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 rigourous 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 monocularity 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.