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 post-synaptic 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 map 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.

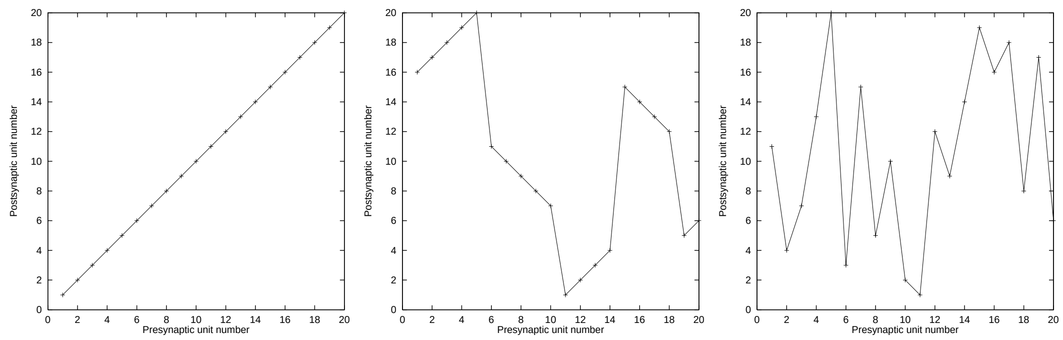

(a) Continuous (b) Patchy (c) Scattered

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).

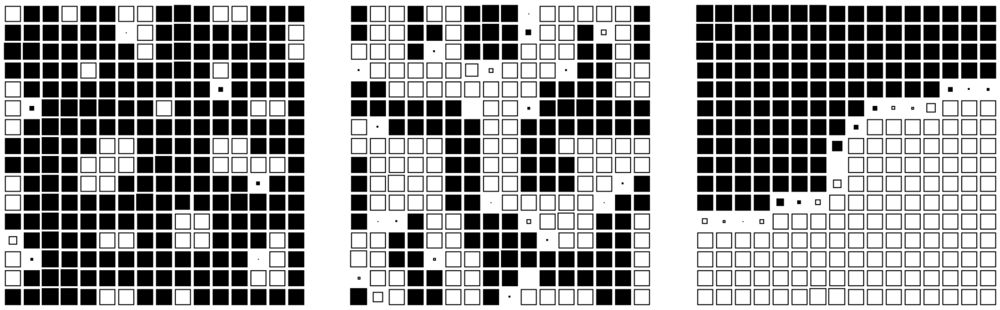

(a) Patches (b) Stripes (c) Layers

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 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.

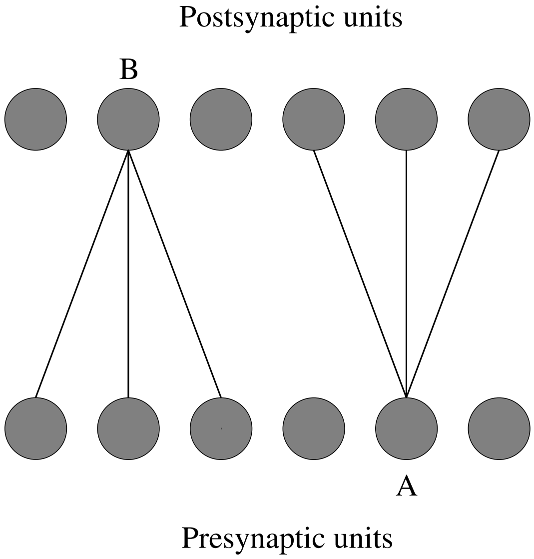

Figure 3.3: The definition of presynaptic and postsynaptic normalisation within a two-layer network. Presynaptic normalisation constrains the sum of the weights from each presynaptic unit (such as unit A). Postsynaptic normalisation constrains the sum of the weights to each postsynaptic unit (such as 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).

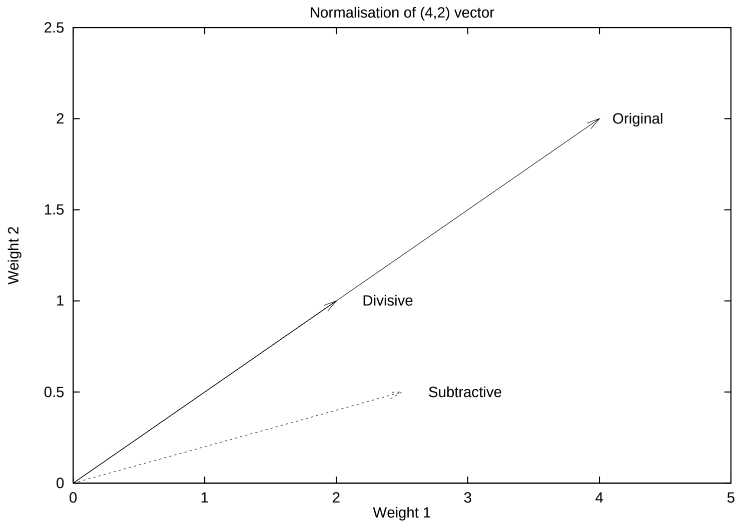

Figure 3.4: A simple comparison of divisive and subtractive normalisation. The original weight vector $(4\ 2)^T$ is to be normalised such that $w_1 + w_2 = 3$. Under divisive normalisation, the vector becomes $(2\ 1)^T$, which lies in the same direction as the original vector. Under subtractive normalisation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

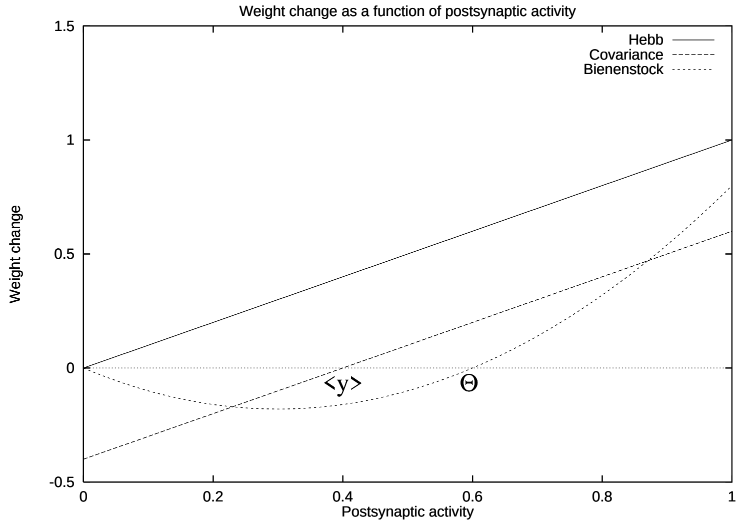

Figure 3.5: Covariance modification rules. The amount of weight change as a function of the postsynaptic activity is plotted for three different correlational rules assuming fixed presynaptic activity ($x_i > \langle x_i \rangle$). The Hebbian rule only specifies weight increments (no negative values 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:

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

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.

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

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.

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

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 $\mathbf{w}_j$. 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:

$$\begin{aligned}
y_j &= \exp\left(\frac{-(\sum_{i=1}^{N_{\text{pre}}}(w_{ij}-x_i)^2)^{0.5}}{2\sigma^2}\right) \\
\Delta w_{ij} &= \alpha \frac{y_j}{\sum_{k\in A} y_k}(x_i - w_{ij}) + \beta k \sum_{k\in N_j}(w_{ik} - w_{ij}) \\
N_j &= \text{set of nearest neighbours of postsynaptic unit } j \\
A &= \text{set of all postsynaptic units}
\end{aligned}$$

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 translation 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 post-synaptic 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.