Chapter 6

A New Competitive Model

6.1 Introduction

So far in this thesis we have studied in some detail two different models of map formation. To recapitulate:

- Miller's model is based on a correlational rule, and has been presented by Miller as a reasonably realistic model: he links some model parameters with biological parameters. Although accounting for ocular dominance stripe formation when correlations within each eye are identical and there are non-positive correlations between the eyes, we have shown that (a) the model is unstable to the addition of perturbative correlations making the two eyes different, and (b) addition of small positive correlations between the eys or small negative correlations within each eye can favour binocular over monocular development. In addition, it has not been shown that the same mechanism can generate topography.
- The elastic net model is based on a more abstract developmental rule, and its parameters are harder to interpret biologically. However, it does form topography and ocular dominance from the same mechanism, and has the virtue that its behaviour can be understood in terms of optimizing an objective function.

Neither of these models, or any of the models we reviewed in chapter 3, has been shown to satisfy *all* of the criteria presented in the introduction. As a reminder, these are as follows: (a) forms topography and ocular dominance from the same mechanisms, (b) is expressed in terms of mechanisms and parameters that can be naturally expressed biologically, (c) develops appropriate mappings in the presence of positive correlations between the two eyes, and (d) produces stripe patterns with similar characteristics to natural stripes.

In this chapter we introduce a model which goes further to simultaneously satisfying these criteria than any model yet presented. This new model consists (as usual) of two presynaptic sheets of units innervating a single postsynaptic sheet of units, and is based on the following principles. At each time step, *distributed* patterns of activity are presented simultaneously in both presynaptic sheets. The postsynaptic units compete for the right to respond to the input pattern. The strongest-responding unit and its neighbours have the weights on their input lines updated. Weights are then normalized by a *subtractive* normalization rule.

We first explain the basic mechanisms of the model, and then present some results, drawing particular attention to the predictions that the model makes regarding stripe width. We then discuss in some detail the important issue of the type of normalization used in the model, and present some analysis of this for a simple case. Finally, we explain how this model relates to some of the other models we have discussed. A preliminary account of this work appeared in [Goodhill 1991], and an earlier version of the model was discussed in [Goodhill 1990].

6.2 Formulation of the model

6.2.1 The framework of the model

The model consists of two two dimensional sheets of input cells (indexed here by r) connected to one two-dimensional sheet of output cells (indexed by c) by fibres with variable synaptic weights w_{cr}. It is assumed that the topography of the retina is essentially unchanged by the LGN on its way to the cortex. In addition, the effects of retinal and LGN processing are taken together, and thus for simplicity we refer to the input layers of the model as being "retinae". It should be understood that the model is formulated at a general enough level to be applicable to both the retinocortical and the retinotectal systems.

Both retina and cortex are arranged in square arrays. All weights and cell activities are positive. Lateral interactions exist in the cortical sheet of a circular center-surround excitation/inhibition form [von der Malsburg 1973], although these are not modeled explicitly. The initial pattern of weights is random apart from a small bias that specifies an orientation for the map (similar to the initial conditions in the Tea Trade Model [Willshaw & von der Malsburg 1979]). At each time step, a pattern of activity is presented by setting the activities a_r of retinal cells. Each cortical cell c calculates its total input x_c according to a linear summation rule:

$$x_c = \sum_r w_{cr} a_r$$

We make the following assumptions to speed calculation at each time step (the significance of these assumptions is discussed later):

- Lateral inhibition in the cortex is sufficient to suppress the activity of all but the cell g with maximum response, and cells in a small neighbourhood of g. Cells close to g have an output that decreases as a function of distance from g.
- The "winning" cell g is the cell that has the largest activity when driven purely by the retina, before taking into account the effect of lateral feedback.
- The input-output function of the cortical cells is monotonic, so that g is simply the cortical cell with the greatest input.

This winning cell and its neighbours have their weights updated at each time-step in a Hebbian manner by adding in a small fraction of the input pattern:

$$w_{cr} = w_{cr} + \alpha a_r s(c, g) \tag{6.1}$$

α is a small positive constant, and s is the function that specifies how the activities of cells c near to g decrease with distance from g. We assume s to be a gaussian function of the Euclidean distance between cells in the cortical sheet. So far these assumptions are similar to those made by Kohonen (see e.g. [Kohonen 1988]). However, an important difference is in the form of the inputs to the model, and the normalization rule used to maintain weights within bounds.

6.2.2 The inputs to the model

Input to the model are random dot patterns with short range spatial correlation introduced by convolution with a blurring function. Patterns were generated by assigning the value 0 or 1 to each pixel in the input array with a fixed probability (usually 50%), and then convolving this array with a gaussian function of standard deviation σ_r. This produces patterns of activity with a range of correlation determined by σ_r. It is important to note that these patterns are *distributed*: all cells in the input layer may be active simultaneously.

6.2.3 Efferent normalization

To prevent weights evolving in accordance with the equations above from growing without bounds, some form of constraint on weight values is needed. Here we assume this constraint to be in the form of a normalization rule: each cortical cell has a fixed amount of "weight substance" to share out amongst the efferent fibres that synapse on it. [Miller 1990(a)] and [Miller & MacKay unpublished] have drawn attention within the framework of a correlational model to the fact that there is a sensitive interaction between the statistics of the inputs, the form of the normalization of efferent fibres contacting a postsynaptic cell, and the receptive fields developed by units in such models. They highlight two ways in which either a linear or squared constraint on synaptic strengths may be enforced: *divisively* (as in e.g. [von der Malsburg 1973]) and *subtractively* (as in e.g. [Miller et al 1989]). In the former (which is more common), the normalization is enforced by dividing each weight by the same amount. In the latter, the normalization is enforced by subtracting the same amount from each weight. As one example of the difference in the behaviour of these two rules, [Miller & MacKay unpublished] focus on ocular dominance segregation. They show that a postsynaptic cell evolving under a correlational learning rule as in [Miller et al 1989] with divisive enforcement will only become dominant for one or the other eye if there exist *anti*correlations between the two eyes. They further argue that this constraint does not exist for a subtractive enforcement rule (but see chapter 4).

Although the correlational framework analysed by Miller and Mackay is somewhat different to the competitive model presented in this chapter, the simulation results we present suggest that a similar conclusion applies to the current model with respect to ocular dominance segregation. The same argument applies to receptive field refinement. With divisive enforcement, cortical unit receptive fields remain distributed over the entire retina for the inputs used here. They do not refine to a small localized region unless retinal cells are anticorrelated with each other (an example is shown below). A form of subtractive enforcement was therefore employed in the current model, as follows.

For each cortical unit, the quantity t is calculated:

$$t = \sum_r w_{cr} - N_c$$

Parameter	Value
Number of retinae	2
Number of retinal units in each dimension	16
Number of cortical units in each dimension	32
Width of bounding box for initial weights as proportion of total width	0.8
Width of uniform noise distribution added to each weight	0.1
Learning rate = α	0.01
Constant for efferent normalization = N_c	10.0
Width of cortical interactions σ_c	1.5
Range of positive part of retinal convolution σ_r	1.5
Parameter controlling correlation between two retinae = h	0.15
Number of input patterns presented	350,000

Table 6.1: Parameters for initial results.

where N_c is the total weight available to each cortical cell (a constant). Define t' to be t divided by the total number of retinal units for which $w_{cr} \neq 0$. Then

$$w_{cr} = \begin{cases} w_{cr} - t' & \text{if } w_{cr} - t' > 0 \\ 0 & \text{otherwise} \end{cases}$$

t is now recalculated: if $t \neq 0$ (i.e. some weights have become zero), divisive normalization is applied:

$$w_{cr} = \frac{N_c w_{cr}}{\sum_r w_{cr}}$$

We now present results of the model. This is followed by a more detailed exposition of some aspects of the algorithm. We then discuss the issue of efferent normalization again in more detail, and present some analysis of this in a simple competitive setting.

6.3 Results

We first study a typical case of the model in action. Later we investigate the effect of various parameters on model results.

6.3.1 Parameter values

For our typical case (figures 6.1, 6.2, 6.3, 6.7(c) and 6.8(c)), the parameter values of table 6.1 were used. Further parameters will be given below as they are discussed. Initial bias in the weights, representing a crude polarity matching machanism, was defined similarly to the Tea Trade model [Willshaw & von der Malsburg 1979]. Correlations within an eye were produced as explained above by convolving with a gaussian: all within-eye correlations were positive, and the two eyes were correlated to a certain degree (the meaning of the parameter h will be explained later). Figure 6.1 shows the development of the map, and figures 6.2, 6.3, 6.7(c) and 6.8(c) show the topography and ocular dominance properties of the final state. The meaning of these pictures is now explained: we start with figure 6.2 and return to figure 6.1 later.

6.3.2 Ocularity

In figure 6.2 cortical units are coloured white or black if they are primarily dominant for the left or right eyes respectively, and grey otherwise. The threshold for a unit to be classified as monocular was taken to be that the sum of a unit's weights for one eye is at least 4 times bigger than the sum of its weights for the other. The development of monocularity of each unit is shown on a continuous scale in figure 6.1(d-f). Here the size of the square for each unit is proportional to the monocularity of the unit, and the colour of the square represents

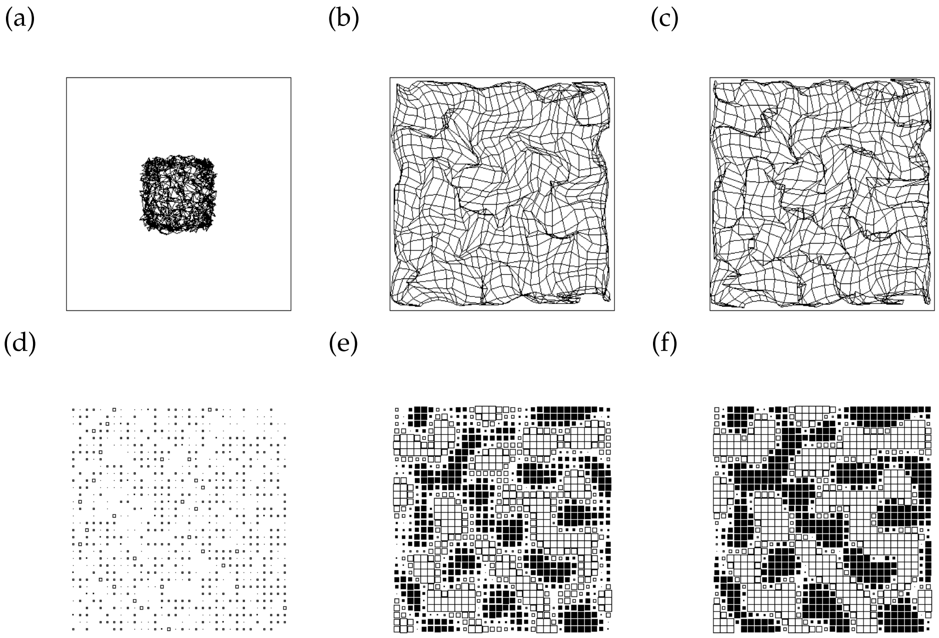

Figure 6.1: Time course of development of the model. (a-c) show the cortical topography averaged over both eyes (see text for further explanation) at 0, 50,000 and 100,000 iterations (i.e. different input patterns presented) respectively. (d-f) show the ocularity of each unit in the cortical array at 0, 50,000 and 100,000 iterations respectively.

for which eye is is dominant. All units are initially binocular, and become gradually more dominant for one eye or the other over time. Longer runs showed that eventually all units become monocular: however the last few units take an extremely long time to saturate and thus runs presented here were terminated when the final properties of the map were clearly visible. Note that all remaining binocular units are on the borders between stripes. Once saturated, weights are stable (see below for further discussion). It can be seen that a pattern of stripes is formed. That is, regions of same-eye dominance are connected into regions that tend to be longer than they are wide, as opposed to blobs.

6.3.3 Topography and its interaction with ocular dominance

In figure 6.2, each monocular cortical unit contains a Hinton diagram of that unit's weights (i.e. efferent connection strengths) for the eye for which it is dominant. This means that there is a 16 by 16 grid within each cortical unit with each grid point representing a retinal unit, and the size of the box at each grid point encodes the strength of the connection between each retinal unit and the cortical unit. For binocular (grey) units, the larger of the two corresponding weights in the two eyes is drawn at each position, coloured white for the left eye and black for the right eye.

It can be seen that neighbouring positions in each eye tend to be represented by neighbouring cortical units, apart from discontinuities across stripe boundaries. For instance, the bottom right corner of the right retina is represented by the bottom right cortical unit, but the bottom left corner of the right retina is represented by cortical unit (3,3) (counting along and up from the bottom left corner of the cortex), since unit (1,1) is representing the left eye.

The topography can be seen more clearly in figures 6.3, 6.7(c) and 6.8(c). In figure 6.3 the receptive fields in the cortex of retinal units are represented. For both eyes, each retinal unit is shown as a square containing a Hinton diagram of its weights to the cortex. Firstly, note that each retinal unit is represented in the cortex. Secondly, we see that the map is broadly topographic: moving across each retina, one observes a continuous shift in the cortical position at which that point in the retina is represented, apart from discontinuities at stripe boundaries. Thirdly, retinal units occasionally arborize to contact two stripes. Fourthly, we see that corresponding units in the two eyes are usually represented close to one another in the cortex.

In figures 6.1(a-c), 6.7(c) and 6.8(c), the centre of mass of the weights of each cortical or retinal unit is represented

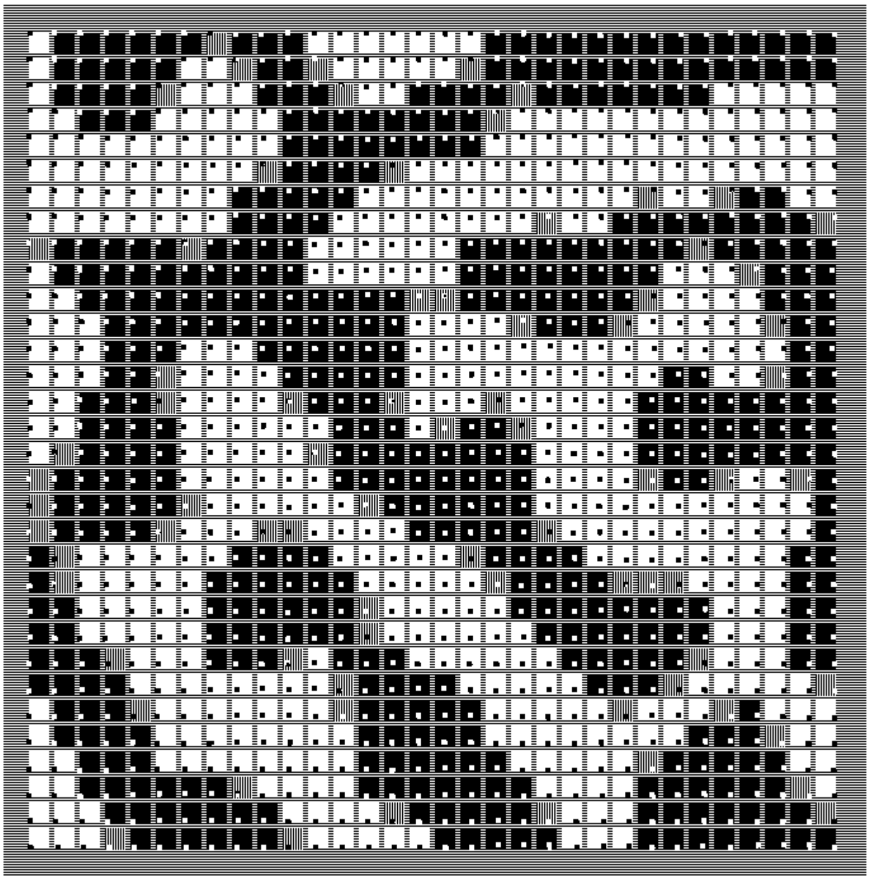

Figure 6.2: The receptive fields of cortical units. Units are coloured white if they are strongly dominant for the left eye, black if they are strongly dominant for the right eye, and grey if they are primarily binocular. Within each unit is a representation of its receptive field (see text for further details).

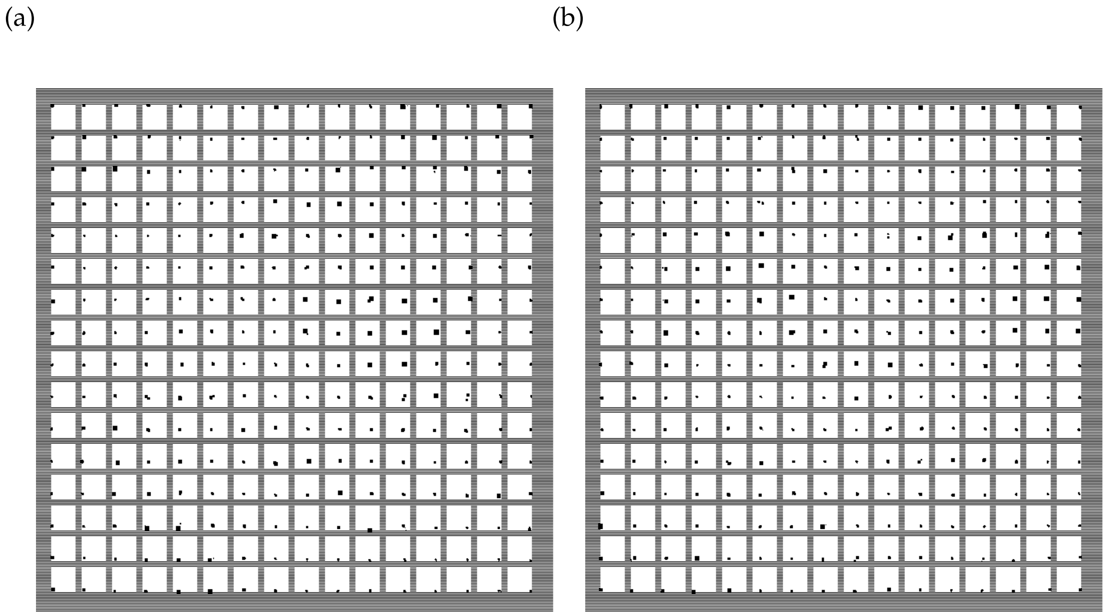

Figure 6.3: The receptive fields in the cortex of retinal units. (a) Left eye. (b) Right eye. Each white box represents a retinal unit, and inside each box is a representation of that unit's connections to the cortex.

as a point in retinal or cortical space respectively, and neighbouring units are connected by lines to form a grid (similar to [Willshaw & von der Malsburg 1976]), except that now we have two sets of pictures, one for each eye. There are three types of picture:

- Retinal topography for each eye (figure 6.8(c)). For each retinal unit, the centre of mass of the weights for each retinal unit is plotted as a point in cortical unit space. It can be seen that for both eyes this map is locally continuous, with gaps (complementary between the two eyes) left for regions of the cortex where the other eye is dominant. However, since we have seen from the retinal weights pictures (figure 6.3) that retinal units can occasionally arborize to contact two stripes, retinal topography pictures are slightly unreliable indicators of topography, since the centre of mass of weight for a retinal unit with contacts evenly balanced between two stripes will be in the middle of the stripe for the other eye.
- Cortical topography for each eye (figure 6.8(c)). For each cortical unit, the centre of mass of the weights for that unit is plotted as a point in retinal unit space. The "holes" in the map arise where cortical units are dominant for the other eye, and are largely complementary between the two eyes. Comparing these picture with the ocularity picture of figure 6.7(c), we see that for each picture topography is continuous within a stripe.
- Cortical topography for both eyes (figures 6.1(a-c), 6.7(c)). Here the centre of mass of weights for each cortical unit is averaged over *both* eyes, imagining the retinae to be lying atop one another. This type of picture reveals where the map is folded to take into account that the cortex is trying to represent both eyes. We see here that discontinuities in terms of folds tend to follow stripe boundaries: first particular positions in one eye are represented, and then the cortex "doubles back" as its ocularity changes in order to represent corresponding positions in the other eye.

Thus the map appears to proceed smoothly for one eye within a stripe, and then double back and "pick up where it left off" to a certain degree for the other eye. This agrees with the experimental data for the natural system (see e.g. [Hubel & Wiesel 1977]).

6.3.4 Development

Various stages of the development of the map are shown in figure 6.1. We see that receptive field refinement and the appearance of topography and the general properties of the stripes occur quite quickly, followed by a

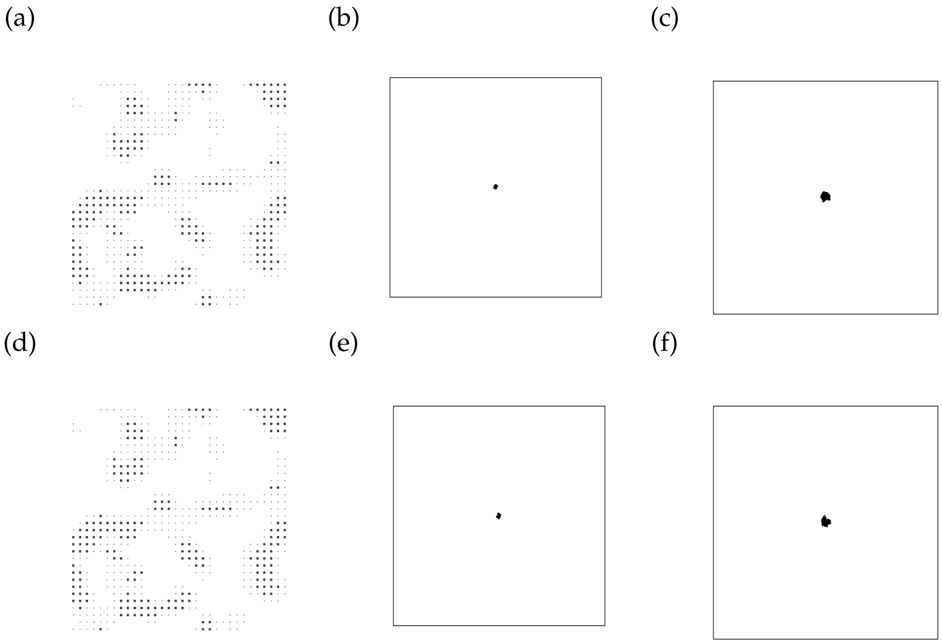

Figure 6.4: Results using a divisive normalization rule: parameters as in table 6.1. (a) Ocularity of cortical units. (b) Retinal topography for left eye. (c) Cortical topography for left eye. (d) Cortical topography averaged over both eyes. (e) Retinal topography for right eye. (f) Cortical topography for right eye.

more gradual process of further refinement.

6.3.5 Divisive normalization

The effect of using divisive rather than subtractive normalization is shown in figure 6.4, for the parameters in table 6.1. It can be seen that receptive fields fail to refine and cortical units remain primarily binocular. This behaviour will be analysed later.

We now examine the effect of changing some of the parameters of the model.

6.3.6 Width of interaction in the cortex

We examined the effect on stripe width of varying the width of cortical interactions. Results are compared in figure 6.5 for a progression of values of σ_c, the width of the gaussian function s in equation 6.1. Also shown in figure 6.5 are the cortical topographies averaged over both eyes, the fourier transforms of the stripe pictures, and the power spectra calculated as described in chapter 5. It can be seen that stripe width increases with the width of the cortical interactions. The retinal and cortical topographies for these simulations are shown in figure 6.6. The parameter value $h = 0.05$ (i.e. a slight amount of correlation between the eyes) was used in these cases.

We also calculated the total length of wire L needed to connect all cortical units representing neighbouring positions within an eye (L_N) or corresponding positions in the two eyes (L_C) as described in chapter 5. Here the cortical unit representing a given retinal unit was taken to be that with the largest weight from that unit. Results are shown in table 6.2. It can be seen that the neighbourhood term tends to decrease as the stripes get wider.

6.3.7 Cross-eye correlations

As we have already argued, it seems highly likely that for animals such as cats and monkeys there exist strong positive correlations between the two eyes. We investigated the behaviour of the competitive model when the two eyes are correlated to varying degrees. Cross-eye correlations were produced in the following way. Once