

Topography And Ocular Dominance: A Model Exploring Positive Correlations

Geoffrey J. Goodhill
University of Edinburgh
Centre for Cognitive Science
2 Buccleuch Place
Edinburgh EH8 9LW
UNITED KINGDOM
gjj@cns.ed.ac.uk

Abstract

The map from eye to brain in vertebrates is topographic, i.e. neighbouring points in the eye map to neighbouring points in the brain. In addition, when two eyes innervate the same target structure, the two sets of fibres segregate to form ocular dominance stripes. Experimental evidence from the frog and goldfish suggests that these two phenomena may be subserved by the same mechanisms. We present a computational model that addresses the formation of both topography and ocular dominance. The model is based on a form of competitive learning with subtractive enforcement of a weight normalization rule. Inputs to the model are distributed patterns of activity presented simultaneously in both eyes. An important aspect of this model is that ocular dominance segregation can occur when the two eyes are *positively correlated*, whereas previous models have tended to assume zero or negative correlations between the eyes. This allows investigation of the dependence of the pattern of stripes on the degree of correlation between the eyes: we find that *increasing correlation leads to narrower stripes*. Experiments are suggested to test this prediction.

1 Introduction

1.1 Striped and Topographic Projections

Projections from cells in one region of the nervous system to those in another are very often arranged so that there is a continuous map of one surface on to the other. In addition, when two equivalent input structures innervate the same target structure it is common for the two projection patterns to be segregated, often interdigitated in stripes or blobs (see for instance [Constantine-Paton & Law 1982]). Perhaps the most thoroughly investigated example of both these phenomena is the mapping from the eye to the brain. In mammals such as cats, monkeys and humans, each retina maps via the lateral geniculate nucleus (LGN) in a topographic fashion across primary visual cortex (V1). Although cortical cells in this region are initially binocularly innervated (i.e. receive connections from both eyes), during the course of development connections from one or the other eye are lost so that in the adult a substantial proportion of cortical cells are monocularly innervated. Looking down on the cortex from above, regions of the cortex dominated by left and right eyes alternate in a pattern somewhat reminiscent of zebra stripes: the so-called "ocular dominance" stripes (e.g. [Hubel & Wiesel 1977, Hendrickson 1985, Horton et al 1990]). In vertebrates such as frogs and goldfish, the projections from the two eyes are normally completely crossed to innervate only the contralateral tectum. However, experimental manipulations can be performed (such as implant of a third eye) that artificially create a region of tectum innervated by both eyes. In this unnatural case the two projections

again segregate to form an interdigitating pattern of stripes (e.g. [Constantine-Paton & Law 1978]; for reviews see [Constantine-Paton 1983, Udin & Fawcett 1988]).

This result has prompted suggestions that there is not a special biological mechanism for forming ocularly dominant cells in visual cortex or optic tectum: rather, monocularity arises naturally as a by-product of mechanisms concerned with map formation [Constantine-Paton & Law 1982, Constantine-Paton 1983]. However, so far no computational model has been proposed that can satisfactorily account for both phenomena simultaneously (see [Goodhill 1991(b)] for discussion). The first motivation for the work presented here was to formulate a model for topographic map formation that can also account for ocular dominance stripe segregation without further assumptions.

1.2 Input Correlations

It is known experimentally that neighbouring cells within the retina have correlated activities when firing spontaneously in the goldfish [Arnett 1978, Ginsburg et al 1984], and in the cat [Mastronade 1989, Meister et al 1991]. Although correlations *between* the two retinæ have not been measured, it seems reasonable to assume that for animals such as cats and monkeys that use stereo vision the following are true:

- Between-eye correlations are approximately zero before eye opening, and for the abnormal cases of dark-rearing, binocular deprivation, and artificial strabismus.
- Between-eye correlations are normally positive after eye opening.

Previous models have tended to address only the cases of zero or negative correlations between the eyes (see Discussion). However, the majority, if not all, of the development of monocular from initially binocular receptive fields in the cat for instance appears during weeks 3 to 6 after birth, whereas eye opening occurs between postnatal days 7 and 10. (see e.g. [Hubel 1988]). The second motivation for the work presented here was therefore to account for the development of cortical cell monocularity in the presence of positive correlations between the eyes.

2 Outline of the Model

The model is formulated at a general enough level to be applicable to both the retinocortical and the retinotectal systems. It consists of either one (monocular case) or two (binocular case) two dimensional sheets of input units (indexed by r) connected to one two-dimensional sheet of output units (indexed by c) by fibres with variable synaptic weights w_{cr} . It is assumed in the retinocortical case 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 retinæ and the output layer as being the cortex. Earlier versions of the model appeared in [Goodhill 1990, Goodhill 1991(a), Goodhill 1993].

Both retina and cortex are arranged in square arrays. All weights and unit 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 (see later). At each time step, a pattern of activity is presented by setting the activities a_r of retinal units. Each cortical unit c calculates its total input x_c according to a linear summation rule:

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

The following assumptions are made to speed calculation at each time step:

- Lateral inhibition in the cortex suppresses the activity of all but the unit g with maximum response, and the units in a small neighbourhood of g . Units close to g have an output that decreases as a function of distance from g .
- The “winning” unit g is the unit 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 units is monotonic, so that g is simply the cortical unit with the greatest input.

This winning unit 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:

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

α is a small positive constant, and s is the function that specifies how the activities of units c near to g decrease with distance from g . s is assumed to be a gaussian function of the Euclidean distance between units in the cortical sheet, with standard deviation σ_c . We refer to s as the ‘‘cortical neighbourhood function’’ (CNF).¹ So far these assumptions are similar to those made by Kohonen [Kohonen 1982, 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.

2.1 Inputs to the Model

Inputs 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 each eye with a fixed probability (50% for the results presented here), and then convolving each eye 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 input units have in general non-zero activity. Between-eye correlations were produced in the following way. Once each eye has been convolved individually with a gaussian function, activity a_j of each unit j in each eye is replaced with $h a_j + (1 - h) a'_j$, where a'_j is the activity of the corresponding unit to j in the other eye, and h specifies the similarity between the two eyes. Thus by varying h it is possible to vary the degree of correlation between the eyes: if $h = 0$ they are uncorrelated, and if $h = 0.5$ they are perfectly correlated (i.e. the pattern of activity is identical in the two eyes).

The correlations existing in the biological system will clearly be rather different to this. However, the simple correlational structure described above aims to capture the *key* features of the biological system: on average, cells in each retina are correlated to an extent that decreases with distance between cells, and corresponding positions in the two eyes are also on the average somewhat correlated.

2.2 Subtractive Normalization

The sum of the weights for each postsynaptic unit is maintained at a constant fixed value, as in for instance [von der Malsburg 1973]. However, whereas in [von der Malsburg 1973] this type of constraint is enforced by dividing each weight by the sum of the weights for that postsynaptic unit (‘‘divisive’’ enforcement), it is enforced here by initially subtracting a constant amount from each weight (‘‘subtractive’’ enforcement), as follows.

For each cortical unit, the quantity t is calculated:

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

where N_c is the total weight available to each cortical unit (a constant), and n_c is 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 enforcement is applied:

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

This is similar to the normalization used in [Miller et al 1989], and is considered further in the Discussion.

¹Later the ‘‘cortical interaction function’’ (CIF) will also be referred to. By this is meant the actual distribution of excitatory and inhibitory lateral connections between cortical units. For a discussion of the relation between the CNF and the CIF see [Kohonen 1988, Dayan 1992].

A constraint was also applied to the sum of the weights for each retinal unit. A rough biological motivation for such a rule is the idea that afferent fibres are supported by a flow of nutrients, and each presynaptic cell generates a fixed quantity of these, which must be divided between the fibres emerging from that cell. A linear normalization constraint for each retinal unit, divisively enforced, was employed in the model at each step after the aforementioned normalization constraint on cortical units:

$$w_{cr} = \frac{N_r w_{cr}}{\sum_c w_{cr}}$$

where N_r is the total weight available to each retinal unit.

2.3 Dead Units, Stability, and Saturation

To avoid the problem of “dead” units (that is, units which do not capture any patterns and hence whose weights do not progress beyond their initial values), we adopt a form of “conscience” mechanism [Hertz et al 1991, page 221]. Here the activity of each cortical unit in response to each input pattern is divided by the number of times it has won the competition so far. This serves to roughly equalize the number of times each cortical unit wins the competition, and thus ensures that dead units do not occur. This could be implemented biologically by a mechanism whereby a cell adjusts its threshold or gain so as to keep its average activity roughly constant. Stability was achieved in the model by introducing a mechanism whereby, once weights reach their maximum or minimum values, they are “frozen” and not allowed to change further (similarly to [Miller et al 1989]). Thus once all connection strengths have saturated the complete map is frozen.

2.4 Initial Bias

There is biological evidence for a crude activity-independent mechanism acting in the early stages of visual map formation that sets up an initial polarity for the map (see e.g. [Udin & Fawcett 1988]). The imposition of such an initial condition is necessary for most models of map formation to generate global order (e.g. [Willshaw & von der Malsburg 1976, Willshaw & von der Malsburg 1979]: for an interesting discussion and analysis see [von der Malsburg & Singer 1988] and [Häussler & von der Malsburg 1983] respectively). Such an initial bias in connectivity is imposed in the present model by generating connection strengths randomly within a fixed range which varies depending on the match in topographic position between each pair of retinal and cortical units, as explained in figure 1. This is similar to the bias used in [Willshaw & von der Malsburg 1979].

3 Results

3.1 Single Eye Case

Figure 2 shows the formation of a topographic map from an initially disordered state (a,d) for the case of a single eye innervating the cortex. Gradually receptive fields refine over the course of development (figure 2(b,e)). Finally, cortical units receive non-zero connections from only a small region of the retina, and neighbouring retinal units are connected to neighbouring parts of the cortex (figure 2(c,f)). Map formation and ocular dominance segregation fail to occur if the normalization constraint for cortical units is enforced divisively rather than subtractively [Goodhill 1991(a), Goodhill 1991(b)].

3.2 Two Eye Case

Typical results for the case of two positively correlated eyes are shown in figure 3 (see table 1 for parameter values). Initial conditions are as before, except that now each unit in *both* eyes has a connection of random strength to every cortical unit. The initial ocular dominance of each cortical unit is thus random (figure 3(a)). A similar overall bias is given to the map for both eyes, so that there is a slight initial registration in the two maps. Development proceeds as in the single eye case, except that now, simultaneous with the refinement of receptive fields, cortical units gradually




Figure 1: The initial bias used to impose an overall orientation for the map. Each retinal sheet was first notionally scaled to the same size as the cortical sheet, and then the normalized distance between each pair of retinal and cortical units calculated. For a unit in the top left corner of the cortical sheet, the distance to the unit in the top left corner of a retinal sheet is zero, and the distance to the bottom right corner is unity. Weights were then chosen from a uniform random distribution in a range determined by the topographic distance between each pair of retinal and cortical units, as shown. The parameter β determines the amount of initial bias: if $\beta = 0$ there is no bias, and if $\beta = 1$ there is a entirely monotonic relationship between topographic distance and size of initial weight. A value of $\beta = 0.5$ was used for the simulations reported here.




Figure 2: Results for a single eye of 16×16 units innervating a cortex of the same size. There are two types of pictures. In (a-c) each large white square represents a cortical unit. Each of these contains a map of the connection strengths between that cortical unit and every retinal unit. The size of a small black square represents the size of that weight, and its position within the white square represents its position in the retina. In (d-f) the “centre of mass” of the weights for each cortical unit is plotted as a point in retinal space. This highlights information about the topography of the map (but does not show the extent to which receptive fields have refined). (a, d) Initial weights. (b, e) After 6000 iterations. (c, f) After 100,000 iterations.

Parameter	Value
Number of retinae	2
Number of retinal units in each dimension	16
Number of cortical units in each dimension	32
Learning rate = α	0.01
Number of iterations	350,000
Amount of initial bias = β	0.5
Constant for cortical unit normalization = N_c	10.0
Constant for retinal unit normalization = N_r	20.0
Width of cortical neighbourhood function = σ_c	1.5
Width of retinal convolution = σ_r	1.5
Parameter controlling correlation between two retinae = h	0.15

Table 1: Parameters for two-eye case.

lose connections from one or the other eye (figure 3(b)). After a large number of input patterns have been presented, cortical units are almost entirely monocular, and units dominant for the left and right eyes are laid out in a pattern of alternating stripes (figure 3(c)). In addition, maps from the two eyes are in register and topographic (figure 3(d-h)). The map of cortical receptive fields (figure 4) confirms that, as described for the natural system in [Hubel & Wiesel 1977], there is a smooth progression of retinal position across a stripe, followed by a doubling back at stripe boundaries for the cortex to “pick up where it left off” in the other eye.

3.3 Width of Cortical Neighbourhood Function

Figure 5 shows the effect of varying the width of the cortical neighbourhood function (CNF) on stripe width. It is apparent that a wider CNF leads to wider stripes. This is analogous to the increase in stripe width seen with increasing width of the cortical interaction function in for instance [Swindale 1980, Miller et al 1989].

3.4 Between-Eye Correlation

An important aspect of the model is that the effect on degree of monocularity and stripe width of the extent of positive correlation between the two eyes can be investigated. This has not been done in previous models (see Discussion). Figure 6 shows a series of results of the model for increasingly strong positive correlations between the two eyes. An effect on stripe width can be observed: stronger between-eye correlations lead to narrower stripes.

4 Discussion

4.1 Relation to Other Models

A number of previously proposed models for topographic map formation have been extended to also account for ocular dominance stripe formation (e.g. [Willshaw & von der Malsburg 1976] extended in [von der Malsburg & Willshaw 1976], [von der Malsburg & Willshaw 1977, Willshaw & von der Malsburg 1979] extended in [von der Malsburg 1979], [Fraser 1980] extended in [Fraser 1985, Fraser & Perkel 1990], [Whitelaw & Cowan 1981] extended in [Cowan & Friedman 1991]). However, these extensions have tended to involve biologically implausible assumptions such as the following:

- There is a global chemical difference between the two eyes (e.g. [von der Malsburg 1979]), a hypothesis which is ruled out most conclusively by the isogenic eye experiment of [Ide et al 1983].

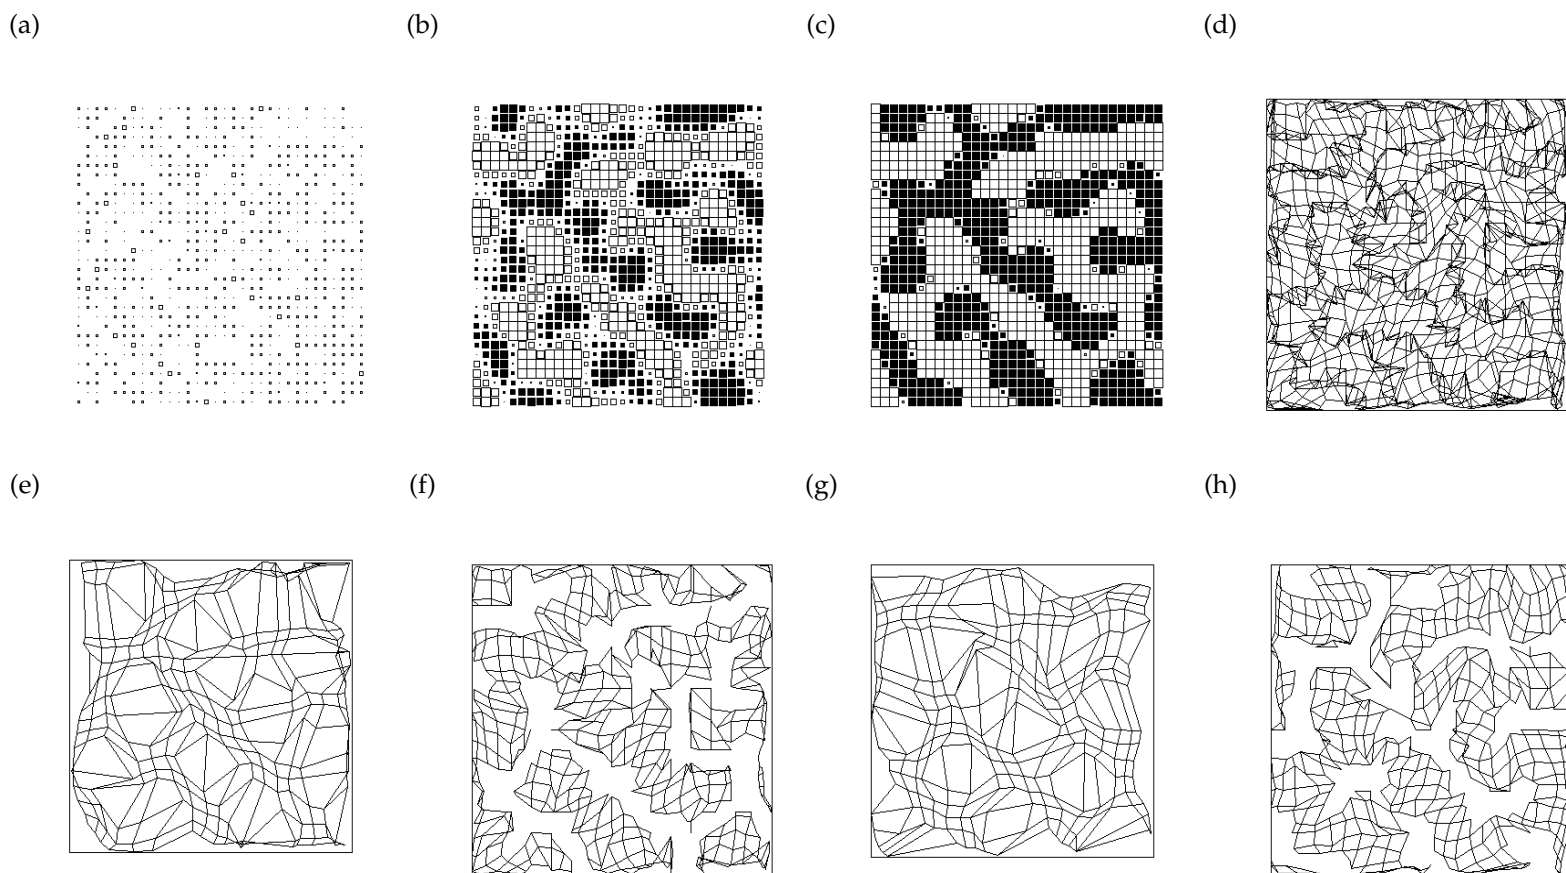


Figure 3: Results for two eyes for parameters of table 1. (a-c) show the ocular dominance of cortical units after 0, 50,000, and 350,000 iterations respectively. Each cortical unit is represented by a square with colour showing the eye for which it is dominant (black for right eye, white for left eye), and size showing the degree to which it is dominant. In (d-h) the centre of mass of the weights of each cortical or retinal unit is represented as a point in retinal or cortical space respectively, and neighbouring units are connected by lines to form a grid (similar to figure 2, except that now we have two sets of pictures, one for each eye). There are three types of picture: (e,g) represent retinal topography for the right and left eyes respectively. For each retina, 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) corresponding to regions of the cortex where the other eye is dominant. (f, h) represent cortical topography for right and left eyes respectively. 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 pictures with the ocularity in (c), we see that for each eye topography is largely continuous within a stripe. (d) represents cortical topography for both eyes. 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 now represents 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.




Figure 4: The receptive fields of cortical units, showing topography and eye preference. 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. “Strongly dominant” was taken to mean that at least 80% of the total weight available to a cortical unit was concentrated in one eye. Within each unit is a representation of its receptive field: 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 or black according to which eye that weight belongs. 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) represents the left retina.

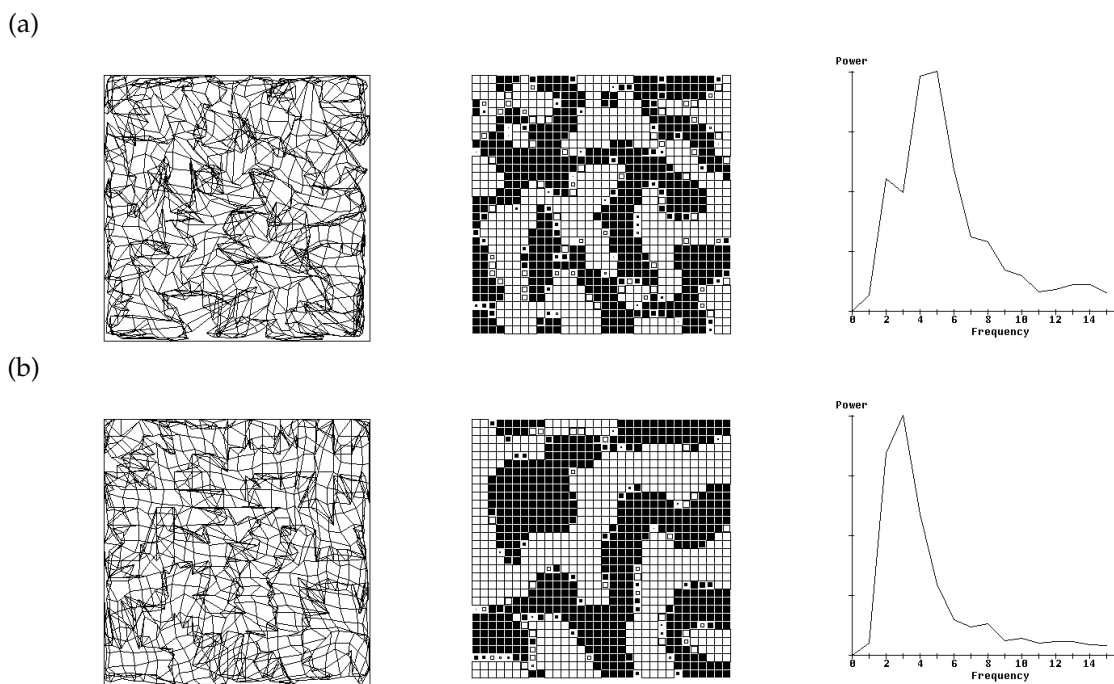


Figure 5: Effect on stripe width of varying the width σ_c of the cortical interaction function. The cortical topography averaged over both eyes, the stripe pattern, and the power spectrum of the fourier transform of the stripe pattern are shown for each case. (a) $\sigma_c = 1.0$. (b) $\sigma_c = 2.0$. Note that stripe width tends to increase with σ_c , as indicated by increasing power at the low-frequency end of the power spectrum of the fourier transform.

- Activity in the two eyes is not correlated (e.g. [Cowan & Friedman 1991]), or even (after being processed by the LGN) negatively correlated [von der Malsburg & Willshaw 1976].

In addition several models have been proposed to account only for ocular dominance, for instance [Swindale 1980]. From the perspective of this paper these are limited since they assume a pre-existing fixed topography, and it has not been shown that they can also account for the development of topography. [Miller et al 1989] present results for the case of non-positive correlations between the two eyes: recent work [Dayan & Goodhill 1992, Goodhill 1991(b), Aguilar-Chongtay 1992] has analysed the breakdown of monocularly in this model with increasing strength of between-eye correlations.

[Obermayer et al 1990, Obermayer et al 1991] presented a model of map formation similar to Kohonen's algorithm, which they applied to the simultaneous development of topography, ocular dominance and orientation columns. Although the results they show are impressive, they are based on the notion of presenting a succession of single points from an abstract "feature space" to a cortical layer of units, rather than development being driven by distributed patterns of activity. This means their model cannot be extended to the case of real-world inputs, unlike the model presented here. In addition, the model presented here uses fixed lateral connections in the cortical sheet, whereas Obermayer et al gradually shrink the extent of lateral connections from an initially large width to zero. There is no biological evidence for such a shrinking process.

Although the "elastic net" model of map and stripe formation of [Goodhill & Willshaw 1990] is based on mechanisms that are rather different to the present model and harder to interpret biologically, it is similar in that it also predicts that stripe width should depend on the form of the input correlations in a similar way. In the elastic net model the cortex attempts to find a "short path" connecting retinal units in an abstract space with distances representing the average correlations between these units. Recent work [Yuille et al 1991] suggests that there may be a mathematical relationship between this model and that of [Miller et al 1989]. A more abstract optimization formulation, addressing the overall pattern of stripes in cats and monkeys in terms of boundary constraints, has recently been proposed in [Jones et al 1991].

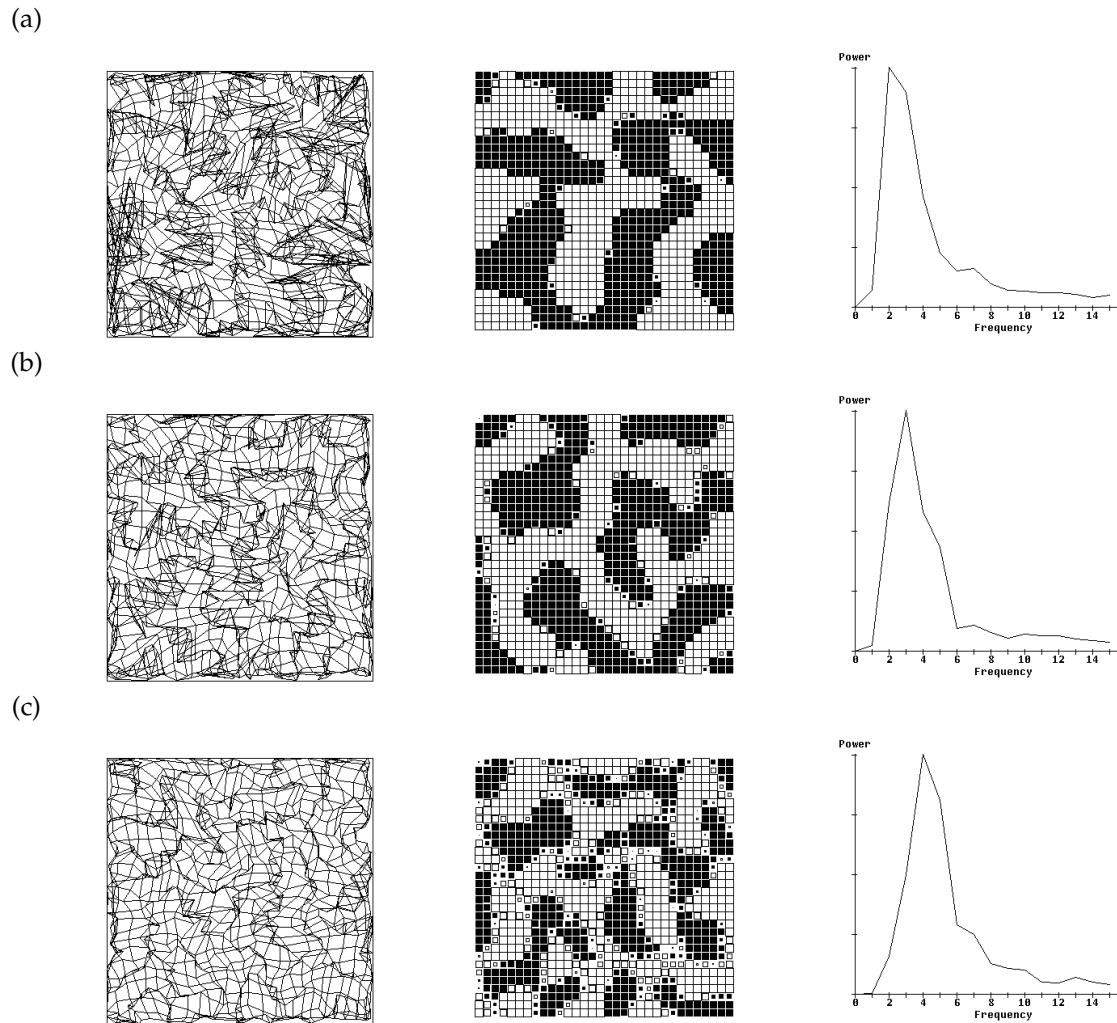


Figure 6: Effect on stripe width of varying h , which controls the degree of correlation between the two eyes. The cortical topography averaged over both eyes, the stripe pattern, and the power spectrum of the fourier transform are shown for each case. (a) $h = 0.0$. (b) $h = 0.1$. (c) $h = 0.2$. Note that stripe width tends to decrease as h increases (indicated by increasingly more power appearing in the high frequency end of the power spectrum of the fourier transform), and that the topography becomes smoother. It can also be seen that as h increases (i.e. correlations between the two eyes become stronger), an increasing proportion of cortical units fail to become fully monocular by the time the simulation was terminated.

4.2 Normalization

An important aspect of the model is that it uses subtractive rather than the more common divisive enforcement of the normalization constraint for cortical units. With divisive enforcement receptive fields fail to refine and segregation into ocular dominance stripes does not occur. An intuitive explanation of this behaviour (see [Miller & MacKay 1992]) is that divisive enforcement leads to a *graded* response of cortical units to the correlations in their inputs: since all input units are correlated to some degree, all weights remain non-zero. Subtractive enforcement however is “less sympathetic” and causes all weights to reach their minimum values (zero in this case), except for the few best-correlated inputs. An analogy is that of taxation: divisive normalization corresponds to income tax, where everybody pays the same percentage of their earnings, whereas subtractive normalization corresponds to a “poll tax”, where everybody pays the same fixed sum.

A mathematical analysis of the effect of these normalization rules has been carried out for a very simple competitive learning case in [Goodhill 1991(b), Goodhill & Barrow 1993]. Here weight vectors evolving in the positive quadrant of a two dimensional space are considered: one axis representing the left eye and the other axis the right eye. It is shown that under divisive enforcement of the normalization constraint weight vectors become evenly distributed through the space, whereas under subtractive enforcement weight vectors saturate at one or other axis, corresponding to ocular dominance segregation.

Competitive learning rules are in general hard to analyse, due to the strong non-linearity of picking the winner. This problem does not arise in Hebbian rules with linear weight development operators, for instance the “correlational” rule of [Miller et al 1989]. The issue of subtractive as compared to divisive enforcement is extensively analysed and discussed in this case for general dimension in [Miller & MacKay 1992]. This paper also considers ways in which such normalization rules may be implemented biologically.

4.3 Stripe Width

The model presented here makes a different prediction from most other models as to what determines ocular dominance stripe width. In for instance [von der Malsburg & Willshaw 1976, Swindale 1980] stripe width is set primarily by the extent of the cortical interaction function: longer-range excitation between cortical units leads to wider stripes. In these models the degree to which (positive) correlations between the eyes affect stripe width has not been investigated: segregation can fail to occur for these models in this case. An intuitive argument can be given to explain why correlations might be expected to have an effect, as follows.

Suppose that cortical cells are locally interconnected by excitatory lateral connections. Given a Hebb-type learning rule, the most stable state will be when neighbouring cortical cells receive inputs that are highly correlated. Cells within each retina will tend on average to be correlated to an extent that decreases with distance. This accounts for the overall topography of the map from the two eyes. However, for animals with stereo vision, corresponding regions of the two retinæ will also in general be correlated. Neighbouring cortical cells thus have competing tendencies to connect with neighbouring cells in the same retina and corresponding cells in the other retina, and the pattern of stripes is a compromise between these two tendencies. The stronger the correlation between the two eyes, the more often the latter tendency will win, leading to narrower stripes.

The form and extent of the cortical interaction function (CIF) also clearly plays a role in setting the stripe width. If the CIF is long-range and purely excitatory, then stripe width will be determined by the correlations between eyes, as described above. However, if the CIF has an excitatory center and an inhibitory surround (as implicitly assumed in this model, and explicitly included in some other models e.g. [Swindale 1980, Miller et al 1989]), then it will not be favourable for an individual stripe to be wider than the width of the excitatory center (assuming within-eye correlations are positive). The CIF in this case provides an upper bound on stripe width.

4.4 Experimental Predictions

It has been predicted that stripe width could be affected by the strength of the correlations between the two eyes. Perhaps the simplest experiment to test this prediction would be to look for changes in stripe width in the cat after artificially induced strabismus, which severely reduces the correlations between the two eyes. Although the

effect of strabismus on the degree of monocularity of cortical cells has been extensively investigated (e.g. [Hubel & Wiesel 1965]), the effect on stripe width has not been examined (David Hubel, personal communication). The model presented here predicts that reduced between-eye correlations should lead to wider stripes. One exception to this would be where, in the normal case, stripe width is already at the upper bound allowed by the cortical interaction function, in which case stripe width will be unchanged.

Acknowledgements

This work was funded by an SERC postgraduate studentship (at Sussex University), and a Joint Councils Initiative in CogSci/HCI postdoctoral training fellowship. I thank Harry Barrow for advice relating to this work, and also David Willshaw and David Price for useful comments on earlier drafts of this paper. I am grateful to Joe Levy for a brief loan of his computing resources.

Bibliography

- Aguilar-Chongtay, R. (1992). The role of correlations in the development of the visual system. Unpublished M.Sc. thesis, University of Edinburgh.
- Arnett, D.W. (1978). Statistical dependence between neighbouring retinal ganglion cells in goldfish. *Exp. Brain. Res.*, **32**, 49-53.
- Constantine-Paton, M. (1983). Position and proximity in the development of maps and stripes. *Trends Neurosci.* **6**, 32-36.
- Constantine-Paton, M. & Law, M.I. (1978). Eye-specific termination bands in tecta of three-eyed frogs. *Science*, **202**, 639-641.
- Constantine-Paton, M. & Law, M.I. (1982). The development of maps and stripes in the brain. *Sci. Am.*, **247**, 54-62.
- Cowan, J.D. & Friedman, A.E. (1991). Studies of a model for the development and regeneration of eye-brain maps. In D.S. Touretzky, ed, *Advances in Neural Information Processing Systems*, **III**, 3-10.
- Dayan, P.S. (1993). Arbitrary elastic topologies and ocular dominance. *Neural Computation*, **5**, 392-401.
- Dayan, P.S. & Goodhill, G.J. (1992) Perturbing Hebbian rules. *Advances in Neural Information Processing Systems*, **4**, 19-26, eds. J.E. Moody, S.J. Hanson and R.P. Lippman, Morgan Kaufmann, San Mateo, CA.
- Fraser, S.E. (1980). A differential adhesion approach to the patterning of neural connections. *Dev. Biol.*, **79**, 453-464.
- Fraser, S.E. (1985). Cell interactions involved in neural patterning. In *Molecular Bases Of Neural Development*, eds. G.M. Edelman, W.E. Gall & W.M. Cowan, 481-507, Wiley, New York.
- Fraser, S.E. & Perkel, D.H. (1990). Competitive and positional cues in the patterning of nerve connections. *Jou. of Neurobiol.*, **21**, 51-72.
- Ginsburg, K.S., Johnsen, J.A. & Levine, M.W. (1984). Common noise in the firing of neighbouring ganglion cells in goldfish retina. *Jou. Physiol.*, **351**, 433-450.
- Goodhill, G.J. (1990). The development of topography and ocular dominance. In Touretzky, D.S., Elman, J.L., Sejnowski, T.J. & Hinton, G.E. (eds.) *Proceedings of the 1990 Connectionist Models Summer School*, San Mateo, CA: Morgan Kaufmann, 338-349.
- Goodhill, G.J. (1991a). Topography and ocular dominance can arise from distributed patterns of activity. *International Joint Conference on Neural Networks, Seattle, July 1991*, **II**, 623-627.
- Goodhill, G.J. (1991b). *Correlations, Competition and Optimality: Modelling the Development of Topography and Ocular Dominance*. PhD Thesis, Sussex University.

- Goodhill, G.J. (1993). Topography and ocular dominance with positive correlations. *Advances in Neural Information Processing Systems*, **5**, eds. C.L. Giles, S.J. Hanson and J.D. Cowan, Morgan Kaufmann, San Mateo, CA.
- Goodhill, G.J. & Barrow, H.G. (1993). The role of weight normalization in competitive learning. Manuscript submitted to *Neural Computation*.
- Goodhill, G.J. & Willshaw, D.J. (1990). Application of the elastic net algorithm to the formation of ocular dominance stripes. *Network*, **1**, 41-59.
- Häussler, A.F. & Malsburg, C. von der (1983). Development of retinotopic projections: an analytical treatment. *Jou. Theoret. Neurobiol.*, **2**, 47-73.
- Hendrickson, A.E. (1985). Dots, stripes and columns in monkey visual cortex. *Trends Neurosci.*, **8**, 406-410.
- Hertz, J., Krogh A. & Palmer, R.G. (1991). Introduction to the Theory of Neural Computation. Lecture notes in the Santa Fe Institute Studies in the sciences of complexity: Addison Wesley.
- Horton, J.C., Dagi, L.R., McCrane, E.P. and de Monasterio, F.M. (1990). Arrangement of ocular dominance columns in human visual cortex. *Arch-Ophthalmol.*, **108**, 1025-1031.
- Hubel, D.H. (1988). *Eye, Brain, and Vision*. Scientific American Library, W.H. Freeman and Company, New York.
- Hubel, D.H. & Wiesel, T.N. (1965). Binocular interaction in striate cortex of kittens reared with artificial squint. *Journal of Neurophysiology*, **28**, 1041-1059.
- Hubel, D.H. & Wiesel, T.N. (1977). Functional architecture of the macaque monkey visual cortex. *Proc. R. Soc. Lond. B*, **198**, 1-59.
- Ide, C.F., Fraser, S.E. & Meyer, R.L. (1983). Eye dominance columns formed by an isogenic double-nasal frog eye. *Science*, **221**, 293-295.
- Jones, D.G., Van Sluyters, R.C. & Murphy, K.M. (1991). A computational model for the overall pattern of ocular dominance. *J. Neurosci.*, **11**, 3794-3808.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biol. Cybern.*, **43**, 59-69.
- Kohonen, T. (1988). *Self-organization and associative memory* (3rd Edition). Springer, Berlin.
- Malsburg, C. von der (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, **14**, 85-100.
- Malsburg, C. von der (1979). Development of ocularity domains and growth behaviour of axon terminals. *Biol. Cybern.*, **32**, 49-62.
- Malsburg, C. von der & Singer, W. (1988). Principles of cortical network organization. In *Neurobiology of Neocortex*, eds. P. Rakic & W. Singer, John Wiley and Sons, 69-99.
- Malsburg, C. von der & Willshaw, D.J. (1976). A mechanism for producing continuous neural mappings: ocularity dominance stripes and ordered retino-tectal projections. *Exp. Brain. Res. Supplementum 1*, 463-469.
- Malsburg, C. von der & Willshaw, D.J. (1977). How to label nerve cells so that they can interconnect in an ordered fashion. *Proc. Nat. Acad. Sci., U.S.A.*, **74**, 5176-5178.
- Mastronade, D.N. (1989). Correlated firing of retinal ganglion cells. *Trends Neurosci.*, **12**, 2, 75-80.
- Meister, M., Wong, R.O.L., Baylor, D.A. & Shatz, C.J. (1991). Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina. *Science*, **252**, 939-943.
- Miller, K.D., Keller, J.B. & Stryker, M.P. (1989). Ocular dominance column development: Analysis and simulation. *Science*, **245**, 605-615.
- Miller, K.D. & MacKay, D.J.C. (1992). The role of constraints in Hebbian learning. Caltech Computation and Neural Systems Memo 19; Also *Neural Computation*, to appear.

- Obermayer, K., Ritter, H. & Schulten, K. (1990). A principle for the formation of the spatial structure of cortical feature maps. *Proc. Nat. Acad. Sci., U.S.A.*, **87**, 8345-8349.
- Obermayer, K., Ritter, H. & Schulten, K. (1991). Development and spatial structure of cortical feature maps: a model study. In *Neural information Processing Systems 3*, eds. R.P. Lippmann, J. Moody & D.S. Touretzky, Morgan Kaufmann, CA.
- Swindale, N.V. (1980). A model for the formation of ocular dominance stripes. *Proc. R. Soc. Lond. B*, **208**, 243-264.
- Udin, S.B. & Fawcett, J.W. (1988). Formation of topographic maps. *Ann. Rev. Neurosci.*, **11**, 289-327.
- Whitelaw, V.A. & Cowan, J.D. (1981). Specificity and plasticity of retinotectal connections: a computational model. *Jou. Neurosci.*, **1**, 1369-1387.
- Willshaw, D.J. & Malsburg, C. von der (1976). How patterned neural connections can be set up by self-organization. *Proc. R. Soc. Lond. B*, **194**, 431-445.
- Willshaw, D.J. & Malsburg, C. von der (1979). A marker induction mechanism for the establishment of ordered neural mappings: its application to the retinotectal problem. *Phil. Trans. Roy. Soc. B*, **287**, 203-243.
- Yuille, A.L., Kolodny, J.A. & Lee, C.W. (1991). Dimension reduction, generalized deformable models and the development of ocularity and orientation. *International Joint Conference on Neural Networks, Seattle, July 1991*, **II**, 597-602. See also Harvard Robotics Laboratory Technical Report 91-3. 1991.