. Given the promising initial results of polarity segregation using a BCM rule with one postsynaptic unit (Lee & Wong, 1996), it is worth investigating performance in the network with both the BCM rule and competitive rules. Competitive rules such as the Obermayer/Goodhill rule use a form of neighbourhood weight update which, unlike the growth rule used in this thesis, should only update the weights of active presynaptic units. This should encourage neighbouring postsynaptic units to develop similar polarity preferences and possibly form polarity sublaminae. This sublamination could also be enhanced by implementing either of the possible forms of polarity bias mentioned in Chapter 5.

7.2.4 Temporal correlations

The spontaneous waves of activity correlate the activity of retinal units both in space and time. In this thesis we have concentrated on the spatial aspects of this correlated activity, without considering the short-term temporal correlations created by the waves moving across the retina. The models presented in this thesis can therefore be regarded as ‘snapshot’ models, since there is no information stored from one iteration to the next. The temporal coincidence of pre- and postsynaptic activity is crucial to weight adaptation (Markram et al., 1997), and exploiting these correlations may be important for segregation of inputs. For example, X and Y ganglion units have considerably different temporal signatures after the onset of vision which could be detected by using temporal information in the weight modification rule (Földiák, 1991; Barrow & Bray, 1992; Wallis, Rolls, & Földiák, 1993).

7.3 Outstanding biological issues

During the course of this thesis, several biological issues have been raised which have not yet been addressed experimentally. Some of these experiments may be possible with current techniques, but it is likely that most have not been performed because there are no suitable techniques currently available. Some topics which deserve future attention are now described.

7.3.1 Retinal waves

More detailed measurements of the patterns of the retinal waves will allow us to investigate the roles of waves in other aspects of retinogeniculate development. For example, are there any differences in firing patterns of X and Y retinal ganglion cells, and does this drive LGN units to become responsive to just X or Y inputs? Measurements of the direction in which waves tend to travel will help to test the hypothesis that wave directions are related to retinal eccentricity and can guide overall retinotopic layout (Trevelyan, Thompson, & Willshaw, 1993). Some of these questions may also be answered as more details of the mechanisms responsible for generation and propagation of the waves are discovered (Wong, 1995; Feller et al., 1996).

7.3.2 Development of topography

Most topographic models show a progressive refinement of both individual receptive field size and overall retinotopic order. How does model refinement correspond to development in the natural system? Most work on this topic has focused on the development of unit responses rather than overall map formation. To answer this question, maps of visual space should be taken at different

times during map development. In the case of retinogeniculate development, projection column maps should also be taken. Pictures of the initial map layout will indicate the amount of order in the system before activity-dependent processes begin to reorganise the pathway. Subsequent maps will then show the amount of reorganisation and refinement. These maps could be generated using retrograde labelling of neighbouring geniculate cells (I. Thompson, personal communication). Producing such maps of visual space in the LGN during the period of retinogeniculate development is complicated however due to the large increase in LGN volume at this time (Elgeti et al., 1976). This may mean that it is possible to create these maps only after development of the retinogeniculate pathway and the LGN has stabilised. Topographic maps of the mature LGN under altered conditions, such as activity blockade and monocular deprivation, would, however, be useful for comparison with both control maps and model predictions.

7.3.3 Development of on- and off-centre units

The segregation of LGN layers into polarity-specific sublaminae is believed to be activity-dependent (Cramer et al., 1996; Cramer & Sur, 1997). There are two outstanding questions related to this segregation. First, why do on-centre cells always go into the dorsal region of a lamina and off-centre cells to the ventral region? This could indicate some form of activity-independent mechanism at work, although it has yet to be found. Second, how does the relative firing rates of on- and off-centre units compare to the number of on- and off-centre geniculate units? Wong and Oakley (1996) report that during P16–24, off-centre cells fire around three to four times more often than on-centre cells. If polarity segregation can be regarded as analogous to ocular segregation, do the results from monocular deprivation also apply, namely that the more often a presynaptic unit fires, the more postsynaptic units connect to it? On this basis, we would expect to find many more off-centre than on-centre geniculate cells. Visual inspection of existing maps however seems to show that the sublayers are approximately the same size (Stryker & Zahs, 1983).

7.3.4 Propagation of waves to other areas

The aim of this thesis has been to see how activity-driven processes can guide development of the retinogeniculate pathway. A central assumption of this work is that these processes not only apply to retinogeniculate development, but also to neural development in general. With reference to the visual system, we would therefore expect that the general principles guiding retinogeniculate development also underlie geniculocortical development. It would be therefore be interesting to see the effects of modifying or blocking retinal wave activity upon the development of cortical cell receptive fields. Given the existence of spontaneous waves in other areas (Wong, 1993), similar blockade experiments could be performed in other sensory modalities. The results of these experiments will help to understand the extent to which neural development is guided by general activity-driven processes.

Bibliography

- Andrade, M. A., & Moràn, F. (1996). Structural study of the development of ocularity domains using a neural network model. *Biological Cybernetics*, *74*, 243–254.
- Angelucci, A., Clasca, F., Bricolo, E., Cramer, K. S., & Sur, M. (1997). Experimentally induced retinal projections to the ferret auditory thalamus: development of clustered eye-specific patterns in a novel target. *Journal of Neuroscience*, *17*, 2040–2055.
- Archer, S. M., Dubin, M. W., & Stark, L. A. (1982). Abnormal development of kitten retinogeniculate connectivity in the absence of action potentials. *Science*, *217*, 743–745.
- Artola, A., Bröcher, S., & Singer, W. (1990). Different voltage-dependent thresholds for the induction of long-term depression and long-term potentiation in slices of the rat visual-cortex. *Nature*, *347*, 69–72.
- Atick, J. J., & Redlich, A. N. (1992). What does the retina know about natural scenes? *Neural Computation*, *4*, 196–210.
- Barrow, H. G. (1987). Learning receptive fields. In *IEEE 1st International Conference on Neural Networks*, IV, pp. 115–121.
- Barrow, H. G., & Bray, A. J. (1992). A model of adaptive development of complex cortical cells. In Aleksander, I., & Taylor, J. (Eds.), *Artificial Neural Networks, 2: Proceedings of the International Conference on Artificial Neural Networks*. North-Holland.
- Bauer, H.-U., Brockmann, D., & Geisel, T. (1997). Analysis of ocular dominance pattern formation in a high-dimensional self-organizing-map model. *Network: Computation in Neural Systems*, *8*, 17–33.
- Berman, N. E. J., Naporn, A., & Payne, B. R. (1987). Shape and arrangement of ‘on’ and ‘off’ modules in cat dorsal lateral geniculate nucleus. In *Society for Neuroscience Abstracts*, Vol. 13, p. 202.
- Berman, N. E. J., & Payne, B. R. (1989). Modular organization of on and off responses in the cat lateral geniculate nucleus. *Neuroscience*, *32*, 721–737.
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, *2*, 32–48.
- Blakemore, C., & Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature*, *228*, 477–478.
- Bliss, T. V., & Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology*, *232*, 331–356.
- Bodnarenko, S. R., & Chalupa, L. M. (1993). Stratification of on and off ganglion cell dendrites depends on glutamate mediated afferent activity in the developing retina. *Nature*, *364*, 144–146.

- Bodnarenko, S. R., Jeyarasasingam, G., & Chalupa, L. M. (1995). Development and regulation of dendritic stratification in retinal ganglion-cells by glutamate-mediated afferent activity. *Journal of Neuroscience*, *15*, 7037–7045.
- Bowling, D. B. (1989a). Timing differences between the light responses of X cells recorded simultaneously in cat lateral geniculate nucleus. *Visual Neuroscience*, *2*, 383–389.
- Bowling, D. B. (1989b). Evidence for sublaminar organization of the cells of origin of the geniculostriate pathway in the cat: a novel method for cortical deposition of HRP. *Brain Research*, *502*, 180–188.
- Bowling, D. B., & Caverhill, J. I. (1989). On/off organization in the cat lateral geniculate nucleus: sublaminae vs columns. *Journal of Comparative Neurology*, *283*, 161–168.
- Bowling, D. B., & Michael, C. R. (1984). Terminal patterns of single, physiologically characterized optic tract fibres in the cat's lateral geniculate nucleus. *Journal of Neuroscience*, *4*, 198–216.
- Bowling, D. B., & Wieniawa-Nakiewicz, E. (1986). The distribution of on- and off- centre X- and Y-like cells in the A layers of the cat's lateral geniculate nucleus. *Journal of Physiology*, *375*, 561–572.
- Casagrande, V. A., & Brunso-Bechtold, J. K. (1985). Development of lamination in the lateral geniculate nucleus: critical factors. In Aslin, R. N. (Ed.), *Advances in Neural and Behavioral Development*, Vol. 1, pp. 33–69. Ablex.
- Cleland, B. G., & Dubin, M. W. (1971). Simultaneous recording of input and output of lateral geniculate neurones. *Nature, New Biology*, *231*, 191–192.
- Cleland, B. G., & Dubin, M. W. (1976). The intrinsic connectivity of the LGN of the cat. *Experimental Brain Research, Suppl. 1*, 493–496.
- Cline, H. T., & Constantine-Paton, M. (1989). NMDA receptor antagonists disrupt the retinotectal topographic map. *Neuron*, *3*, 413–426.
- Constantine-Paton, M., Cline, H. T., & Debski, E. (1990). Patterned activity, synaptic convergence, and the NMDA receptor in developing visual pathways. *Annual Review of Neuroscience*, *13*, 129–154.
- Cottrell, M., & Fort, J. C. (1986). A stochastic model of retinotopy: a self-organizing process. *Biological Cybernetics*, *53*, 405–411.
- Cowan, J. D., & Friedman, A. E. (1991). Studies of a model for the development and regeneration of eye-brain maps. In Lippmann, R. P., Moody, J. E., & Touretzky, D. S. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 3, pp. 3–10. Morgan Kaufmann, San Mateo.
- Cramer, K. S., Angelucci, A., Hahm, J. O., Bogdanov, M. B., & Sur, M. (1996). A role for nitric oxide in the development of the ferret retinogeniculate projection. *Journal of Neuroscience*, *16*, 7995–8004.
- Cramer, K. S., & Sur, M. (1997). Blockade of afferent impulse activity disrupts on/off sublamination in the ferret lateral geniculate nucleus. *Developmental Brain Research*, *98*, 287–290.
- Da Silva Filho, A. C. R. (1992). Investigation of a generalised version of Amari's continuous model for neural networks. DPhil thesis, School of Cognitive and Computing Sciences, Sussex University, UK.

- Dale, H. H. (1935). Pharmacology and nerve endings. *Proceedings of the Royal Society of Medicine*, 28, 319–332.
- Dan, Y., Atick, J. J., & Reid, R. C. (1996). Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *Journal of Neuroscience*, 16, 3351–3362.
- Dayan, P. S., & Goodhill, G. J. (1992). Perturbing Hebbian rules. In Moody, J. E., Hanson, S. J., & Lippmann, R. P. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 4, pp. 19–26. Morgan Kaufmann, San Mateo.
- Dowling, J. E. (1987). *The Retina: An approachable part of the brain*. Harvard University Press.
- Dubin, M. W., Stark, A., & Archer, S. M. (1986). A role for action-potential activity in the development of neuronal connections in the kitten retinogeniculate pathway. *Journal of Neuroscience*, 6, 1021–1036.
- Durbin, R., & Mitchison, G. (1990). A dimension reduction framework for understanding cortical maps. *Nature*, 343, 644–647.
- Durbin, R., & Willshaw, D. (1987). An analogue approach to the travelling salesman problem using an elastic net method. *Nature*, 326, 689–691.
- Eglen, S. J. (1995). Modelling the development of the cat lateral geniculate nucleus with Hebbian learning. Tech. rep. CSRP 383, Cognitive and Computing Sciences, Sussex University.
- Eglen, S. J. (1996). Modelling the prenatal development of the lateral geniculate nucleus. In Silva, F. L., Principe, J. C., & Almeida, L. B. (Eds.), *Frontiers in Artificial Intelligence and Applications*, Vol. 37, pp. 33–41. IOS Press, Amsterdam.
- Elgeti, H., Elgeti, R., & Fleischhauer, K. (1976). Postnatal growth of the dorsal lateral geniculate nucleus of the cat. *Anatomy and Embryology*, 149, 1–13.
- Elliot, T., Howarth, C. I., & Shadbolt, N. R. (1996). Neural computation and statistical mechanics. *Proceedings of the Royal Society of London Series B*, 263, 601–606.
- Elliot, T., & Shadbolt, N. R. (1996). A mathematical model of activity-dependent, anatomical segregation induced by competition for neurotrophic support. *Biological Cybernetics*, 75, 463–470.
- Erwin, E., Obermayer, K., & Schulten, K. (1995). Models of orientation and ocular dominance columns in the visual-cortex — a critical comparison. *Neural Computation*, 7, 425–468.
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47.
- Feller, M. B., Wellis, D. P., Stellwagen, D., Weblin, F. S., & Shatz, C. J. (1996). Requirement for cholinergic synaptic transmission in the propagation of spontaneous retinal waves. *Science*, 272, 1182–1187.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical-cells. *Journal of the Optical Society of America A*, 4, 2379–2394.
- Földiák, P. (1991). Learning invariance from transformation sequences. *Neural Computation*, 3, 194–200.
- Frank, E. (1987). The influence of neuronal activity on patterns of synaptic connections. *Trends in Neuroscience*, 10, 188–190.

- Friedlander, M. J., Frégnac, Y., & Burke, J. P. (1993). Temporal covariance of postsynaptic membrane potential and synaptic input — role in synaptic efficiency in visual cortex. *Progress in Brain Research*, *95*, 207–223.
- Galli, L., & Maffei, L. (1988). Spontaneous impulse activity of rat retinal ganglion cells in prenatal life. *Science*, *242*, 90–91.
- Garraghty, P. E., Shatz, C. J., & Sur, M. (1988). Prenatal disruption of binocular interactions creates novel lamination in the cat's lateral geniculate nucleus. *Visual Neuroscience*, *1*, 93–102.
- Gaze, R. M., & Sharma, S. C. (1970). Axial differences in the reinnervation of the goldfish optic tectum by regenerating optic nerve fibres. *Experimental Brain Research*, *10*, 171–181.
- Goodhill, G. J. (1992). Correlations, competition and optimality: modelling the development of topography and ocular dominance. DPhil thesis, School of Cognitive and Computing Sciences, Sussex University, UK.
- Goodhill, G. J. (1993). Topography and ocular dominance with positive correlations. In Hanson, S. J., Cowan, J. D., & Giles, C. L. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 5, pp. 985–992. Morgan Kaufmann, San Mateo.
- Goodhill, G. J., & Barrow, H. G. (1994). The role of weight normalization in competitive learning. *Neural Computation*, *6*, 255–269.
- Goodhill, G. J., & Löwel, S. (1995). Theory meets experiment — correlated neural activity helps determine ocular dominance column periodicity. *Trends in Neuroscience*, *18*, 437–439.
- Goodhill, G. J., & Willshaw, D. J. (1990). Application of the elastic net algorithm to the formation of ocular dominance stripes. *Network: Computation in Neural Systems*, *4*, 41–59.
- Goodhill, G. J., & Willshaw, D. J. (1994). Elastic net model of ocular dominance — overall stripe pattern and monocular deprivation. *Neural Computation*, *6*, 615–621.
- Goodman, C. S., & Shatz, C. J. (1993). Developmental mechanisms that generate precise patterns of neuronal connectivity. *Cell*, *72*, 77–98.
- Graybiel, A. M. (1975). Anatomical organisation of retinotectal afferents in the cat: an autoradiographic study. *Brain Research*, *96*, 1–23.
- Greiner, J. V., & Weidman, T. A. (1980). Histogenesis of the cat retina. *Experimental Eye Research*, *30*, 439–453.
- Guido, W., Lu, S. M., & Sherman, S. M. (1992). Relative contributions of burst and tonic responses to the receptive-field properties of lateral geniculate neurons in the cat. *Journal of Neurophysiology*, *68*, 2199–2211.
- Guillery, R. W., Lamantia, A. S., Robson, J. A., & Huang, K. (1985). The influence of retinal afferents upon the development of layers in the dorsal lateral geniculate-nucleus of mustelids. *Journal of Neuroscience*, *5*, 1370–1379.
- Hahn, J. O., Langda, R. B., & Sur, M. (1991). Disruption of retino geniculate afferent segregation by antagonists to NMDA receptors. *Nature*, *351*, 568–570.
- Hancock, P. J. B., Smith, L. S., & Phillips, W. A. (1991). A biologically supported error-correcting learning rule. *Neural Computation*, *3*, 201–212.

- Hankin, M., & Lund, R. (1991). How do retinal axons find their targets in the developing brain? *Trends in Neuroscience*, *14*, 224–228.
- Häussler, A. F., & von der Malsburg, C. (1983). Development of retinotopic projections: an analytical treatment. *Journal of Theoretical Neurobiology*, *2*, 47–73.
- Hayes, W. P., & Meyer, R. L. (1988a). Retinotopically inappropriate synapses of subnormal density formed by surgically misdirected optic fibers in goldfish tectum. *Developmental Brain Research*, *38*, 304–312.
- Hayes, W. P., & Meyer, R. L. (1988b). Optic synapse number but not density is constrained during regeneration onto surgically halved tectum in goldfish — HRP-EM evidence that optic fibers compete for fixed numbers of postsynaptic sites on the tectum. *Journal of Comparative Neurology*, *274*, 539–559.
- Hebb, D. O. (1949). *The organization of behaviour*. New York: Wiley.
- Hirsch, H. V. B., & Spinelli, D. N. (1970). Visual experience modifies distribution of horizontally and vertically oriented receptive fields in cats. *Science*, *168*, 869–871.
- Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, *117*, 500–544.
- Hope, R. A., Hammond, B. J., & Gaze, R. M. (1976). The arrow model: retinotectal specificity and map formation in the goldfish visual system. *Proceedings of the Royal Society of London Series B*, *194*, 447–466.
- Hubel, D. H., Levay, S., & Wiesel, T. N. (1975). Mode of termination of retinotectal fibres in macaque monkey: an autoradiographic study. *Brain Research*, *96*, 25–40.
- Hubel, D. H., & Wiesel, T. N. (1961). Integrative action in the cat's lateral geniculate body. *Journal of Physiology*, *155*, 383–398.
- Hubel, D. H., & Wiesel, T. N. (1963). Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *Journal of Neurophysiology*, *26*, 994–1001.
- Hubel, D. H., & Wiesel, T. N. (1972). Laminar and columnar distribution of geniculocortical fibres in macaque monkey. *Journal of Comparative Neurology*, *146*, 421–450.
- Hubel, D. H., & Wiesel, T. N. (1977). Functional architecture of the macaque monkey visual cortex. *Proceedings of the Royal Society of London Series B*, *198*, 1–59.
- Hubel, D. H., Wiesel, T. N., & LeVay, S. (1977). Plasticity of ocular dominance columns in monkey striate cortex. *Philosophical Transactions of the Royal Society of London Series B*, *278*, 377–409.
- Humphrey, A. L., & Weller, R. E. (1988). Functionally distinct groups of X-cells in the lateral geniculate nucleus of the cat. *Journal of Comparative Neurology*, *268*, 429–447.
- Jeffery, G. (1989). Shifting retinal maps in the development of the lateral geniculate-nucleus. *Developmental Brain Research*, *46*, 187–196.
- Johnson, J. K., & Casagrande, V. A. (1993). Prenatal development of axon outgrowth and connectivity in the ferret visual system. *Visual Neuroscience*, *10*, 117–130.
- Jones, E. C. (1985). *The Thalamus*. Plenum Press: New York.

- Kaas, J. H., Guillery, R. W., & Allman, J. M. (1972). Some principles of organization in the dorsal lateral geniculate nucleus. *Brain Behaviour and Evolution*, *6*, 253–299.
- Kaas, J. H., Nelson, R. J., Sur, M., & Lin, C.-S. (1979). Multiple representations of the body within the primary somatosensory cortex of primates. *Science*, *204*, 521–523.
- Kalil, R. E. (1980). A quantitative study of the effects of monocular enucleation and deprivation on cell growth in the lateral geniculate nucleus of the cat. *Journal of Comparative Neurology*, *189*, 438–524.
- Kalil, R. E., Dubin, M. W., Scott, G., & Stark, L. A. (1986). Elimination of action potentials blocks the structural development of retinogeniculate synapses. *Nature*, *323*, 156–158.
- Katz, L. C., & Shatz, C. J. (1996). Synaptic activity and the construction of cortical circuits. *Science*, *274*, 1133–1138.
- Keesing, R., Stork, D. G., & Shatz, C. J. (1992). Retinogeniculate development: the role of competition and correlated retinal activity. In Moody, J. E., Hanson, S. J., & Lippmann, R. P. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 4, pp. 91–97. Morgan Kaufmann, San Mateo.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, *43*, 59–69.
- Kohonen, T. (1988). *Self-Organization and Associative Memory* (third edition). Springer-Verlag.
- Lee, C. W., & Wong, R. O. L. (1996). Developmental patterns of on/off retinal ganglion cell activity lead to segregation of their afferents under a Hebbian synaptic rule. In *Society for Neuroscience Abstracts*, Vol. 22, p. 477.6.
- Lee, D. Y., & Malpeli, J. G. (1994). Global form and singularity — modeling the blind spots role in lateral geniculate morphogenesis. *Science*, *263*, 1292–1294.
- LeVay, S., & Ferster, D. (1979). Proportion of interneurons in the cat's lateral geniculate nucleus. *Brain Research*, *164*, 304–308.
- LeVay, S., Hubel, D. H., & Wiesel, T. N. (1975). The pattern of ocular dominance columns in macaque visual cortex revealed by a reduced silver stain. *Journal of Comparative Neurology*, *159*, 559–576.
- Leventhal, A. G., & Schall, J. D. (1983). Structural basis of orientation sensitivity of cat retinal ganglion-cells. *Journal of Comparative Neurology*, *220*, 465–475.
- Levick, W. R., & Thibos, L. N. (1982). Analysis of orientation bias in cat retina. *Journal of Physiology*, *329*, 243–261.
- Levine, M. W., Cleland, B. G., Mukherjee, P., & Kaplan, E. (1996). Tailoring of variability in the lateral geniculate nucleus of the cat. *Biological Cybernetics*, *75*, 219–227.
- Linden, D. C., Guillery, R. W., & Cucchiari, J. (1981). The dorsal lateral geniculate-nucleus of the normal ferret and its postnatal-development. *Journal of Comparative Neurology*, *203*, 189–211.
- Linsker, R. (1986a). From basic network principles to neural architecture: emergence of spatial-opponent cells. *Proceedings of the National Academy Of Sciences of the U.S.A.*, *83*, 7508–7512.

- Linsker, R. (1986b). From basic network principles to neural architecture: emergence of orientation-selective cells. *Proceedings of the National Academy Of Sciences of the U.S.A.*, 83, 8390–8394.
- Linsker, R. (1986c). From basic network principles to neural architecture: emergence of orientation columns. *Proceedings of the National Academy Of Sciences of the U.S.A.*, 83, 8779–8783.
- Lippe, W. R. (1994). Rhythmic spontaneous activity in the developing avian auditory-system. *Journal of Neuroscience*, 14, 1486–1495.
- Löwel, S. (1994). Ocular dominance column development: strabismus changes the spacing of adjacent columns in cat visual cortex. *Journal of Neuroscience*, 14, 7451–7468.
- Löwel, S., & Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neural activity. *Science*, 255, 209–211.
- MacKay, D. J. C., & Miller, K. D. (1990a). Analysis of Linsker's simulations of Hebbian rules. *Neural Computation*, 2, 173–187.
- MacKay, D. J. C., & Miller, K. D. (1990b). Analysis of Linsker's application of Hebbian rules to linear networks. *Network: Computation in Neural Systems*, 1, 257–297.
- Maffei, L., & Galli, L. (1990). Correlation in the discharges of neighbouring rat retinal ganglion cells during prenatal life. *Proceedings of the National Academy Of Sciences of the U.S.A.*, 87, 2861–2864.
- Markram, H., Lubke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275, 213–215.
- Masland, R. H. (1996). Processing and encoding of visual information in the retina. *Current Opinion in Neurobiology*, 6, 467–474.
- Maslim, J., & Stone, J. (1988). Time course of stratification of the dendritic fields of ganglion cells in the retina of the cat. *Developmental Brain Research*, 44, 87–93.
- Mastrorarde, D. N. (1983a). Correlated firing of cat retinal ganglion-cells. 1. Spontaneously active inputs to X-cell and Y-cell. *Journal of Neurophysiology*, 49, 303–324.
- Mastrorarde, D. N. (1983b). Correlated firing of cat retinal ganglion-cells. 2. Responses of X-cell and Y-cell to single quantal events. *Journal of Neurophysiology*, 49, 325–349.
- Mastrorarde, D. N. (1983c). Interactions between ganglion-cells in cat retina. *Journal of Neurophysiology*, 49, 350–365.
- Mastrorarde, D. N. (1987a). 2 classes of single-input X-cells in cat lateral geniculate-nucleus. 1. Receptive-field properties and classification of cells. *Journal of Neurophysiology*, 57, 357–380.
- Mastrorarde, D. N. (1987b). 2 classes of single-input X-cells in cat lateral geniculate-nucleus. 2. Retinal inputs and the generation of receptive-field properties. *Journal of Neurophysiology*, 57, 381–413.
- Mastrorarde, D. N. (1989). Correlated firing of retinal ganglion cells. *Trends in Neuroscience*, 12, 75–79.
- Meister, M., Wong, R. O. L., Baylor, D. A., & Shatz, C. J. (1991). Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina. *Science*, 252, 939–943.

- Métin, C., & Frost, D. O. (1989). Visual responses of neurons in somatosensory cortex of hamsters with experimentally induced retinal projections to somatosensory thalamus. *Proceedings of the National Academy Of Sciences of the U.S.A.*, *86*, 357–361.
- Meyer, R. L. (1979). Retinotectal projection in goldfish to an inappropriate region with a reversal in polarity. *Science*, *205*, 819–821.
- Miller, K. D. (1994). A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between on- and off-center inputs. *Journal of Neuroscience*, *14*, 409–441.
- Miller, K. D. (1996). Synaptic economics: competition and cooperation in synaptic plasticity. *Neuron*, *17*, 371–374.
- Miller, K. D., Keller, J. B., & Stryker, M. P. (1989). Ocular dominance column development — analysis and simulation. *Science*, *245*, 605–615.
- Miller, K. D., & MacKay, D. J. C. (1992). The role of constraints in Hebbian learning. Tech. rep., CNS Memo 19, California Institute of Technology.
- Miller, K. D., & Mackay, D. J. C. (1994). The role of constraints in Hebbian learning. *Neural Computation*, *6*, 100–126.
- Montague, P. R., Gally, J. A., & Edelman, G. M. (1991). Spatial signalling in the development and function of neural connections. *Cerebral Cortex*, *1*, 199–220.
- Montero, V. M. (1991). A quantitative study of synaptic contacts on interneurons and relay cells of the cat lateral geniculate-nucleus. *Experimental Brain Research*, *86*, 257–270.
- Mooney, R., Madison, D. V., & Shatz, C. J. (1993). Enhancement of transmission at the developing retinogeniculate synapse. *Neuron*, *21*, 815–825.
- Mooney, R., Penn, A. A., Gallego, R., & Shatz, C. J. (1996). Thalamic relay of spontaneous retinal activity prior to vision. *Neuron*, *17*, 863–874.
- Mooney, R., Penn, A. A., & Shatz, C. J. (1995). Periodic synaptic currents in the neonatal LGN are generated by retinal activity. In *Society for Neuroscience Abstracts*, Vol. 21, p. 591.4.
- Morgan, J., & Thompson, I. D. (1993). The segregation of on-center and off-center responses in the lateral geniculate-nucleus of normal and monocularly enucleated ferrets. *Visual Neuroscience*, *10*, 303–311.
- Mumford, D. (1991). On the computational architecture of the neocortex. I. The role of the thalamo-cortical loop. *Biological Cybernetics*, *65*, 135–145.
- Murphy, P. C., & Sillito, A. M. (1987). Corticofugal feedback influences the generation of length tuning in the visual pathway. *Nature*, *329*, 727–729.
- Murphy, P. C., & Sillito, A. M. (1996). Functional-morphology of the feedback pathway from area-17 of the cat visual-cortex to the lateral geniculate-nucleus. *Journal of Neuroscience*, *16*, 1180–1192.
- Nelson, M. E., & Bower, J. M. (1990). Brain maps and parallel computers. *Trends in Neuroscience*, *13*, 403–408.
- Nelson, R., Famiglietti, E. V., & Kolb, H. (1978). Intracellular staining reveals different levels of stratification for on-centre and off-centre ganglion cells in the cat retina. *Journal of Neurophysiology*, *41*, 472–483.

- Norden, J. J., & Constantine-Paton, M. (1994). Dynamics of retinotectal synaptogenesis in normal and 3-eyed frogs: evidence for the postsynaptic regulation of synapse number. *Journal of Comparative Neurology*, *348*, 461–479.
- Obermayer, K., Blasdel, G. G., & Schulten, K. (1991). A neural network model for the formation and for the spatial structure of retinotopic maps, orientation- and ocular dominance columns. In Kohonen, T., Mäkisara, K., Simula, O., & Kangas, J. (Eds.), *Proceedings of the International Conference on Artificial Neural Networks Helsinki*.
- Obermayer, K., Ritter, H., & Schulten, K. (1990). A principle for the formation of the spatial structure of cortical feature maps. *Proceedings of the National Academy Of Sciences of the U.S.A.*, *87*, 8345–8349.
- Obermayer, K., Ritter, H., & Schulten, K. (1991). Development and spatial structure of cortical feature maps: a model study. In Lippmann, R. P., Moody, J. E., & Touretzky, D. S. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 3, pp. 11–17. Morgan Kaufmann, San Mateo.
- Oja, E. (1982). A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, *15*, 267–273.
- Overton, K. J., & Arbib, M. A. (1982). The extended branch-arrow model of the formation of retino-tectal connections. *Biological Cybernetics*, *45*, 157–175.
- Pallas, S. L., & Finlay, B. L. (1991). Compensation for population-size mismatches in the hamster retinotectal system: alterations in the organization of retinal projections. *Visual Neuroscience*, *6*, 271–281.
- Penn, A. A., Gallego, R., Mooney, R., & Shatz, C. J. (1995). Spontaneous retinal inputs drive postsynaptic action potentials in the LGN. In *Society for Neuroscience Abstracts*, Vol. 21, p. 591.5.
- Perrett, D. I., Rolls, E. T., & Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. *Experimental Brain Research*, *47*, 329–342.
- Prestige, M. C., & Willshaw, D. J. (1975). On a role for competition in the formation of patterned neural connexions. *Proceedings of the Royal Society of London Series B*, *190*, 77–98.
- Reiter, H. O., & Stryker, M. P. (1988). Neural plasticity without postsynaptic action-potentials — less-active inputs become dominant when kitten visual cortical-cells are pharmacologically inhibited. *Proceedings of the National Academy Of Sciences of the U.S.A.*, *85*, 3623–3627.
- Rétaux, S., & Harris, W. A. (1996). Engrailed and retinotectal topography. *Trends in Neuroscience*, *19*, 542–546.
- Ritter, H., & Schulten, K. (1988). Convergence properties of Kohonen's topology conserving maps: fluctuations, stability and dimension selection. *Biological Cybernetics*, *60*, 59–71.
- Rochester, N., Holland, J. H., Haibt, L. H., & Duda, W. L. (1956). Tests on a cell assembly theory of the action of the brain using a large digital computer. *IRE Transactions on Information Theory*, *2*, 80–93.
- Roe, A. W., Garraghty, P. E., Esguerra, M., & Sur, M. (1993). Experimentally-induced visual projections to the auditory thalamus in ferrets — evidence for a W cell pathway. *Journal of Comparative Neurology*, *334*, 263–280.

- Sabel, B. A., & Schneider, G. E. (1988). The principle of conservation of total axonal arborizations — massive compensatory sprouting in the hamster subcortical visual-system after early tectal lesions. *Experimental Brain Research*, *73*, 505–518.
- Sanderson, K. J. (1971a). The projection of the visual field to the lateral geniculate and medial interlaminar nuclei in the cat. *Journal of Comparative Neurology*, *143*, 101–118.
- Sanderson, K. J. (1971b). Visual field projection columns and magnification factors in the lateral geniculate nucleus of the cat. *Experimental Brain Research*, *13*, 159–177.
- Saul, A. B., & Humphrey, A. L. (1990). Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate-nucleus. *Journal of Neurophysiology*, *64*, 206–224.
- Schlaggar, B. L., & O’Leary, D. D. M. (1991). Potential of visual-cortex to develop an array of functional units unique to somatosensory cortex. *Science*, *252*, 1556–1560.
- Schmidt, J. T., & Buzzard, M. (1993). Activity-driven sharpening of the retinotectal projection in goldfish — development under stroboscopic illumination prevents sharpening. *Journal of Neurobiology*, *24*, 384–399.
- Schmidt, J. T., & Edwards, D. L. (1983). Activity sharpens the map during the regeneration of the retinotectal projection in goldfish. *Brain Research*, *209*, 29–39.
- Sejnowski, T. J. (1977). Storing covariance with nonlinearly interacting neurons. *Journal of Mathematical Biology*, *4*, 303–321.
- Sejnowski, T. J., Koch, C., & Churchland, P. S. (1988). Computational neuroscience. *Science*, *241*, 1299–1306.
- Sernagor, E., & Grzywacz, N. M. (1996). Influence of spontaneous activity and visual experience on developing retinal receptive-fields. *Current Biology*, *6*, 1503–1508.
- Shatz, C. J. (1977). A comparison of visual pathways in Boston and Midwestern Siamese cats. *Journal of Comparative Neurology*, *171*, 205–228.
- Shatz, C. J. (1983). The prenatal development of the cats retinogeniculate pathway. *Journal of Neuroscience*, *3*, 482–499.
- Shatz, C. J. (1994). Role for spontaneous neural activity in the patterning of connections between retina and LGN during visual system development. *International Journal of Developmental Neuroscience*, *12*, 531–546.
- Shatz, C. J. (1996). Emergence of order in visual-system development. *Proceedings of the National Academy Of Sciences of the U.S.A.*, *93*, 602–608.
- Shatz, C. J., & Kirkwood, P. A. (1984). Prenatal development of functional connections in the cats retinogeniculate pathway. *Journal of Neuroscience*, *4*, 1378–1397.
- Shatz, C. J., & Stryker, M. P. (1978). Ocular dominance in layer IV of the cat’s visual cortex and the effects of monocular deprivation. *Journal of Physiology*, *281*, 267–283.
- Shatz, C. J., & Stryker, M. P. (1988). Prenatal tetrodotoxin infusion blocks segregation of retinogeniculate afferents. *Science*, *242*, 87–89.
- Sherman, S. M. (1985). Development of retinal projections to the cat’s lateral geniculate nucleus. *Trends in Neuroscience*, *8*, 350–355.

- Sherman, S. M., & Guillery, R. W. (1996). Functional-organization of thalamocortical relays. *Journal of Neurophysiology*, *76*, 1367–1395.
- Sherman, S. M., & Koch, C. (1986). The control of retinogeniculate transmission in the mammalian lateral geniculate nucleus. *Experimental Brain Research*, *63*, 1–20.
- Sherman, S. M., & Koch, C. (1990). Thalamus. In Shepherd, G. M. (Ed.), *The synaptic organization of the brain* (3rd edition), chap. 8, pp. 246–278. Oxford University Press, Oxford.
- Shou, T. D., & Leventhal, A. G. (1989). Organized arrangement of orientation-sensitive relay cells in the cat's lateral geniculate nucleus. *Journal of Neuroscience*, *9*, 4287–4302.
- Shou, T. D., Leventhal, A. G., Thompson, K. G., & Zhou, Y. F. (1995). Direction biases of X-type and Y-type retinal ganglion-cells in the cat. *Journal of Neurophysiology*, *73*, 1414–1421.
- Sidman, R. L., & Wislocki, G. B. (1954). Histochemical observations on rods and cones in retinas of vertebrates. *Journal of Histochemistry and Cytochemistry*, *2*, 413–433.
- Sillito, A. M., Jones, H. E., Gerstein, G. L., & West, D. C. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature*, *369*, 479–482.
- Simon, D. K., Prusky, G. T., O'Leary, D. D. M., & Constantine-Paton, M. (1992). N-methyl-d-aspartate receptor antagonists disrupt the formation of a mammalian neural map. *Proceedings of the National Academy Of Sciences of the U.S.A*, *89*, 10593–10597.
- Singer, W. (1977). Control of thalamic transmission by corticofugal and ascending reticular pathways in the visual system. *Physiological Reviews*, *57*, 386–420.
- Singer, W. (1994). A new job for the thalamus. *Nature*, *369*, 444–445.
- Sirosh, J., & Miikkulainen, R. (1994). Cooperative self-organization of afferent and lateral connections in cortical maps. *Biological Cybernetics*, *71*, 65–78.
- Sirosh, J., & Miikkulainen, R. (1995). Ocular dominance and patterned lateral connections in a self-organizing model of the primary visual cortex. In Tesauro, G., Touretzky, D. S., & Leen, T. K. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 7, pp. 109–116. The MIT Press, Cambridge.
- Smetters, D. K., Hahm, J., & Sur, M. (1994). An n-methyl-d-aspartate receptor antagonist does not prevent eye-specific segregation in the ferret retinogeniculate pathway. *Brain Research*, *658*, 168–178.
- Somers, D. C., Nelson, S. C., & Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortex simple cells. *Journal of Neuroscience*, *15*, 5448–5465.
- Spear, P. D., Kim, C. B. Y., Ahmad, A., & Tom, B. W. (1996). Relationship between numbers of retinal ganglion-cells and lateral geniculate neurons in the rhesus-monkey. *Visual Neuroscience*, *13*, 199–203.
- Sperry, R. W. (1963). Chemoaffinity in the orderly growth of nerve fibre patterns and connections. *Proceedings of the National Academy Of Sciences of the U.S.A*, *50*, 703–710.
- Sretavan, D. W., & Shatz, C. J. (1986). Prenatal development of retinal ganglion-cell axons — segregation into eye-specific layers within the cats lateral geniculate-nucleus. *Journal of Neuroscience*, *6*, 234–251.

- Sretavan, D. W., Shatz, C. J., & Stryker, M. P. (1988). Modification of retinal ganglion cell axon morphology by prenatal infusion of tetrodotoxin. *Nature*, *336*, 468–471.
- Stone, J. (1978). The number and distribution of ganglion cells in the cat's retina. *Journal of Comparative Neurology*, *180*, 753–771.
- Stone, J. (1983). *Parallel processing in the visual system: the classification of retinal ganglion cells and its impact on the neurobiology of vision*. Plenum Press, New York.
- Stryker, M. P., & Harris, W. A. (1986). Binocular impulse blockade prevents the formation of ocular dominance columns in cat visual-cortex. *Journal of Neuroscience*, *6*, 2117–2133.
- Stryker, M. P., & Zahs, K. R. (1983). On and off sublaminae in the lateral geniculate nucleus of the ferret. *Journal of Neuroscience*, *3*, 1943–1951.
- Suarez, H., Koch, C., & Douglas, R. (1995). Modeling direction selectivity of simple cells in striate visual-cortex within the framework of the canonical microcircuit. *Journal of Neuroscience*, *15*, 6700–6719.
- Sur, M. (1995). Visual activity and cortical development. In *Cerebral Cortex: Function and Development*, (Lyon, France).
- Sur, M., Esguerra, M., Garraghty, P. E., Kritzer, M. F., & Sherman, S. M. (1987). Morphology of physiologically identified retinogeniculate X- and Y- axons in the cat. *Journal of Neurophysiology*, *58*, 1–32.
- Sur, M., Garraghty, P. E., & Roe, A. W. (1988). Experimentally induced visual projections into auditory thalamus and cortex. *Science*, *242*, 1437–1441.
- Swindale, N. V. (1980). A model for the formation of ocular dominance stripes. *Proceedings of the Royal Society of London Series B*, *208*, 243–264.
- Swindale, N. V. (1981). Rules for pattern formation in mammalian visual cortex. *Trends in Neuroscience*, *4*, 102–104.
- Swindale, N. V. (1996). The development of topography in the visual cortex: a review of models. *Network: Computation in Neural Systems*, *7*, 161–247.
- Thompson, K. G., Zhou, Y., & Leventhal, A. G. (1994). Direction-sensitive X and Y cells within the A laminae of the cat's LGNd. *Visual Neuroscience*, *11*, 927–938.
- Thurlow, G. A., Bowling, D. B., & Cooper, R. M. (1993). ON and OFF activity gradients in the lateral geniculate nucleus of the cat: a combined ¹⁴C 2-deoxyglucose and d,l-2-amino-4-phosphonobutyric acid study. *Visual Neuroscience*, *10*, 1027–1033.
- Trevelyan, A., Thompson, I., & Willshaw, D. (1993). Waves of retinal activity: implications for the development of retinotopicity. In *Society for Neuroscience Abstracts*, Vol. 19, p. 52.
- Tzonev, S., Schulten, K., & Malpeli, J. G. (1995). Morphogenesis of the lateral geniculate nucleus: how singularities affect global structure. In Tesauro, G., Touretzky, D. S., & Leen, T. K. (Eds.), *Advances in Neural Information Processing Systems*, Vol. 7, pp. 133–140. The MIT Press, Cambridge.
- von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striata cortex. *Kybernetik*, *14*, 85–100.
- von der Malsburg, C. (1979). Development of ocularity domains and growth behaviour of axon terminals. *Biological Cybernetics*, *32*, 49–62.

- von der Malsburg, C., & Singer, W. (1988). Principles of cortical network organization. In Rakic, P., & Singer, W. (Eds.), *Neurobiology of Neocortex*, pp. 69–99. John Wiley & Sons Ltd.
- von der Malsburg, C., & Willshaw, D. J. (1976). A mechanism for producing continuous neural mappings: ocularity dominance stripes and ordered retino-tectal projections. *Experimental Brain Research, Supplement 1*, 463–469.
- von der Malsburg, C., & Willshaw, D. J. (1977). How to label nerve cells so that they can interconnect in an ordered fashion. *Proceedings of the National Academy Of Sciences of the U.S.A.*, *74*, 5176–5178.
- Wallis, G., Rolls, E., & Földiák, P. (1993). Learning invariant responses to the natural transformations of objects. In *International Joint Conference on Neural Networks 2*, pp. 1087–1090.
- Wassle, H., Boycott, B. B., & Illing, R. B. (1981a). Morphology and mosaic of on-beta and off-beta cells in the cat retina and some functional considerations. *Proceedings of the Royal Society of London Series B*, *212*, 177–195.
- Wassle, H., Peichl, L., & Boycott, B. B. (1981b). Morphology and topography of on-alpha and off-alpha cells in the cat retina. *Proceedings of the Royal Society of London Series B*, *212*, 157–175.
- Wehmeier, U., Dong, D., Koch, C., & Van Essen, D. (1989). Modeling the mammalian visual system. In Koch, C., & Segev, I. (Eds.), *Methods in Neuronal Modeling*, chap. 10. MIT Press.
- Whitelaw, V. A., & Cowan, J. D. (1981). Specificity and plasticity of retinotectal connections: a computational model. *Journal of Neuroscience*, *1*, 1369–1387.
- Wiesel, T. N., & Hubel, D. H. (1963a). Single cell responses in striate cortex of kittens deprived of vision in one eye. *Journal of Neurophysiology*, *26*, 1003–1017.
- Wiesel, T. N., & Hubel, D. H. (1963b). Effects of visual deprivation on morphology and physiology of cells in the cats lateral geniculate body. *Journal of Neurophysiology*, *26*, 978–993.
- Williams, R. W., Bastiani, M. J., Lia, B., & Chalupa, L. M. (1986). Growth cones, dying axons, and developmental fluctuations in the fiber population of the cat's optic nerve. *Journal of Comparative Neurology*, *246*, 32–69.
- Willshaw, D. J., & Dayan, P. (1990). Optimal plasticity from matrix memories: what goes up must come down. *Neural Computation*, *2*, 85–93.
- Willshaw, D. J., & von der Malsburg, C. (1976). How patterned neural connections can be set up by self-organization. *Proceedings of the Royal Society of London Series B*, *194*, 431–445.
- Willshaw, D. J., & von der Malsburg, C. (1979). A marker induction mechanism for the establishment of ordered neural mappings: its application to the retinotectal problem. *Philosophical Transactions of the Royal Society of London Series B*, *287*, 203–243.
- Wingate, R. J. T., & Thompson, I. D. (1995). The morphological development of mammalian retinal ganglion cells. *Progress in retinal and eye research*, *14*, 413–435.
- Wong, R. O. L. (1993). The role of spatio-temporal firing patterns in neuronal development of sensory systems. *Current Opinion in Neurobiology*, *5*, 595–601.
- Wong, R. O. L. (1995). Use, disuse, and growth of the brain. *Proceedings of the National Academy Of Sciences of the U.S.A.*, *92*, 1797–1799.

- Wong, R. O. L., Chernjavsky, A., Smith, S. J., & Shatz, C. J. (1995). Early functional neural networks in the developing retina. *Nature*, *374*, 716–718.
- Wong, R. O. L., Herrmann, K., & Shatz, C. J. (1991). Remodeling of retinal ganglion-cell dendrites in the absence of action-potential activity. *Journal of Neurobiology*, *22*, 685–697.
- Wong, R. O. L., & Hughes, A. (1987). Developing neuronal populations of the cat retinal ganglion cell layer. *Journal of Comparative Neurology*, *262*, 473–495.
- Wong, R. O. L., Meister, M., & Shatz, C. J. (1993). Transient period of correlated bursting activity during development of the mammalian retina. *Neuron*, *11*, 923–938.
- Wong, R. O. L., & Oakley, D. M. (1996). Changing patterns of spontaneous bursting activity of on and off retinal ganglion cells during development. *Neuron*, *16*, 1087–1095.
- Xiong, M. J., Pallas, S. L., Lim, S., & Finlay, B. L. (1994). Regulation of retinal ganglion-cell axon arbor size by target availability — mechanisms of compression and expansion of the retinotectal projection. *Journal of Comparative Neurology*, *344*, 581–597.
- Yuasa, J., Hirano, S., Yamagata, M., & Noda, M. (1996). Visual projection map specified by topographic expression of transcription factors in the retina. *Nature*, *382*, 632–635.
- Yuste, R., Peinado, A., & Katz, L. C. (1992). Neuronal domains in developing neocortex. *Science*, *257*, 666–669.
- Zahs, K. R., & Stryker, M. P. (1985). The projection of the visual-field onto the lateral geniculate-nucleus of the ferret. *Journal of Comparative Neurology*, *241*, 210–224.

Appendix A

Mathematical details

A.1 The derivation of $\frac{d\mathbf{w}}{dt} = \mathbf{C}\mathbf{w}$ from the covariance rule

Some models of visual system development assume that the covariance rule described by Sejnowski (1977) can be reduced to a rule of the form $\frac{d\mathbf{w}}{dt} = \mathbf{C}\mathbf{w}$ (Linsker, 1986a; Miller et al., 1989). In the descriptions of these models, the details of this derivation are omitted. The derivation is non-trivial, and so is included here since no other full derivation has been published in the literature to the author's knowledge.

Consider a set of n presynaptic units $x_k, k = 1 \dots n$, connected by weights w_k to one postsynaptic unit. The activation of the postsynaptic unit, y , is:

$$y = \sum_k w_k x_k \quad (\text{A.1})$$

The covariance rule (Sejnowski, 1977) for updating weights upon presentation of one input vector is:

$$\begin{aligned} \Delta w_i &= \alpha(y - y_0)(x_i - x_0) \quad (\alpha, x_0, y_0 \text{ are constants}) \quad (\text{A.2}) \\ \Delta w_i &= \alpha y x_i - \alpha y_0 x_i - \alpha y x_0 + \alpha y_0 x_0 \end{aligned}$$

This rule can be averaged over a set of inputs, assuming that the weights do not change during presentation of the inputs:

$$\begin{aligned} \langle \Delta w_i \rangle &= \alpha \langle y x_i \rangle - \alpha y_0 \langle x_i \rangle - \alpha x_0 \langle y \rangle + \alpha y_0 x_0 \quad (\text{A.3}) \\ \langle \Delta w_i \rangle &= \alpha \left\langle \sum_k w_k x_k x_i \right\rangle - \alpha y_0 \langle x_i \rangle - \alpha x_0 \left\langle \sum_k w_k x_k \right\rangle + \alpha y_0 x_0 \quad (\text{Substituting for } y \text{ from A.1}) \\ &= \alpha \sum_k w_k \langle x_i x_k \rangle - \alpha y_0 \langle x_i \rangle - \alpha x_0 \sum_k w_k \langle x_k \rangle + \alpha y_0 x_0 \end{aligned}$$

Assuming that the average value of each input is the same:

$$\langle x_i \rangle = \bar{x}, \quad \forall i \quad (\text{A.4})$$

$$\langle \Delta w_i \rangle = \alpha \sum_k w_k \langle x_i x_k \rangle - \alpha y_0 \bar{x} - \alpha x_0 \bar{x} \sum_k w_k + \alpha y_0 x_0 \quad (\text{A.5})$$

The covariance matrix of the inputs, \mathbf{C} , is defined as:

$$\begin{aligned} \mathbf{C}_{ij} &= \langle (x_i - \bar{x}_i)(x_j - \bar{x}_j) \rangle \\ &= \langle x_i x_j \rangle - \langle x_i \bar{x}_j \rangle - \langle \bar{x}_i x_j \rangle + \langle \bar{x}_i \bar{x}_j \rangle \\ &= \langle x_i x_j \rangle - \bar{x}_j \langle x_i \rangle - \bar{x}_i \langle x_j \rangle + \bar{x}_i \bar{x}_j \\ &= \langle x_i x_j \rangle - \bar{x}_i \bar{x}_j - \bar{x}_i \bar{x}_j + \bar{x}_i \bar{x}_j \\ \mathbf{C}_{ij} &= \langle x_i x_j \rangle - \bar{x}^2 \quad (\text{using A.4}) \\ \langle x_i x_j \rangle &= \mathbf{C}_{ij} + \bar{x}^2 \end{aligned}$$

This value of $\langle x_i x_j \rangle$ can be substituted into equation A.5:

$$\begin{aligned} \langle \Delta w_i \rangle &= \alpha \sum_k w_k (\mathbf{C}_{ik} + \bar{x}^2) - \alpha y_0 \bar{x} - \alpha x_0 \bar{x} \sum_k w_k + \alpha y_0 x_0 \\ &= \alpha \sum_k w_k \mathbf{C}_{ik} + \alpha \sum_k w_k \bar{x}^2 - \alpha y_0 \bar{x} - \alpha x_0 \bar{x} \sum_k w_k + \alpha y_0 x_0 \\ &= \alpha \sum_k w_k \mathbf{C}_{ik} + \sum_k w_k (\bar{x}^2 \alpha - \alpha x_0 \bar{x}) - \alpha y_0 \bar{x} + \alpha y_0 x_0 \\ &= \alpha \sum_k w_k \mathbf{C}_{ik} + \bar{x} \alpha \sum_k w_k (\bar{x} - x_0) - \alpha y_0 (\bar{x} - x_0) \end{aligned}$$

Assuming that the average input activity \bar{x} is equal to x_0 , $\langle \Delta w_i \rangle$ reduces to:

$$\begin{aligned} \langle \Delta w_i \rangle &= \alpha \sum_k w_k \mathbf{C}_{ik} \\ \text{or } \frac{d\mathbf{w}}{dt} &= \mathbf{C}\mathbf{w}, \mathbf{w} = (w_1, w_2, \dots, w_n)^T \end{aligned}$$

Hence, given the two assumptions that weight changes occur on a slower timescale than presentation of inputs and that the average activation of all input units is x_0 , the covariance rule reduces to $\frac{d\mathbf{w}}{dt} = \mathbf{C}\mathbf{w}$. This form of the covariance rule is much simpler because it relies only upon presynaptic, and not postsynaptic, activity levels.

A.2 Why eigenvectors dominate development in correlational-based modification rules

Modification rules of the form $\frac{d\mathbf{w}}{dt} = \mathbf{C}\mathbf{w}$ are often analysed by eigenvector analysis. Let us assume that the eigenvectors of \mathbf{C} are \mathbf{e}_a with eigenvalues λ_a . If \mathbf{C} is real and symmetric, there are n real

eigenvectors. Writing \mathbf{w} in terms of the eigenvectors of \mathbf{C} :

$$\mathbf{w} = \sum_a \mathbf{e}_a w_a \quad \text{where } w_a = \mathbf{e}_a \cdot \mathbf{w} \quad (\text{A.6})$$

$$\frac{d\mathbf{w}}{dt} = \mathbf{C}\mathbf{w} = \mathbf{C} \sum_a \mathbf{e}_a w_a \quad (\text{A.7})$$

$$= \mathbf{C}\mathbf{e}_1 w_1 + \mathbf{C}\mathbf{e}_2 w_2 + \cdots + \mathbf{C}\mathbf{e}_n w_n \quad (\text{A.8})$$

$$\frac{d\mathbf{w}}{dt} = \lambda_1 w_1 \mathbf{e}_1 + \lambda_2 w_2 \mathbf{e}_2 + \cdots + \lambda_n w_n \mathbf{e}_n \quad (\text{A.9})$$

Therefore the rate of growth of the weight vector in the direction of each eigenvector is determined by its eigenvalue. Any component of the weight vector in the direction of an eigenvector with negative eigenvalue is quickly removed. Components of the weight vector in the direction of an eigenvector with positive eigenvalue grows exponentially, with the eigenvector with highest eigenvalue quickly dominating development. Assuming that the maximum eigenvector dominates weight development before any constraint limits are met (such as normalisation), the weight vector will develop to become similar to the maximum eigenvector.

A.3 Implementation of subtractive normalisation

The task of subtractive normalisation is to normalise the set of weights $w_i, i = 1 \dots n$, so that $\sum_{i=1}^n w_i = t$, where t is the target value. If the sum is less than the target before normalisation, it is a simple case of adding $(t - \sum_{i=1}^n w_i)/n$ to each weight (assuming there is no upper limit on individual weights). However, if the sum is bigger than the target, the following iterative procedure must be used to prevent weights from taking negative values:

1. Initialise $\Delta = \sum_{i=1}^n w_i - t, N = n$.
2. Set $\delta = \Delta/N$ and then $\Delta = 0$.
3. For each weight i :

(a) Set the new value of w_i :

$$w_i = \begin{cases} w_i - \delta & \text{if } w_i > 0 \\ 0 & \text{otherwise} \end{cases}$$

(b) If $w_i < 0$, then set $\Delta = \Delta - w_i, w_i = 0, N = N - 1$.

4. Go back to step 2 unless $\Delta = 0$.