Chapter 9

Cortical Maps

9.1 Introduction

V1 cells are selective to many different aspects of the input, including at least location on the retina, ocularity (favoring input from one eye over the other), and orientation, spatial frequency, direction and speed of movement of bars or gratings. These selectivities are arranged across cortex in a regular manner. For instance, figure 2.4 shows the topographic nature of the map from one eye to the cortex – *ie* neighboring cells in the cortex have receptive fields in neighboring (and overlapping) parts of the retina. Figure 9.1 shows the result of an optical imaging experiment that investigated how some of these selectivities are arranged across a region of the primary visual cortex of a macaque monkey. The thick lines show the boundaries of the ocular dominance stripes, showing a part of a characteristic fingerprint-like pattern (seen more completely later in the chapter in figure 9.15A). The thin lines show iso-orientation contours, *ie* locations where the preferred orientations are roughly constant. These indicate, by the regions they enclose, that whole neighborhoods of cells favor similar orientations, and show how these neighborhoods are arranged with respect to each other and with respect to the ocular dominance stripes. Large scale order is evident, as in the singularities in the orientation map (called pinwheels), where the patches for all orientations meet, and which tend to occur at the centers of ocular dominance stipes, and as in the linear zones, where the iso-orientation domains run parallel to each other, and which tend to occur at, and to run perpendicular to, the boundaries of the ocular dominance stripes. These relationships between orientation and ocular dominance concern local order; the maps also have more global order in macaques in the sense that the two dimensional power spectra of the patterns of each across extended patches of cortex are elliptical, and the major axes of the ellipses for ocular dominance and orientation are orthogonal to each other.

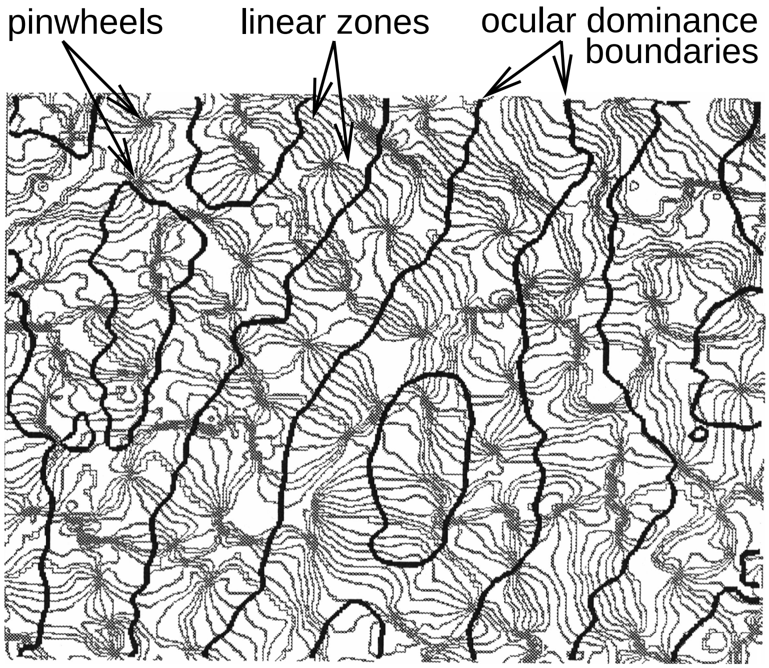

Figure 9.1: Orientation domains and ocular dominance. Contour map showing iso-orientation contours (grey lines) and the boundaries of ocular dominance stripes (black lines) in a 1.7mm patch of macaque primary visual cortex. Iso-orientation contours are drawn in intervals of 11.25°. Pinwheels are singularities in the orientation map where all the orientations meet; linear zones are extended patches over which the iso-orientation contours are parallel. From Obermayer & Blasdel (1993); Erwin (1994).

Higher visual areas and also auditory and somatosensory projections to cortex are also systematically ordered (tonotopic maps and the somatosensory homunculus for the latter two). Although it is believed that many aspects of the nature and development of these maps can be characterised using methods similar to those we discuss for V1, they have not been so extensively modeled, and so we do not discuss them further. Other regularities can also be found in V1, such as a grid-like pattern of blobs of high activity of an enzyme called cytochrome oxidase, which are reported to coincide with connections to superficial layers of cortex from a particular class of cells inbetween the main layers of the LGN and correspond to cells that are poorly selective for orientation, but are sensitive to color. These are have not been well modeled.

Abstractly, all the maps show that cortical cells with similar selectivities tend to be nearby, and, conversely, that nearby cortical cells have similar selectivities. However, making this absolutely true is impossible, since cells are arranged on an essentially two dimensional cortical sheet, but are selective in many more dimensions. Actual cortical maps show regularities as to how these two general rules are violated, and it is the regularities that models of the selectivity maps must capture. Some of the regularities arise from basic mathematical facts. For instance, it is impossible to have a continuous map between spaces with different topological characteristics (called homotopy classes) without having singularities. This constrains the nature of the map of orientation (which lives in a space that is characterised as the unit circle) onto cortex (treated as a two dimensional sheet). Given that pinwheels are the singularities, it is possible to draw further conclusions about the order in which orientation changes around each (Tal & Schwartz, 1997; Swindale, 1996).

There are three underlying tasks for the models. The first is to characterise the maps quantitatively, understanding what structures they contain, and how much variability there is between different parts of the map and the maps in different animals. The second task is to understand how these maps develop over the course of development of the animal as a whole, and how this is affected by manipulations to the input or nature of the developing system. The third task is to understand what purpose, if any, might be served by the observed arrangements.

There are few compelling suggestions for interpreting the maps, but there is an embarrassment of mathematical models to account for their development. Unfortunately, experimental data on the arrangement and development of selectivities are currently in substantial flux, significantly outpacing all the models. Further, the models do not fit neatly into a set of logical categories. We will therefore study rather abstract versions of the problem of maps, focusing on topography, and the way that the breakdown of topography is engendered by ocular dominance. As we will see, whereas some models are only intended to explain the maps, others also capture the development of the selectivities of the individual cells that comprise the maps. In general, the models focus only on a few aspects of development culled from a very large collection. Most aspects of chemical markers, the growth and retraction of axons and dendrites and development in three rather than two dimensions of cortex have received little attention.

There is an intricate developmental interaction between innate specification and environmental influence, and the development of the maps is determined by both activity-independent and activity-dependent means. Crudely speaking, activity-independent mechanisms are believed to control initial axon targeting and the appropriate layer in which to make connections and to establish a coarse topographic order in projections. Then other activity independent and activity-dependent mechanisms refine this order and help to create the regular receptive fields of individual cells, and the patterns of arrangement of multiple cells, such as ocular dominance stripes. Although activity dependent development has been a particularly seductive target for modeling because it fits so well with the extensive study of learning rules, it is likely that a large majority of neural developmental processes are completely unaffected by activity. Indeed, there is an active experimental debate about the true extent of activity-dependence,

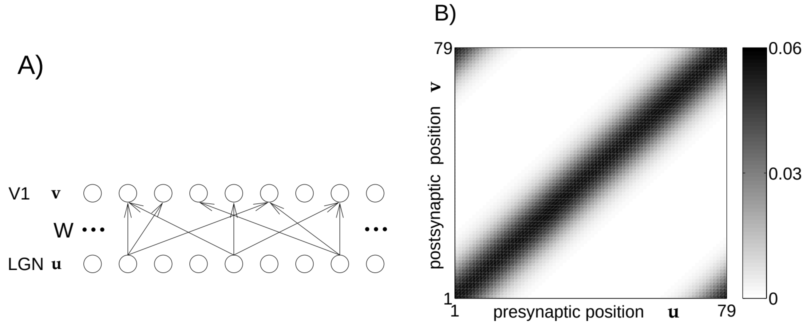

Figure 9.2: A) Abstract version of the task for creating or refining a topographic map from the LGN to V1. Only a few of the connections are shown. B) Example of topographically ordered weights **W** from **u** to **v**. Wrap-around (toroidal) boundary conditions have been assumed (unit u_1 is connected to v_{79} as well as v_2), which eliminate edge effects.

even for such complex maps as the orientation map. Note that there is *no* necessary equivalence between activity-dependence and *environmental* influence – the patterns of activity that drive adaptation (for instance waves of activity moving slowly across the retinas of ferrets during early development) can be created by internal mechanisms and reflect any external milieu only indirectly.

9.2 Topographic Maps

The most fundamental characteristic of the map shown in figure 2.4 is topography. We consider the problem shown in figure 9.2A of specifying and creating in a set of weights **W** a topographic map between input units **u** (modeling the retina or the lateral geniculate nucleus) and output units **v** (modeling the tectum or the primary visual cortex). We focus on the case that the topographic map is dimension preserving, *ie* the input and output units are arranged in the same number of (spatial) dimensions. It is straightforward to characterise topographic maps in this case, and therefore we focus on algorithms that are intended to model their development and refinement. As a further simplification, we treat **u** and **v** as if they are arranged in just one spatial dimension. Figure 9.2B shows an example of a set of artificially constructed weights from **u** to **v** that exhibit topographic order. Although this identity map might seem trivial to create, making it adapt appropriately to the (possibly different and changing) shapes of the input and output structures is tricky.

Activity-independent mechanisms for initial targeting and for establishing coarse topographic order have been identified. Establishing topographic order requires some explicit report of neighborliness in both innervating and innervated structures. Some possible mechanisms are long range (possibly diffusion-based) chemoattraction and chemorepulsion and short range (possibly membrane-bound) contact adhesion and repulsion. Families of guidance molecules (such as ephrins) have been identified which exist in concentration gradients over target structures, are recognised by receptors in innervating axons (such as Eph receptors), which themselves exist in gradients over input structures, and the interaction between the two affects the growth of axons (in this case, usually being repulsive). Exactly which collection of molecules are involved is still under experimental investigation – the same molecules appear to be repulsive and attractive to different sets of axons and/or at different times of development. Another mechanism is axon neighborliness – in space and/or time. Axons from neighboring cells can stick together (fasiculate) as they grow. Further, if neighboring cells are born at similar times, the temporal order of outgrowth of their projections can be used to control order in the maps. All these mechanisms are only just starting to attract quantitative models.

Modeling has been much more popular for the refinement of the (presumably coarse) topography left by such mechanisms. Two classes of suggestion have been made. One comes from study of the topographic map from the retina to the optic tectum of lower vertebrates such as frogs and toads; and the other is a version of the sort of Hebbian algorithms discussed in chapter 8. Both work in similar ways – information about neighborliness in the input (defined by chemical markers or correlated activities) is propagated to the output (by diffusion or through spikes). Neighborliness in the output is enforced by allowing each output unit to influence proximal units (by diffusion or through explicit connections between the output cells). Finally, a form of positive feedback comes from strengthening connections between input cells and output cells that match.

The Tea-Trade Model

The map from the retina to the optic tectum in lower vertebrates such as frogs and goldfish has long been the target of experimental investigation. The map is mostly contralateral, although in some species there is a small ipsilateral projection via an intermediate nucleus. In normal animals the map is topographically ordered. In some species, there is also some order in the optic nerve and the optic tract, which are the portions of the pathways from the retina to the tectum, upstream and downstream of the optic chiasm respectively. During normal development, the retina grows cells radially whereas the tectum grows cells rectilinearly, and since the maps are constantly in register, substantial continual plasticity must occur.

Activity-dependent and activity-independent effects are thought to be important for refining the map from what may be an initially crude topographic state. The activity-independent tea-trade model (von der Malsburg & Willshaw, 1977; Willshaw & von der Malsburg, 1979; so called because of a complicated analogy between topographic maps and the importation, blending and distribution of varieties of tea) was one of the first quantitative accounts of map formation to be developed. Various qualitative theories had incorporated the principles that it adopts, namely that chemical markers should provide a signature for locations in the retina and tectum and that a matching process can be based on some measure of the similarity between the sets of markers. Its particular feature is that the markers in the tectum are modeled as arising by molecular transport along the retinotectal axons from the markers in the retina rather than being created *ab initio*. There is no direct evidence for the mechanisms it uses.

The tea-trade model employs a handful of distinct chemical markers that are released in spatially disparate spots in the retina. Let $\mathrm{c}^u_{b\rho}$ be the concentration of marker ρ in cell b. The superscript u (and later v) indicates that the concentration is in a retinal (and later tectal) cell. At retinal location b, the concentration of marker ρ changes according to the diffusion equation:

$$\frac{\partial \mathrm{c}^u_{b\rho}(t)}{\partial t} = D\frac{\partial^2 \mathrm{c}^u_{b\rho}(t)}{\partial b^2} - \gamma \mathrm{c}^u_{b\rho}(t) + s^u_{b\rho} \tag{9.1}$$

where γ is a decay rate, D a diffusion constant, a discrete version of the spatial derivative is used, and

$$s^u_{b\rho} = \begin{cases} q & \text{if } b \text{ is the source for marker } \rho \\ 0 & \text{otherwise} \end{cases} \tag{9.2}$$

Figure 9.3A shows an example of the resulting stable concentrations for the different markers (*ie* $\partial \mathrm{c}^u_{b\rho}(t)/\partial t = 0$). Clearly, neighboring cells have similar signatures in terms of the concentrations of markers. To make the marker signature for a cell independent of the total concentration of the markers that it sees, normalised values:

$$\hat{\mathrm{c}}^u_{b\rho} = \mathrm{c}^u_{b\rho} / \sum_{\rho'} \mathrm{c}^u_{b\rho'} \tag{9.3}$$

are actually used.

A measure s for judging the similarity between the marker concentrations in two cells b and b' is:

$$\phi(u_b, u_{b'}) = \sum_{\rho} \min\left(\hat{\mathrm{c}}^u_{b\rho}, \hat{\mathrm{c}}^u_{b'\rho}\right) . \tag{9.4}$$

This expression is appropriate since the concentrations in both cells are normalised – cells b and b' will only be seen as being similar if $\hat{\mathrm{c}}^u_{b\rho}$ is large for those ρ for which $\hat{\mathrm{c}}^u_{b'\rho}$ is also large. Figure 9.3B shows the resulting matrix of similarities for the case in figure 9.3A, showing how the markers report on neighborliness in the retina.

The markers are transported through the optic nerve, and cell v_a in the tectum receives them through synapses of weight $\mathbf{W}_{ab}$ from retinal cell u_b. For activity-based models, the weights govern the effect on a postsynaptic

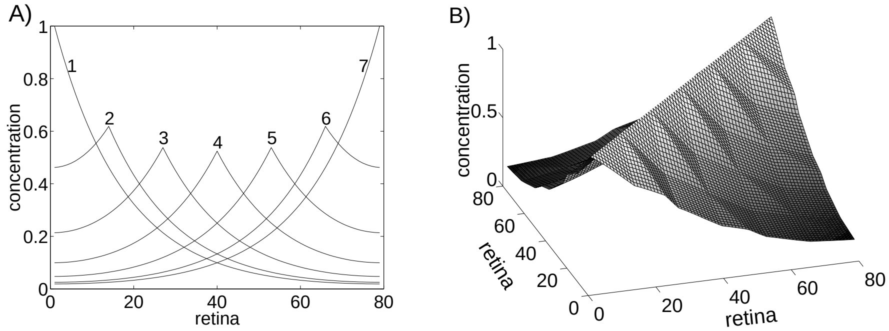

Figure 9.3: Marker diffusion and correlation in the tea-trade model. Sources for 7 different markers were placed equidistantly along a one-dimensional retina of 79 cells. The diffusion constant was $D = 0.85$; the decay rate was $\gamma = 0.003$. A) Scaled stable concentrations ($\mathbf{c}^u$) for the 7 marker types. B) Similarity matrix between retinal cells using the measure in equation 9.4. This is not translation invariant because of the normalisation. The value along the diagonal is always 1. After Willshaw & von der Malsburg (1979), though with different parameters.

cell of a presynaptic spike. For the tea-trade model, the weight also governs the degree of release of the markers into the tectum. There is as yet no evidence for this (or indeed for the involvement of diffusible markers in creating or refining the topographic map). Let $\mathrm{c}^v_{a\rho}$ be the concentration of marker ρ in tectal cell v_a. Then $\mathrm{c}^v_{a\rho}$ evolves according to the same equation 9.1, except that rather than coming from intrinsic sources, the term $s^v_{a\rho}$ represents the effects of diffusion from the retina:

$$s^v_{a\rho} = \sum_b \mathbf{W}_{ab} \hat{\mathrm{c}}^u_{b\rho}. \tag{9.5}$$

Except for the case of regeneration that we discuss below, sufficient time is allowed for equation 9.1 to reach a stable state. The diffusion is what makes nearby cells on the tectum look similar to each other. A Hebbian rule for neural activity would change weight $\mathbf{W}_{ab}$ according to the extent of coactivation of cells v_a and u_b. Here, the equivalent of coactivation comes from the similarities, defined by $\phi(v_a, u_b)$ of the (normalised) marker signatures. Thus, the synaptic strength is modified according to:

$$\Delta \mathbf{W}_{ab} \propto \phi(v_a, u_b) - \overline{\phi}(u_b) + \kappa \tag{9.6}$$

where $\overline{\phi}(u_b) = \frac{1}{n} \sum_a \phi(v_a, u_b)$ is the mean similarity over the cells to which u_b connects (making something like one of the Hebbian covariance rules) and $\kappa = 0.003$ is a small constant which represents an intrinsic tendency for connections to strengthen and, empirically, is required to prevent the map from degenerating. The weights are then multiplicatively normalised presynaptically, keeping constant at W the sum for each retinal axon. Very weak connections ($\mathbf{W}_{ab} < .005W$) are eliminated, and a retino-tectal axon can also sprout a *new* connection onto a tectal cell neighboring one where

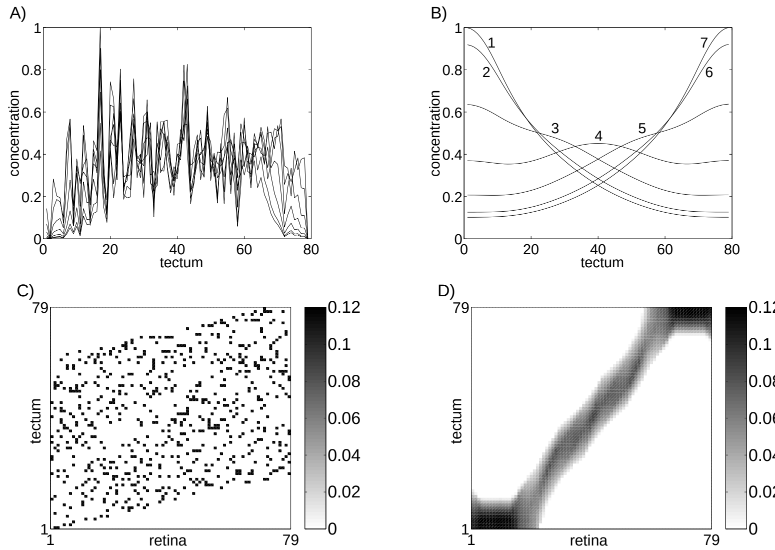

Figure 9.4: Topographic development in the tea-trade model. The top plots show scaled versions of the concentrations $\mathbf{s}^v$ transported to the tectum; the bottom plots show the weights, with the gray-scale indicating the magnitude. A;C) initial configuration, with the coarse initial topography and equal weights. The concentrations in (A) are too confused to label. B;D) terminal configuration, with a smooth topographic map. After Willshaw & von der Malsburg (1979), though with different parameters.

it already has a sufficiently strong weight ($\mathbf{W}_{ab} > .02W$), provided that the connection does not already exist.

The gross orientation of the map is specified by a very weak bias as to which connections initially exist, which may be partly created by the interaction of Eph receptors existing in retinal axons and ephrin ligands in the tectum, both in topographically ordered concentration gradients. The tea-trade model is designed to account for the refinement of the crude map, not the crude map itself. Figure 9.4C;D show the initial and final configurations of the weights in the one-dimensional case – plainly an ordered map forms (although with significant edge effects). Because of the ability to sprout and retract synapses, the population of extant weights changes significantly. However, a similar map could have been produced with more graded weights, as in the Hebbian models we considered in chapter 8. Figure 9.4A;B show the (scaled) transported marker concentrations $\mathbf{s}^v$ at the beginning and end of adaptation. Toroidal boundary conditions are not used in this case.

The topographic map is particularly easy to manipulate experimentally, since growth and regrowth of connections are extremely robust to large manipulations, including ablating part of the retina or tectum at various points during development, and crushing, and allowing a regrowth of the optic nerve. Revealing findings from these experiments include the facts that half a retina will eventually regenerate ordered connections across the whole tectum, and a whole retina will compress its projection onto half a tectum. Also, regeneration is not random – when a half retina regrows connections back onto the whole tectum, the initial regrowth to the tectum roughly reconnects retinal cells with the tectal cells to which they were formerly attached, but then the whole map readjusts so that the retina is smoothly represented across the whole tectum. If half the retina and the non-matching half of the tectum are both removed, then on regrowth, one of two topographic maps is established. The map can run either forwards or backwards with respect to its normal orientation, depending on aspects of the experimental technique.

When the optic nerve is crushed, the synapses made by the retinal ganglion cells in the tectum do not survive. Since the experiments indicate that, after regeneration, various aspects of the map survive this operation, models that store everything about the map in the retino-tectal synaptic weights alone cannot be wholly correct. A persistence of the concentrations of the markers in the tectum after the optic nerve is crushed is the non-synaptic memory that allows the tea-trade model to account for the various experiments implying that there is something other than the retino-tectal connections that stores information about the map. As an example, figure 9.5 shows what happens in the model when the retino-tectal axons are crushed, non-matching halves of the retina and tectum are removed, and the axons are allowed to regrow.

Crushing and regrowth have various effects in the model. First, the original synapses are lost, and the set that regrow is determined by exactly the same rules as for the initial growth. In this case, the weights shown in figure 9.5D come from a different random sample from the distribution that led to the upper left quadrant of figure 9.4C. Note, in particular that they are not symmetrical, since the crude initial mechanism is not adapted to the loss of parts of the retina and tectum. The second effect is that diffusion and decay of the makers in the tectum is assumed to happen whilst regrowth occurs. Since no marker diffuses from the retina until the axons make contact with the tectum, the longer it takes, the flatter the concentrations of the markers – as is evident in figure 9.5B;C, which model cases in which regrowth takes different lengths of time. Here, flattening is a form of molecular forgetting. The final effect is that a small amount of synaptic adaptation from the regrowing axons is assumed to happen before substantial diffusion of the markers from the retina occurs. This (and the original fact that diffusion is much faster than synaptic change) are extreme simplifications. Finally, the concentrations of markers in the remaining cells in the retina are assumed to be unchanged – the top plot of figure 9.5A is just the normalised version of the portion of figure 9.3A corresponding to cells $u_{40} - u_{79}$.

Figure 9.5E shows the consequences if regrowth happens before significant flattening of the concentrations of the markers in the tectum. Comparing figure 9.5A;B, one can see that the match of the markers favors an inverted

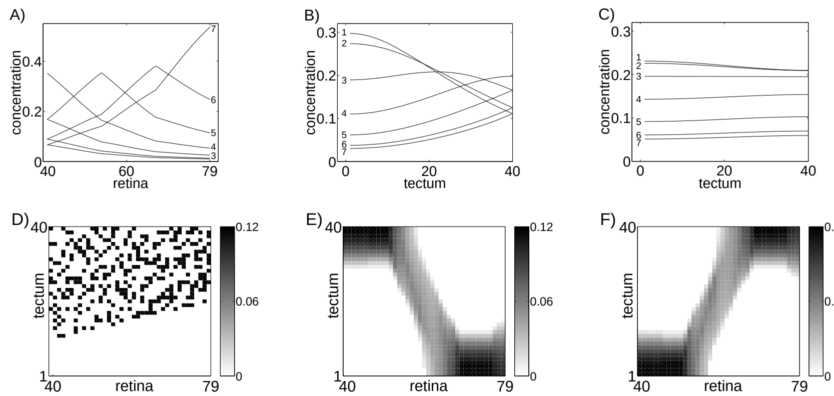

Figure 9.5: Crushing and regrowth between non-matching halves of retina and tectum. A) Concentrations of markers ($\mathbf{c}^u$) in the retina after cells $1 - 39$ are removed. D) initial connections that regrow in this non-matching case. The numbers are the marker labels. B;C) normalised marker concentrations ($\mathbf{c}^v$) in the remaining part of the tectum before adaptation begins after no (B) or substantial (C) diffusion and decay. 10 steps of synaptic adaptation happen before subsequent diffusion of markers from the retina. E;F) terminal configuration of the weights in these two cases.

map – the first remaining cell in the retina (u_{40}) is best matched to the *last* remaining cell in the tectum (v_{40}). The weights at the end of subsequent development shown in figure 9.5E, indicate that this bias can outweigh the converse bias that comes from the initial weights (shown in figure 9.5D), which favor a normally oriented map. Figure 9.5C;F shows that if substantial molecular forgetting occurs first, then the bias of the initial weights is more powerful, and the final map is conventionally oriented. Both outcomes are observed experimentally.

The activities of nearby retinal ganglion cells are known to be correlated, and indeed, activity-dependent mechanisms are also known to exert control over the retino-tectal map. For instance, if a experimentally introduced third eye forms connections to one of the tecta, then it competes with the projection from the normal eye, leading to a pattern of stripes. This pattern is disrupted if the NMDA receptors in the tectum are blocked. The tea-trade model lacks any activity-dependent component – and even its activity-independent components require updating in the light of recent evidence about molecular guidance mechanisms. Nevertheless, the tea-trade model established important concepts that underlie later models, and it also investigated factors such as sprouting that are significant biological possibilities that many models ignore. Sprouting is clearly necessary if the bias that comes from the initial weights is harmful, as in the case that half a retina grows back onto a whole tectum.

Hebbian Learning

A standard way to produce a refined topographic map is to use a form of activity-dependent Hebbian learning, as described in chapter 8. Correlations in the activities of neighboring cells and Hebbian synaptic plasticity replace the mechanisms based on chemical markers employed by the tea-trade model.

Hebbian learning can only be used to create topography and other aspects of the maps if the input correlations are appropriate. The activity of retinal ganglion cells (RGCs) has been well characterised in goldfish, cats, and particularly the ferret. The ferret is an attractive mammal to study because it is born after much less intrauterine development than similar animals such as cats. It can therefore provide a window on earlier developmental processes. Around the time of birth in ferrets, there are known to be waves of activity slowly moving across the retina, that are observable in the activity of the RGCs. Later on in development, the OFF-center RGCs participate in additional waves amongst themselves that the ON-center RGCs ignore. The waves will certainly induce correlations amongst neighboring cells in each retina, although not between retinas, since there is no known mechanism by which waves in each could be coupled. The output of the visual thalamus, the lateral geniculate nucleus (LGN) at just these times has not been characterised. There could be interactions between the projections from the two eyes in the LGN or in the reticular nucleus of the thalamus (a structure that surrounds the thalamus and inhibits thalamic relay cells), or indeed through thalamic interneurons, although there is no evidence on this point. Although blocking all activity in the retinas disrupts development, it is not known whether the patterns in the activity that comes from the waves is critical.

Figure 9.6 shows an example of the use of Hebbian learning to create a topographic map. This model is very closely related to the models of cortical maps discussed in the first part of chapter 8, using a coarsely topographic and fixed arbor function which indicates how many synapses there are between an input and an output cell (the solid line in figure 9.6A), fairly broad, translation invariant correlations between nearby input cells (the dashed line in figure 9.6A), and crudely topographic initial weights (figure 9.6D). These aspects of the model are in addition to standard subtractive normalisation (keeping the sum of all the weights into a single postsynaptic cell at 1), and synaptic saturation (preventing the weights from being negative or being greater than a maximum value) The interaction weights between output units i and j (which are located at positions i and j in a one-dimensional grid) take the form of a Mexican hat coming from a difference of two Gaussians:

$$\mathbf{K}_{ij} = e^{-(i-j)^2/\sigma_+^2} - \tfrac{\sigma_+}{\sigma_-} e^{-(i-j)^2/\sigma_-^2} \tag{9.7}$$

Ttwo different scales of Mexican hat are shown in figures 9.6B;C), with $\boldsymbol{\sigma} = (\sigma_+, \sigma_-) = (6, 24)$ for figure 9.6B and $\boldsymbol{\sigma} = (3, 12)$ for figure 9.6C. Wrap-

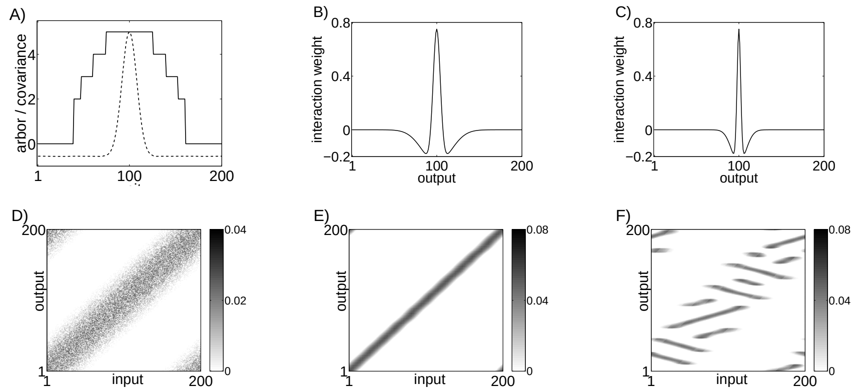

Figure 9.6: Hebbian topographic refinement. A) solid: coarsely topographic arbor function between input (presynaptic) unit 100 and the 200 *output* (postsynaptic) units; dashed: covariance between input unit 100 and the 200 *input* units. The covariances have been scaled to match the arbor function (by a factor of 22000). D) starting values for the weights (all weights are shown taking the arbor function into account). B;C) Interaction weights between output unit 100 and all the other outputs. B;E) Broad output interactions give rise to a characteristically refined topographic map. C;F) Narrow output interactions produce a characteristically segmented topographic map.

around boundary conditions are used – making the overall matrix of connections Töplitz. As in chapter 8, the activities $\mathbf{v}$ of the cortical units are determined as $\mathbf{v} = \mathbf{KWu}$, where $\mathbf{W}$ are the adaptive bottom-up weights.

Many of these features (notably the subtractive normalisation and the negative part of the cortical interaction function) are not required merely to create a topographic map (and may even be harmful) but they will be required later. The output interaction function plays the critical role that diffusion on the tectum plays for the tea-trade model, *ie* allowing neighborliness on the cortex to influence the adaptation of the synaptic weights.

The result of making all these various assumptions is shown in figure 9.6E in terms of the terminal stable values of the weights for the case of broad cortical interactions. Sharp and precise topography results. To demonstrate that topography is not completely trivial, figure 9.6F shows terminal values of the weights for the case of narrow cortical interactions. In this case, although the map is still topographic at a very coarse scale (this is enforced by the fixed arbor function in any case), at a finer scale, the map breaks down into separated patches with fine-scale topography, in some of which the map actually runs *backwards* with respect to the global topography (an outcome that may correspond to fine details of the structure in area V2). The breakdown of the topography can be understood intuitively by considering the mismatch between the correlation and cortical interaction functions (Takeuchi & Amari, 1979). If the map were precise as in figure 9.6E, then, for instance, the activities of cortical units 100 and 106 would be positively (though moderately) correlated through the input. However, these cells inhibit each other through the cortical interaction function, which would then weaken such weights. Patchy solutions such as the one seen in figure 9.6F tend to make cortical cells that inhibit each other receive inputs from better separated parts of the input, and therefore be less correlated. The broader the cortical interaction function, the less this effect.

9.3 Ocular Dominance and Topography

A more challenging case comes from considering ocular dominance in addition to topography. In cats and some species of monkeys, V1 appears to be parceled out into regions whose inputs primarily come from just one or the other eye. Only in layer IV do monocular cells (*ie* ones that are strongly selective to just one eye) predominate, but the other cells in a column tend to have similar biases to the cells in layer IV, which is why ocular dominance is visible in optical imaging experiments whose results only weakly reflect direct layer IV activity. In cats and macaques, ocular dominance stripes are roughly 0.5 mm wide. The thick lines in figure 9.1 show the boundaries of ocular dominance stripes. The similarity between the ocular dominance stripes and other biological patterns, such as fingerprints or the stripes on a zebra, has long been noted, and various of the mathematical models employ very similar mechanisms to reaction-diffusion equations, which were first suggested by Turing (1952) as a general mechanism for biological pattern formation.

Monocularity forms over the course of pre-natal and post-natal development, in a way that is sensitive to manipulations that affect activity in the afferent structures or in the cortex, or that affect plasticity in the cortical synapses. This is at least true during a critical period for plasticity of a few weeks. For instance, kittens that are made strabismic (*ie* unable to converge the two eyes on the same visual target) develop to have a greater proportion of monocular cells and also wider ocular dominance stripes in V1 (though not in V2, which also receives a direct projection from the LGN). Further, kittens brought up with artificially highly correlated inputs do not develop monocular dominance. Results such as these, and those on the effects of monocular and binocular deprivation (*ie* denying one or both of the eye normal input), and also the basic findings on the width of the stripes, are what models attempt to account for.

Of course, there is no necessary relationship between the presence of monocular cells and that of ocular dominance stripes. Further, ocular dominance stripes are not completely obligatory – for instance, adult marmosets lack them. Marmosets do have ocular dominance stripes at birth, but then lose them over the course of normal development. If marmosets are monocularly deprived, then the stripes are preserved into adulthood. Conversely, we have already mentioned that whereas certain lower vertebrates such as frogs do not normally have ocular dominance stripes in their optic tecta, striped projections are produced if two equivalent parts of a single eye (*eg* two nasal halves) or two different eyes (as in the famous three-eyed frogs) are forced to innervate the same tectum. This has also led to much speculation about whether the stripes are somehow just byproducts of mechanisms that instantiate other developmental goals.

Figure 9.7A shows one model of the ocularity and topography of connections in a one dimensional idealisation of layer IV of area V1. The pattern of one-dimensional stripes is apparent, as is the way that topography and ocular dominance interact. The map is topographic within each stripe, but breaks at the stripe boundaries, having sharp reversals. Cells near the boundaries of the stripes are binocular; cells in the middle of the stripes are strongly monocular. Toroidal boundary conditions are used for convenience, so the top and bottom of the plot are identified with each other.

Figure 9.7B shows the ocularity aspect of this map in a different way. The top part of the figure shows the difference between the sum of weights into the right eye and that into the left eye as a function of the number of the output unit. Since the weights are normalised to sum to 1 across both eyes, this difference varies between -1 and 1. For the moderately strongly monocular map in figure 9.7A, the values are mostly near ± 0.75, with only a few cells at the edges of the monocular patches being more binocular. The lower part of the figure shows part of the power spectrum of this alternating function. The power spectrum has one strong peak, indicating the frequency of alternations between the two inputs, and hence the width of the stripes. Since stripe width is experimentally observable, this power spectrum plot is useful for comparing models with data.

The calculation leading to figure 9.7B distills the feature of ocularity from the weights to each output unit. One can similarly distill the feature of topography by calculating the weighted average of the location in the input favored by a particular cortical cell. Figure 9.7C shows the features of ocularity and topography together in a two dimensional plot. Maps such as that in figure 9.7A look like sideways 'Z-folds' in this representation, with switchbacks where the ocularity changes. Clearly, not all weight matrices $\mathbf{W}$ can be adequately represented in this manner – only ones with sharp single peaks in the weights.

The task for models of the ocular dominance and topography map is to account for the presence and width of the stripes, and the way that these can be affected by manipulations such as changing the correlational structure of the inputs. We consider two sets of models, one based on the representation in figure 9.7A and the other on the representation in figure 9.7C. Included in the last class are models that attempt to understand what the purpose the map could serve. There are two different competitive aspects to ocular dominance, although some models handle them using a single mechanism. First is competition at the level of single cells for one eye to dominate, *ie* to create monocular cells at all. Second is competition at the level of the maps between input from the two eyes. There is also the same cooperative aspect as in topography so that nearby cortical cells tend to

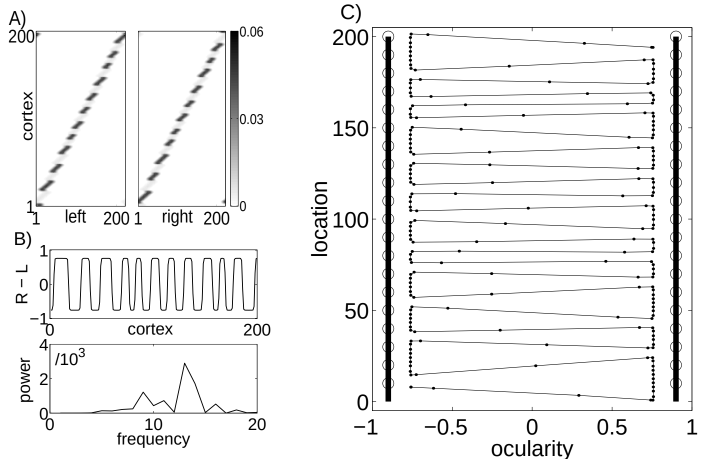

Figure 9.7: The combination of ocular dominance and topography. A) Weights (generated by a competitive Hebbian model, see section 9.3) from 200 input units in each of two projections (eyes) to 200 output units showing a combination of topography and ocular dominance. Topography dominates globally – but its continuity is broken by shifts in ocularity. B) Top: map of ocularity across the cortex, showing the difference between the total weights for the right and left eyes; bottom: the square of the absolute value of the Fourier transform of this map, effectively showing which stripe widths dominate. C) Abstract version of the map in (A), plotting ocularity (from B) *versus* the weighted average location in the input for each output unit. Note the sideways 'Z-fold' stripes where topography reverses. The dots show the locations of the individual cortical units; the thin line joining them shows their order in cortex. The gray bands and open circles depict abstract input activities and are used by feature-based algorithms. Since wrap-around boundary conditions are employed, topographically equivalent locations are substituted to avoid switchbacks between locations 0 and 200.

represent the same eye.

Recent evidence significantly complicates the picture of development of ocular dominance and its relationship with other aspects of the maps such as orientation. The orientation map (which is also shown in figure 9.1) is established extremely early in development and has been shown to remain stable even during the course of the drastic anatomical changes consequent on the development of ocular dominance, during which thalamocortical axonal arbors are undergoing very extensive remodeling. Important aspects of orientation selectivity are stable during monocular deprivation. These results and others suggest the primacy of the orientation map over the map of ocular dominance. Further, there is evidence from cats that there is an initial bias favoring connections from the contralateral eye over those from the ipsilateral eye that goes away during the time that the stripes are forming, provided that the animals are not deprived. These factors and their ramifications have yet to be fully captured in a model.

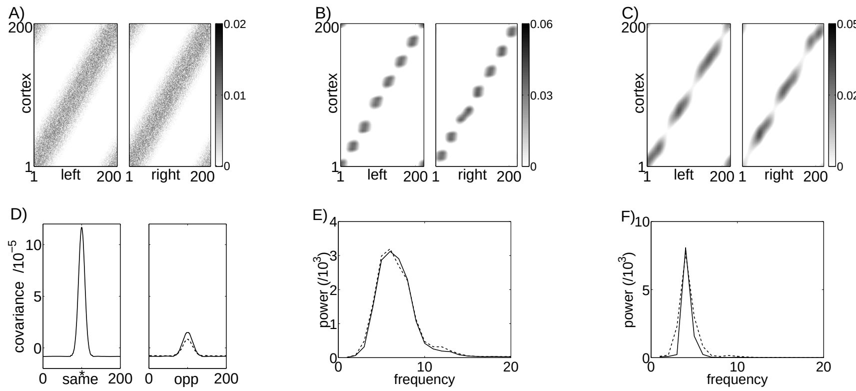

Figure 9.8: Hebbian development of ocular dominance and topography. A) The starting weights are broadly topographic and the two eyes are statistically equivalent. D) covariance between unit 100 in one eye and all the other units for weak (solid) and weaker (dashed) correlations. B) The final weights for a particular case for the cortical interactions shown in figure 9.6B show refined topography and clear ocular dominance. E) The power spectrum of such weights for weak (solid) and weaker (dashed) correlations averaged across many cases. C;F) The final weights for broader cortical interactions are distinctly broader and accordingly have power at lower spatial frequencies.

Hebbian Learning

The Hebbian model of topography described in section 9.2 incorporates aspects (such as subtractive normalisation) that were shown in chapter 8 to be sufficient to create ocular dominance when there is input from two eyes (Miller, Keller & Stryker, 1989). Figure 9.8A shows the initial weights from input to output, for which the two eyes are statistically identical. We refer to inputs as being from the eyes, although they are really from relay cells in the eye-specific layers of the LGN. Figure 9.8D shows covariance functions between all the input cells for cases of weak and very weak correlations between matching units in opposing eyes. For later convenience, we imagine that there is a baseline of negative covariance for distant units.

Sample final weights for the weak correlations and for two cortical interaction functions with $\boldsymbol{\sigma} = (6, 24)$ and $\boldsymbol{\sigma} = (10, 40)$ (in equation 9.7) are shown in figure 9.8B;C. As expected from the case of pure ocular dominance in chapter 8, the broader the cortical interactions, the wider the stripes. The solid lines in figure 9.8E;F confirm this by showing the power in various frequencies of stripes (as in the lower plot of figure 9.7B), averaged across many examples. Note in figure 9.8B how the regular topography of the map even *within* the eye-specific patches is somewhat disturbed in the presence of ocular dominance.

A Competitive Hebbian Model

The dashed lines in figure 9.8E;F show that reducing the correlations between the two eyes does not have a substantial effect on the width of the stripes. This is expected on the basis of the discussion in chapter 8, in which the exclusive role of the cortical interaction function in determining stripe width was established, but is inconsistent with the data showing that strabismic cats (which presumptively have lower correlations between the eyes than normals, at least after eye opening) have wider stripes. In fact, making the two projections moderately positively correlated prevents single cells from being monocular in any case, thus making it hard to investigate the consequences of increasing the degree of correlation. This remains true even with the weak baseline anti-correlations shown in figure 9.8D.

An alternative developmental model (Willshaw & von der Malsburg, 1976; Kohonen, 1982; Goodhill, 1993; Bauer *et al*, 1997) employs the same underlying Hebbian learning rule, but uses a different and non-linear model of cortical interactions to determine the activities of the output units, which in turn determine how the weights change. Chapter 7 describes many possible effects of Mexican hat interactions in a recurrent network, particularly in the face of a threshold non-linearity that makes firing rates non-negative. The competitive model uses an abstraction of one of these in which the cortical interactions find a single cortical unit that has the greatest feedforward activation in response to each input **u**, *ie* they find a value of a^* such that

$$v_{a^*} = \max_a \{v_a\} \quad \text{where} \quad \mathbf{v} = \mathbf{W}\mathbf{u} \tag{9.8}$$

and produce a Gaussian bump of activity centered on a^* and involving a few neighbors of the most activated cell:

$$v_a = e^{-(a-a^*)^2/\sigma^2} \tag{9.9}$$

where the parameter σ determines the width of the Gaussian, and is determined by the width of cortical interactions. A standard Hebbian rule then applies, adapting the weights as:

$$\Delta \mathbf{W}_{ab} \propto v_a u_b = e^{-(a-a^*)^2/\sigma^2} u_b \tag{9.10}$$

subject also to saturation and normalisation. The width of interactions can either be fixed or can be decreased slowly over iterations (a process called *annealing*), which tends to produce more regular maps, but is harder to justify biologically.

The non-linearity in equation 9.8, forcing just one single winner in the area of the cortex simulated in response to a single pattern, is clearly too strong. However, the consequences of weakening it are not known. Unlike standard Hebbian learning, this non-linearity also implies that it is not only the

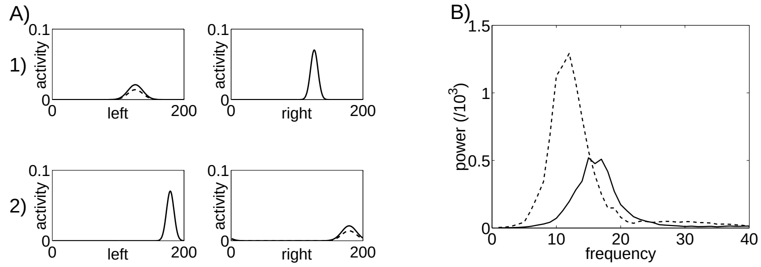

Figure 9.9: The competitive Hebbian model. A) Activities of the units for two different inputs (1;2). These are randomly centered Gaussians with the input to one (randomly chosen) eye a fraction of the input to the other, and coming from a broader underlying Gaussian, where the fraction (solid: less; dashed: more) governs the degree of correlation between the eyes. Inputs of this form, only with weaker correlations, generate the covariances between the input units shown in figure 9.6A. B) Average power spectra for the stripe patterns for stronger (solid) and weaker (dashed) correlations. The lower the correlation, the wider the stripes, as seen experimentally in strabismus.

second order statistics of the patterns of input activity that are critical, but rather the whole patterns. Following Bauer *et al* (1997), we consider inputs that are themselves randomly located Gaussian bumps in each eye, with one (randomly chosen) eye predominant over the other one to a variable degree that determines the total covariance between the eyes. Figure 9.9A shows the activation of the two eyes for two example inputs – the overall shape of the covariances between input units is just as in figure 9.8D, except with rather larger covariances between the eyes. The extra competition inherent in equation 9.8 is enough to favor ocular dominance even for fairly strongly correlated inputs.

Figure 9.7A shows the output of this competitive model for the case of weaker correlations between the eyes. The maps from this model are distinctly more regular than for the Hebbian model and are slightly less strongly monocular (although this depends on the details of normalisation and the annealing of the size of the neighborhoods in the cortex). Figure 9.9B shows the average power spectrum of the stripe pattern for stronger (solid) and weaker (dashed) correlations between the eyes. Unlike the Hebbian model, but as suggested by the experimental data, the weaker the correlations between the eyes, the wider the stripes, *ie* the lower the stripe frequency (Goodhill & Löwel, 1995). This can be justified intuitively by a rather tortuous chain of reasoning. However, it is easier to understand in the elastic net model, which is a more abstract version of a competitive model that we discuss below (section 9.3) and so we defer discussion to that section.

The results in figures 9.7 and 9.9 are based on decreasing σ over iterations. If it is fixed instead, there turns out to be a somewhat complicated relationship between σ and the widths of the stripes. In general, the larger σ, the broader the stripes – however, for some breadths of input correlations, the stripe width is quite insensitive to the size of the cortical interactions.

One interesting difference between the linear and competitive Hebbian models arises if there are inputs that are almost perpendicular to all the weights $\mathbf{Wu} \sim \mathbf{0}$. In the linear model, these inputs lead to almost no weight changes. However, in the competitive model, the unit that wins the competition in equation 9.8 and its neighbors always have the same final activation given in equation 9.9, and therefore the same amount of learning. This makes it comparatively much harder for the competitive model to come to 'ignore' certain classes of inputs.

It has proved difficult to analyse weight-based self-organising maps, since there is provably no energy function that learning rules such as equations 9.8 and 9.10 are minimising in the face of the sharp non-linearity inherent in the selection of a single winning unit. Riesenhuber, Bauer & Geisel (1996) have suggested an heuristic method which captures the behavior of simulated weight-based nets quite well.

Feature-Based Models

The depiction of the end point of development shown in figure 9.7C can be seen as motivating a number of algorithms that are more abstract than the linear and competitive Hebbian models we have so far seen. Their abstraction comes from operating in the low dimensional space of features (in this case, topographic position on the retina and relative ocularity) rather than the high dimensional space of weights, and, in some cases, using feature-based rather than explicit inputs. The advantage of feature-based models is that they are much more tractable, from the perspective of both theory and simulation, since they are based on many fewer dimensions.

The earliest feature-based models were noise models (Swindale, 1980; 1982; Rojer & Schwartz, 1990; Niebur & Wörgötter, 1993; Grossberg & Olson, 1994). These take random feature values or selectivities as their starting points and operate on them to produce the final forms of the maps. For instance, applying a Mexican hat filter to regularly positioned but random ocularities (as in chapter 8), and thresholding the result gives rise to a pattern of ocular dominance stripes whose widths are largely determined by the spectral response of the filter. The noise models offer a simple way to characterise the maps, in the sense that they indicate how the complex patterns seen can be generated by simple rules that may have little or no connection to the underlying mechanisms of development.

Other feature-based models offer alternative characterisations, with different, and sometimes less abstract, sets of simple rules. A feature-based version of the competitive Hebbian algorithm called the self-organising map (Kohonen, 1995) operates as follows. As in figure 9.7C, the weights for each cortical cell a are collectively represented by two numbers $\tilde{\mathbf{w}}_a = (x_a, z_a)$, *ie* the topographic location it prefers (x_a) and the amount by which it favors the right eye over the left eye (z_a). Further, input activity is also represented by two numbers $\tilde{\mathbf{u}} = (x, z)$. This is appropriate for input activities as in figure 9.9A which consist of a topographically compact bump (with center location x in input space) which is stronger in one of the eyes than the other. The ocular preference of the input z is typically taken as being $\pm Z$, where the value of Z reflects the degree of correlation of the eyes. In this low dimensional embedding, distance models dissimilarity, *ie* lack of correlation, so the higher the correlation between the two eyes, the smaller is Z. The two thick gray bands in figure 9.7C show the space of inputs for $Z = 0.9$. Actual inputs are drawn at random with x uniform across the whole topographic range, z being Z with probability 0.5 and x and z being independent.

The idea behind the key equation 9.8 in the competitive Hebbian model is that cortical interactions select the input that is most strongly activated by the input, and create a bump of activity involving just this unit and its cortical neighbors. This intuition is captured in the abstract algorithm by determining the cortical unit whose weights are best aligned with the input pattern. This can be achieved by finding a^* such that:

$$|\tilde{\mathbf{u}} - \tilde{\mathbf{w}}_{a^*}|^2 = \min_a \left\{ |\tilde{\mathbf{u}} - \tilde{\mathbf{w}}_a|^2 \right\} \tag{9.11}$$

This implements a competitive aspect of the model. Similarly, the intuition behind the learning rule in equation 9.10 is to pull the winning unit and its neighbors towards the input pattern presented, making it more likely to respond to this input on subsequent presentations. This can be captured using:

$$\Delta \tilde{\mathbf{w}}_a \propto e^{-(a-a^*)^2/\sigma^2} (\tilde{\mathbf{u}} - \tilde{\mathbf{w}}_a) \,. \tag{9.12}$$

As for the weight-based competitive model, the effect of the cortical interactions is to restrict plasticity to the most closely matching cortical unit and its neighbors, represented by $e^{-(a-a^*)^2/\sigma^2}$. These interactions are the cooperative aspect of the model and ensure that nearby cells will generally end up having similar selectivities. As before, σ is generally annealed, decreasing it slowly over iterations, although one could also use a fixed value.

Figure 9.10 shows the consequence of applying this algorithm for three values of Z and a fixed topographic distance of 200 from one end of the input to the other. The magnitude of Z controls how topography and ocularity interact. Large Z models weak correlations, as in strabismus, and leads to wide stripes; small Z models high correlations, and leads to narrow stripes.

Although there is no energy function for the feature-based self-organising map, it is possible to analyse some aspects of the outcome by analysing the combination of equations 9.11 and 9.12 as a form of Markov process and looking for the conditions about equilibrium states (such as a topographically ordered map with no ocular preference) under which instabilities

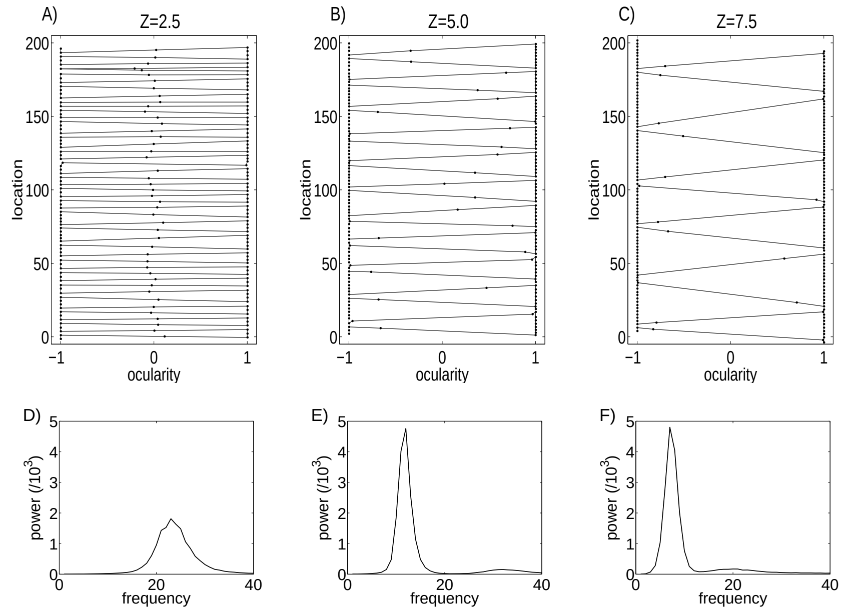

Figure 9.10: Feature-based competitive model for ocular dominance and topography for $Z = 2.5$ (A;D), $Z = 5.0$ (B;E), and $Z = 7.5$ (C;F). The upper plots show sample maps (in the form of figure 9.7C); the lower plots show the average power spectra of the ocularity component.

(such as random perturbations in ocular preference) will grow rather than shrink (Obermayer, Blasdel & Schulten, 1992). This analysis suggests that there is a critical separation between the two eyes, which depends on the width of the cortical interactions (σ), below which the map without ocular preferences is stable and above which it is unstable. This prediction is confirmed in simulations. One can try to use the same technique to predict the preferred spatial frequency of the unstable mode, and hence the width of ocular dominance columns. However, this prediction is reported not to be borne out in practice.

The justification of the feature based competitive model as an abstraction of the full Hebbian competitive model is only appropriate under stringent conditions on the weights and inputs. The algorithm is probably better thought of as a non-reductionist model of the maps that captures the essentials of how changing abstract characteristics of the input such as the degree of correlation between the eyes changes the nature of the map, but does not attempt to offer a mechanistic account of the process by which this happens.

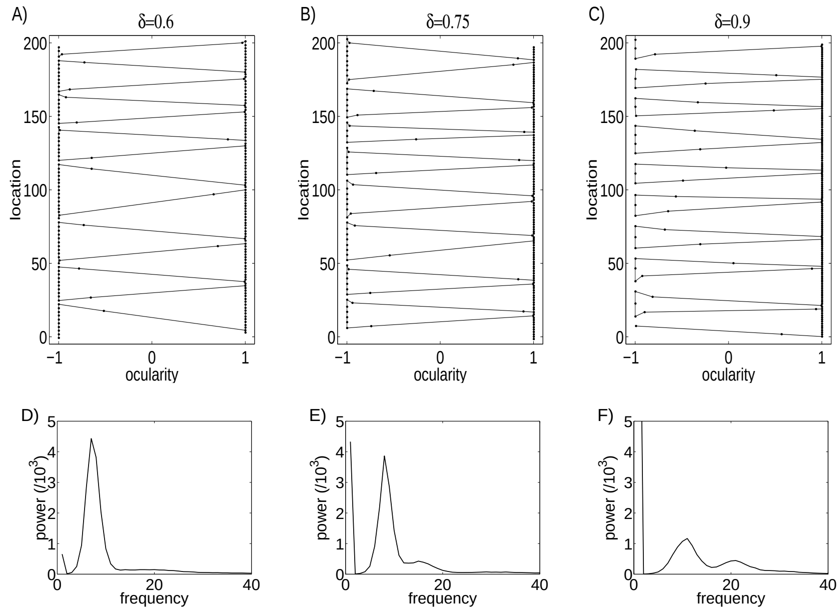

Figure 9.11: Maps and stripe power spectra for the feature-based competitive model in the face of monocular deprivation. The upper plots show example maps; the lower ones average power spectra of the stripes. Apart from the deprivation parameter $\delta = 0.6$ (A;D), $\delta = 0.75$ (B;E), and $\delta = 0.9$ (C;F), the parameters are the same as in figure 9.10C;F. The power in the constant mode in (F) is about 14 times greater than that for the maximal periodic mode.

Monocular Deprivation

If the input to one of the eyes is drastically reduced during the critical period, then, provided that no manipulations are made to activity or plasticity in the cortex, the area of V1 devoted to that eye is severely reduced. In the context of the feature-based competitive model, one straightforward model of monocular deprivation is to change the random distribution of inputs to the model, choosing $z = Z$ (indicating the right eye dominates) with probability δ and $z = -Z$ (favoring the left eye) with probability $1 - \delta$. The closer δ to 1, the more one eye (in this case, the left eye) is deprived. Figure 9.11 shows the consequence of monocular deprivation for $\delta = 0.6, 0.75$ and 0.9 and for $Z = 7.5$. The upper plots show example maps, which indicate clearly that the right eye grabs an increasingly large proportion of the cortical units (the dots), so the stripes get increasingly unbalanced. The lower plots show average power spectra of the stripes – there is a trend for the stripes to get a little wider for increasing δ – just as if this effectively reduces the correlation between the two eyes. The increasing power in the constant (0 frequency) mode results from the increasing imbalance between the eyes.

Weight-based models can be made to capture monocular deprivation in similar ways. The Hebbian model is rather sensitive to reducing the activity in one of the eyes – the competition inherent in subtractive normalisation tends to result in the deprived eye being eliminated altogether. This can be prevented by adding to the postsynaptic normalisation that creates monocular cells a form of presynaptic normalisation, preventing the sum of the cortical weights of any thalamic cell from decreasing too much. All the models need to be altered to take account of the new data mentioned above that there is an imbalance between the connections from the ipsi- and contra-lateral eye that goes away over the course of development in animals that are not binocularly deprived. The effects of this on their accounts of monocular deprivation are unclear.

Optimization Models

The self-organising map (and its peers, such as the elastic net which we discuss below), are important, since they are the only algorithms that have been systematically applied to study topography, ocular dominance and orientation selectivity. They can also be extended to other features. However, their justification as abstractions of full weight-based algorithms such as the competitive Hebbian model is only valid under special circumstances, and there have been various attempts to justify these algorithms at a more abstract level as (possibly approximate) ways of optimising one of a set of reasonable quantitative measures of the quality of the map. The best understood of these approaches is that of minimal wiring. Except for the few proposals to treat information processing as occurring in the full three dimensions of the brain rather than the largely single dimension of dendrites and axons (Montague, Gally & Edelman, 1991), or the apparent two dimensions of cortical maps, there is no reason to group together cells with particular response characteristics. The same connectivity could, in principle, result from different arrangements. However, the volume of the brain that is devoted to wires (*ie* axons and dendrites) is vast, and the structural and metabolic costs of connections are bound to be critical constraints on the size and thus processing power of a system whose computation depends on the communication between a large number of fairly simple processors. Therefore neural computations may be organised so as to minimise wiring (Cowey, 1979; Mitchison, 1991; Cherniak, 1995; Van Essen, 1998).

Consider the case of topography and ocular dominance. One way of looking at figure 9.7C is that the cortical units divide the bands of possible input into separate parts – so each cortical unit is responsible for a particular topographic region of a particular eye. The notional cost of the map from the LGN to the cortex does not depend strongly on the arrangement of the topographic regions, and therefore optimising the organisation of this map requires optimising the *intra-cortical* wiring that results from a particular arrangement. One might expect that cortical computations will be predominantly local in image space, *ie* if two retinal ganglion cells report information from nearby points in the world, then there will have to be substantial connections between the cortical cells that receive information from those cells. It is necessary to choose the right mapping from retinal to cortical space to minimise this wiring. The lower the correlations between the inputs from different eyes (*eg* arising from artificial strabismus), the fewer connections are required between the cortical cells that are connected to units from the different eyes.

Minimising this wiring leads to the qualitative picture of ocular dominance stripes – cells that are nearby in one eye report on similar aspects of the world, but so, to some extent, do cells in the other eye that can be pointed at the same scenes, and ocular dominance stripes can be seen as the compromise solution to these mutual wiring problems. Under this argument, those species that lack ocular dominance stripes would have either to have different characteristic correlations within and between the ganglion cells of the different eyes, or, and this is less likely, to be required to perform different computations with the visual information.

To make this more quantitative, one way to proceed is to specify a simple form of the connections and a simple cost for the wires, and then perform the optimisation (Goodhill & Sejnowski, 1997; Wiskott & Sejnowski, 1998). One might reasonably assume that wires should cost according to their volumes, rather than their lengths. Of course, this is complicated to calculate, given the branching structure of axonal arborisations and dendritic trees. Also, the resulting optimisation problems are extremely difficult to solve exactly – their complexity scales exponentially with the number of points. Heuristic algorithms are therefore in general use, and this is one way that the feature-based competitive models arise (Mitchison, 1995). Although stripes are qualitatively reasonable, it has actually proved quite difficult to find a quantitative cost function for which stripes of a variety of different widths are the globally optimal solution.

The key cases for minimal wiring come when there is a change in the effective dimension between input and output. A simple example of this is shown in figure 9.12A, which shows in the same format as the previous figures one way that two one-dimensional retinas might be mapped onto a one-dimensional cortex. As usual, dimensions here refer to the way that the units are *arranged, ie* to their neighborhood relations, and not to the *number* of units involved. Two different ways of counting wiring costs for such a map have been suggested, and are illustrated using figures 9.12B;C. We also consider the case of mapping a two-dimensional input surface onto a one-dimensional cortex, since this is more complicated and a bit more revealing.

Figure 9.12B;C is actually based on two further simplifications. The first is that instead of there being a whole continuum of possible inputs, a discrete set is used. Each discrete input represents itself and its local region of input space. This is reasonable since an algorithm such as the feature-based competitive model ultimately separates (tessellates) the input into regions that are the province of one cortical unit via the maximisation in

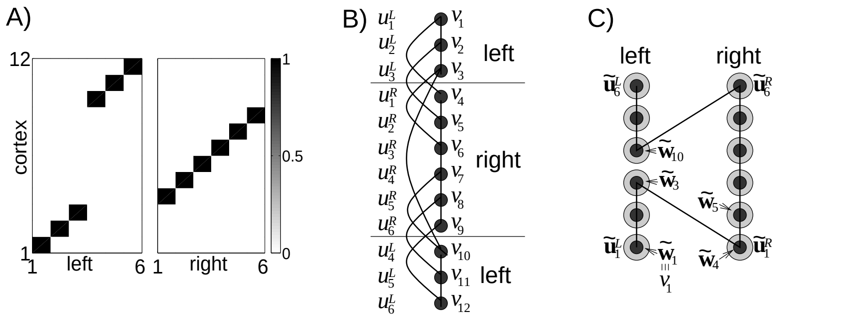

Figure 9.12: Ocular dominance, topography and minimal wiring. A) Symbolic weights for a particular one-one match from two one-dimensional retinas onto a one dimensional cortex. B) The intra-cortical wires (the lines connecting cortical cells) required for a particular match in the case that processing information from each input point requires connections between the cortical cells reporting on each adjoining input point and the matching point in the other projection (*eg* $v_6(\sim u_3^R)$ is connected to $v_5(\sim u_2^R)$, $v_7(\sim u_4^R)$ and $v_3(\sim u_3^L)$). C) The elastic net embeds the input points $\tilde{\mathbf{u}}$, and hence the matched cortical points $\tilde{\mathbf{w}}$, in Euclidean space and the cost is the energy of an elastic band stretched between the cortical points.

equation 9.8. The second simplification is that the discrete inputs are often identified with the input *units* themselves. The open circles in figure 9.7C show (a subset of) the inputs or input units in this representation, and the idea is that the inputs are drawn randomly from these. The second simplification is reasonable for ocular dominance and topography, since the input units also have topographic locations and eyes of origin, and we use it, but it is not so reasonable when features such as orientation are also included.

Dimension Reduction

Figure 9.12B shows the intracortical connections required for a particular match between the input and the cortex for which each cortical unit has a non-zero weight from only one input unit from just one eye of origin. In this simple case, point u_3^R is close in image space to points u_2^R, u_4^R and u_3^L and therefore direct connections are required between the associated cortical units v_6, v_5, v_7 and v_4. The cost of these is governed by the intracortical distances between v_6 and v_5, between v_6 and v_7, and between v_6 and v_3. More generally, if $c(a, a')$ is the cost of cortical wiring between v_a and $v_{a'}$, a wiring cost is

$$C^{\mathrm{MW}} = \sum_b \sum_{b' \in \mathcal{N}_b} c(a(b), a(b')) \tag{9.13}$$

where the sum over b is over all points in both retinas, the sum over b' is over the neighborhood $\mathcal{N}_b$ of point u_b (the neighborhood of u_1^R consists of points u_2^L and u_1^R) and $a(b)$ is the index of the cortical point a that is

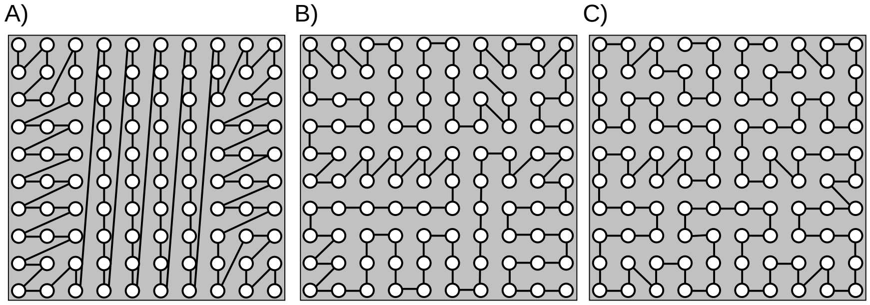

Figure 9.13: Minimal wiring and elastic net maps. All three figures show the mapping of the square input (white circles), which is 10×10 onto the linear output (solid lines), which is 1×100. For A;B, the neighborhood of each internal point consists of its four nearest neighbors under a Manhattan metric. A) Globally minimal wiring map with $s = 1$. B) Approximately minimal wiring map with $s = 0.5$. C) Elastic net map. Clearly, $s = 0.5$, which penalises short distances relatively more, is a better fit for the elastic net. From Durbin & Mitchison (1990)

matched to the input point u_b. One way to derive the cost function is to embed the cortical points in a Euclidean space, with cortical point v_a being at location $\tilde{\mathbf{w}}_a$, and to set

$$c(a, a') = |\tilde{\mathbf{w}}_a - \tilde{\mathbf{w}}_{a'}|^s \tag{9.14}$$

for some exponent s. Minimal wiring mappings are matches between input and output that minimise C^{MW}.

Figure 9.13A;B shows examples of minimal wiring solutions in the case for which the input feature space is a two-dimensional square and the cortex is just a line. The only aspect of the structure of the inputs that is important is the neighborhood of the input locations – here, these consist of the four locations in the Manhattan neighborhood of each input point (*ie* the ones in the direction North, East, South and West; although fewer for the points on the boundaries). Different values of the exponent s clearly lead to radically different maps. Note that the specification of the neighborhood is both key and obscure. For instance, in mapping information from ON-center and OFF-center cells in two eyes to a single area of two-dimensional cortex (as is believed to underlie the orientation selectivity of simple cells), it is not at all clear how to ascribe neighborhood relations amongst the two different classes of inputs.

Unfortunately, minimising even the highly simplified cost function C^{MW} is not easy. Doing so exactly requires solving a combinatorial optimisation problem, because one would have to search over all possible matchings between the input and the cortex. For the case that $s < 1$, Mitchison (1995) used a sophisticated heuristic argument to suggest an algorithm that is closely related to the self-organising map.

The Elastic Net

The wiring measure in equation 9.13 counts the cost of the intra-cortical connections. In the match in figure 9.12B, this charges a substantial amount, for instance, for the connection from u_1^R to u_1^L, since they are far apart in the cortex. An alternative is to count the cost of distortion of the cortical map in the space of inputs, and assume that maps with little distortion require little wire. Figure 9.12C gives an example to make one approach to this clear. Here, as in figure 9.7C, the input points $u_1^L, \ldots, u_N^L$ and $u_1^R, \ldots, u_N^R$ (which represent input regions) are embedded in the two dimensional Euclidean space of features, including topography and ocularity. Again, call $\tilde{\mathbf{u}}_i^L$ the features of the ith input from the left eye. The cortical points, if they match the input points, have locations $\tilde{\mathbf{w}}_a$ in this same feature space (we use $\tilde{\mathbf{w}}_a$ rather than the more natural $\tilde{\mathbf{v}}_a$ since the cortical features are really determined by the weights). Distortions of the cortical map occur when the features of neighboring cortical points are very different (as for v_3 and v_4 in the figure, with features $\tilde{\mathbf{w}}_3$ and $\tilde{\mathbf{w}}_4$ respectively). More formally, if the cost of the cortical distortion for connecting two neighboring cortical points v_a and $v_{a'}$ is $|\tilde{\mathbf{w}}_a - \tilde{\mathbf{w}}_{a'}|^s$, then one form of the distortion cost is

$$C^{EN} = \sum_a \sum_{a' \in \mathcal{N}_a} |\tilde{\mathbf{w}}_a - \tilde{\mathbf{w}}_{a'}|^s, \tag{9.15}$$

where the sum over a is over all points in the cortex, the sum over a' is over the cortical neighborhood $\mathcal{N}_a$ of point v_a (in figure 9.12B, the neighborhood of v_4 is just v_3 and v_5). Note that in figure 9.12C, there is a substantial cost for the connection between v_3 and v_4, since the cortical map is highly distorted at this point, however, there is no link between v_1 and v_4, and therefore no associated cost. These costs are the exact opposites of the measure in equation 9.13. The exponent used in equation 9.15 is usually $s=2$, in which case the cost function C^{EN} is exactly like the energy function for a perfect elastic band stretched (in feature space) between the cortical points.

The cost function C^{EN} is only an appropriate measure of the distortion if there is a map between input and output, *ie* if if every input point $\tilde{\mathbf{u}}$ is matched by a cortical point $\tilde{\mathbf{w}}$. Minimising C^{EN} subject to the constraints of such a map is again a challenging combinatorial optimisation problem. The elastic net (Durbin & Willshaw, 1987) is an approximate method (called a continuation method) for solving this problem. The idea is to define a series of optimisation problems, the first of which is trivial to solve, the last of which is the full combinatorial problem, and for which intermediate problems are of intermediate difficulty. The later, harder, problems are solved with the knowledge of the solution to the earlier, easier ones. The full energy function $E(K, \tilde{\mathbf{W}})$ for the elastic net is given in the appendix. It depends on a parameter K, which governs the strength of the matching constraints and hence the difficulty of the optimisation problem. Here $\tilde{\mathbf{W}}$ is the collection of all the $\tilde{\mathbf{w}}_a$.

For a given value of K, the positions of the cortical points are changed to minimise $E(K, \tilde{\mathbf{W}})$ by a form of gradient descent. In practice, $\tilde{\mathbf{w}}_a$ are altered according to:

$$\Delta\tilde{\mathbf{w}}_a \propto -K\nabla_{\tilde{\mathbf{w}}_a} E(K, \tilde{\mathbf{W}}) \tag{9.16}$$

$$= \alpha \sum_{e,b} p[a|\tilde{\mathbf{u}}_b^e](\tilde{\mathbf{u}}_b^e - \tilde{\mathbf{w}}_a) + \beta K \sum_{a' \in \mathcal{N}_a} (\tilde{\mathbf{w}}_{a'} - \tilde{\mathbf{w}}_a) \tag{9.17}$$

Here,

$$p[a|\tilde{\mathbf{u}}] = \frac{e^{-|\tilde{\mathbf{u}}-\tilde{\mathbf{w}}_a|^2/2K^2}}{\sum_{a'} e^{-|\tilde{\mathbf{u}}-\tilde{\mathbf{w}}_{a'}|^2/2K^2}} \tag{9.18}$$

is the relative responsibility that cortical unit a takes for input $\tilde{\mathbf{u}}$. This is a soft form of the competition in equation 9.8, where the degree of softness is determined by parameter K. Whereas in the weight update equation 9.10 for the self-organising map, similarity between the cortical features of neighboring units a and a' is created by pulling them both towards the input $\tilde{\mathbf{u}}$ if they are sufficiently close to the winning unit a^*, in the elastic net update equation 9.17, similarity comes from pulling cortical feature vectors $\tilde{\mathbf{w}}_a$ and $\tilde{\mathbf{w}}_{a'}$ towards each other directly. Typically, rather more cortical points than input points are used (for one reason, to prevent local minima which might prevent some input cell from having a cortical match). However, this makes the already heuristic distance metric harder to interpret. Further, parameter K is usually annealed over the course of iterations.

Figure 9.14 shows possible outcomes from the use of the elastic net for ocular dominance. The solution in figure 9.14A is generally optimal, and corresponds to a division of the whole of V1 into just two ocular patches. However, the dynamical behavior of the elastic net prevents this solution from forming, since the 'line' of cortical cells always starts off by elongating in the direction of the principal component of the data. This is fortunate, since the *sub-optimal* solution in figure 9.14B is a closer match to the data and the maps such as figure 9.7C. If the global optimum of figure 9.14A is ruled out, then the optimal stripe width is proportional to the ratio l/d – the larger the separation between the retinas (modeling the case of decreased correlation between the inputs), the wider are the stripes (Goodhill & Willshaw, 1990). This, finally, is the explanation of the effects of strabismus, which reduces the correlation between the eyes, thereby increasing l and increasing the widths of the stripes.

Figure 9.13C shows the elastic net map in the same circumstance as figures 9.13A,B, for which a two dimensional input is mapped onto a one-dimensional output. The resulting solution is very close to the minimal wiring map for $s = 0.5$. It is not clear why finding local minima of cortical distortion should be at all related to the true goal of minimising cortical wiring.

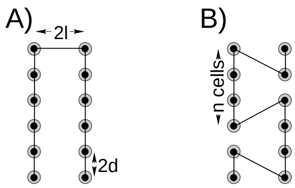

Figure 9.14: Ocular dominance. Two possible mappings in the case based on figure 9.12C in which the input points within each projection are distance $2d$ apart, and the two projections are distance $2l$ apart. A) This is generally the shortest path of all – and is equivalent to having half the cortex devoted to one eye and the other half to the other, with reversed topography in each. Although it is optimal, it almost never forms with the elastic net, because of dynamical constraints. B) Ocular dominance domains with a stripe width of $2n$ cells. Adapted from Goodhill & Willshaw (1990).

The elastic net has been used to study the large-scale, global, properties of the ocular dominance map as well as the small-scale ones that we have considered hitherto. Figure 9.15A shows a computer reconstruction of an almost complete, flattened pattern of ocular dominance stripes from a macaque. Some of the more obvious global properties of this map are the dark black spot in the center, which is due to the optic disk (the blind spot, which is innervated exclusively by the ipsilateral eye), the white patch at the medial (left) end which is called the monocular crescent, representing the part of the visual field that only the contralateral eye can see, the fact that the ocular dominance stripes tend to run perpendicularly into the boundaries of V1, and that in the representation of the foveal region (which is to the right of the blind spot) the stripes are thicker and significantly less regular than elsewhere.

Figure 9.15B shows the output of a simulation of the elastic net designed to match some of these characteristics (Goodhill, Bates & Montague, 1997). Arranging for this map required making the input correlations have spatial non-uniformities (foveal sampling of the visual world is more dense) and anisotropies that might arise from the way that the retinas grow and also from the elliptical shape of the cortex. The more densely packed input points are, the wider the stripes, since d in figure 9.14B decreases and so the ratio l/d increases. Anisotropies in the input correlations are modeled by making the input points $\tilde{\mathbf{u}}_a$ closer together (*ie* the activities to be more correlated) in one spatial dimension than the other. This tends to makes the stripes run in a direction perpendicular to the direction in which the correlations are highest and to be wider. The anisotropy in mapping two circular retinas onto an elliptical cortex tends to make the stripes run perpendicular to the short axis of the ellipse (Goodhill & Willshaw, 1994; Bauer, 1995). These influences are made to oppose each other in the foveal region, which tends to disrupt the regular order of the stripes there. Finally, the foveal region has more densely packed inputs, and so wider stripes. The blind spot and the monocular crescent are not modeled.

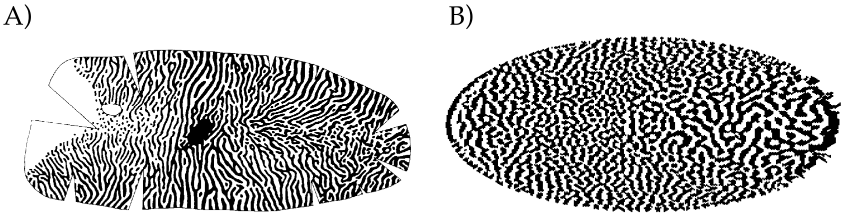

Figure 9.15: Global patterns of ocular dominance. A) Computer-aided reconstruction of the ocular dominance stripes in the right hemisphere of V1 in a macaque. Left is medial; right is lateral (the lateral end represents the fovea); up is dorsal and down ventral. B) Global elastic net model. The fovea has a graduated higher sampling density (making the stripes thicker) and anisotropic correlations (encouraging them to be horizontal); and the cortex is elliptical (encouraging stripes to be vertical). A) From LeVay *et al* 1985; B) From Goodhill, Bates & Montague (1997).

9.4 Orientation Selectivity

We have focused on the simple case of topography and ocular dominance. Adding in selectivity for orientation makes the map substantially more complicated, as seen in figure 9.1. To make things worse, the experimental relationship between orientation and topography is under question. Measuring the rate of change of orientation preference as one moves (with an electrode) along cortex, there is evidence that whereas orientation changes fastest when ocularity changes slowest (*ie* the pinwheels are at peaks of ocularity) and *vice-versa* (*ie* the linear zones for orientation occur at the boundaries of the ocular dominance stripes), the rate of change of ocularity and topography are positively, and not negatively correlated. Models have yet to address these data in full, or indeed others about the consequences on the orientation map of monocular deprivation, and we therefore only sketch how those we described above have been applied to add orientation selectivity.

Note that various models have been advanced to account for the orientation map without consideration of ocular dominance (von der Malsburg, 1973; Swindale, 1982; Obermayer, Ritter & Schulten, 1990; Miller, 1992; Wolf *et al*, 1993; Wolf & Geisel, 1998). These pose important unanswered questions about issues for the orientation map by itself, such as the theoretical and experimental stability of the orientation singularities – in many models, given enough iterations, the pinwheels mutually annihilate.

The main Hebbian, noise, and and feature-based competitive models that we have described have also all been used to model the fully coupled maps. For the Hebbian and noise models, one key issue is the way that the maps of ocular dominance and orientation are coupled (*ie* orientation changes fastest where ocularity changes most slowly) – the natural extension of the models to orientation leaves the selectivities uncoupled. For

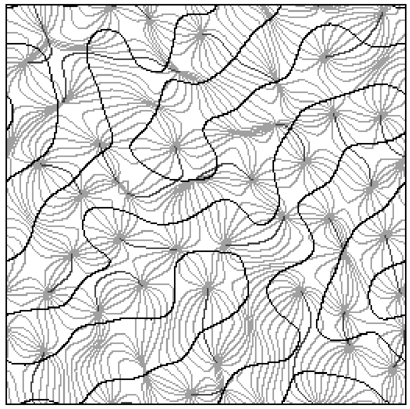

Figure 9.16: Ocular dominance and orientation map produced by the elastic net model. The thick black lines show the boundaries of the ocular dominance stripes; the thin grey lines are iso-orientation contours at intervals of 11.25°; the darker grey lines show the orientation 0° (this coloration is omitted in figure 9.1). From Erwin, Obermayer & Schulten (1995).

the Hebbian model (Erwin & Miller, 1997; Piepenbrock *et al*, 1997), it has been suggested that a form of staggered development might couple the maps, allowing the orientation map to develop first (for which there is some evidence). However the simple picture of development controlled by the principal eigenvectors of the correlation matrix ceases being true, and factors such as the early bias favoring input from the contralateral eye have yet to be incorporated. The noise models (Swindale, 1992; Rojer & Schwartz, 1990), since they are more abstract, can achieve the same effect more directly – by slowing down development of orientation at places where ocular dominance is strongest. If pinwheels are to occur somewhere, they will then prefer to avoid the regions where orientation development happens the fastest, since the development tends to smooth out the sort of local differences in orientation that define pinwheels.

The feature-based models (Durbin & Mitchison, 1990; Obermayer, Blasdel & Schulten, 1992) also capture the general relationships between the features – for instance, the orientation and ocularity maps formed by the elastic net model is shown in figure 9.16. Here, the features of the input are two dimensions of topographic position (x, y), one dimension z for ocularity (again taking the values $\pm Z$), and two dimensions for ocularity, coded as $R(\cos(\theta), \sin(\theta))$ for uniformly distributed θ. Here R models the effective strength of the feature of orientation, albeit in a rather abstract way. Figure 9.16 shows just ocular dominance and orientation selectivity.

Comparing figures 9.1 and 9.16, we see that the elastic net (and similarly the feature-based competitive model) is extremely successful in capturing the general form of the ocular dominance and orientation maps and also the relationship between the two. Understanding why the more abstract models should be more faithful to the data (though admittedly not the most recent data on the relationship between topography and orientation) than the less abstract ones is an important question for the future.

9.5 Appendix - The Elastic Net Energy Function

Development in the elastic net (section 9.3) is based on an energy function:

$$
\begin{aligned}
E(\zeta, \tilde{\mathbf{W}}) \quad = \quad & -\alpha K \sum_{e\in\{L,R\}} \sum_{b=1}^{N} \log \sum_{a=1}^{M} e^{-|\tilde{\mathbf{u}}_b^e - \tilde{\mathbf{w}}_a|^2/2K^2} + \\
& \beta \sum_{a} \sum_{a'\in\mathcal{N}_a} |\tilde{\mathbf{w}}_a - \tilde{\mathbf{w}}_{a'}|^2 ,
\end{aligned}
\tag{9.19}
$$

that balances two factors, the matching between the $(2N)$ input points $\tilde{\mathbf{u}}$ and the (M) cortical points $\tilde{\mathbf{w}}$, and the resulting distortion of the map. Here, α and β are constants, and K is a parameter. The matching constraints (which is what makes the optimisation problem of minimising $E(\zeta, \tilde{\mathbf{W}})$ hard) are imposed in the first term of $E(K, \tilde{\mathbf{W}})$; the requirement to minimise distortion is imposed in the second term. When K is large, the first term is very small, so the matching constraints are only weakly imposed. As K tends to 0, the cost of the first term tends to ∞ unless each input point has a cortical match within a distance of about K. The second term is just C^{EN}, the cost function in equation 9.15 that is of ultimate interest. There is an interaction between K and α to determine the relative significance of these terms. It is common to use the elastic net with either a fixed or annealed K. The case with annealed K is suitable for the case that there is a finite number of retinal points to match with the same (or a greater) number of cortical points, as in the simplification in figure 9.12, and as in the link with minimal wiring. The case with fixed K has been used as an abstraction of competitive Hebbian learning to model the case that there is a whole probability distribution over the input points, and therefore exact matches are impossible.

9.6 Further Reading and References

Further Reading

Erwin, E, Obermayer, K & Schulten, K (1995) Models of orientation and ocular dominance columns in the visual cortex: A critical comparison. *Neural Computation* **7**:425-468.

Flanagan, JG & Vanderhaeghen, P (1998) The ephrins and Eph receptors in neural development. *Annual Reviews of Neuroscience* **21**:309-345.

Goodhill, GJ & Löwel, S (1995) Theory meets experiment: Correlated neural activity helps determine ocular dominance column periodicity. *Trends in Neurosciences* **18**:437-439.

Hubener, M, Shoham, D, Grinvald, A & Bonhoeffer, T (1997) Spatial relationships among three columnar systems in cat area 17. *Journal of Neuroscience* **17**:9270-9284.

Katz, LC & Shatz, CJ (1996) Synaptic activity and the construction of cortical circuits. *Science* **274**:1133-1138.

Miller, KD (1996) Receptive fields and maps in the visual cortex: Models of ocular dominance and orientation columns. In E Domany, JL van Hemmen & K Schulten, editors, *Models of Neural Networks, III*. New York:Springer-Verlag, 55-78.

Murray, JD (1989) *Mathematical Biology*. New York:Springer, 548.

Peters, A & Jones, EG, editors (1984-1995) *Cerebral Cortex*. New York:Plenum Press.

Swindale, NV (1996) The development of topography in the visual cortex: A review of models. *Network: Computation in Neural Systems* **7**:161-247.

Tessier-Lavigne, M & Goodman, CS (1996) The molecular biology of axon guidance. *Science* **274**:1123-1133.

Turing, AM (1952) The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London* **B237**:37-72.

Udin, SB & Fawcett, JW (1988) Formation of topographic maps. *Annual Review of Neuroscience* **11**:289-327.

References

Bauer, H-U (1995) Development of oriented ocular dominance bands as a consequence of areal geometry. *Neural Computation* **7**:36-50.

Bauer, H-U, Brockmann, D & Geisel, T (1997) Analysis of ocular dominance pattern formation in a high-dimensional self-organizing-map model. *Network: Computation in Neural Systems* **8**:17-33.

Cherniak, C (1995) Neural component placement. *Trends in Neurosciences* **18**:522-527.

Cowey, A (1979) Cortical maps and visual perception. *Quarterly Journal of Experimental Psychology* **31**:1-17.

Crair, MC, Gillespie, DC & Stryker MP (1998) The role of visual experience in the development of columns in cat visual cortex. *Science* **279**:566-570.

Crair, MC, Ruthazer, ES, Gillespie, DC & Stryker, MP (1997a) Relationship between the ocular dominance and orientation maps in visual cortex of monocularly deprived cats. *Neuron* **19**307-318.

Crair, MC, Ruthazer, ES, Gillespie, DC & Stryker, MP (1997b) Ocular dominance peaks at pinwheel center singularities of the orientation map in cat visual cortex. *Journal of Neurophysiology* **77**:3381-3385.

Das, A & Gilbert, CD (1997) Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature* **387**:594-598.

Durbin, R & Mitchison, G (1990) A dimension reduction framework for cortical maps. *Nature* **343**:644-647.

Durbin, R & Willshaw, DJ (1987) An analogue approach to the travelling salesman problem using an elastic net method. *Nature* **326**:689-691.

Erwin, E (1994) *Structure and Development in the Macaque Primary Visual Pathway* PhD Thesis, University of Illinois, Urbana.

Erwin, E & Miller, KD (1997) Correlation-based development of combined orientation and ocular dominance maps: Analysis and simulation. *Journal of Neuroscience* submitted.

Gaze, RM (1970) *The Formation of Nerve Connections*. London:Academic Press.

Gödecke, I & Bonhoeffer, T (1996) Development of identical orientation maps for two eyes without common visual experience. *Nature* **379**:251-254.

Goodhill GJ (1993) Topography and ocular dominance: A model exploring positive correlations. *Biological Cybernetics* **69**:109-118.

Goodhill, GJ, Bates, KR & Montague, PR (1997) Influences on the global structure of cortical maps. *Proceedings of the Royal Society, Series B,* **264**:649-655.

Goodhill, GJ & Sejnowski, TJ (1996) A unifying objective function for topographic mappings. *Neural Computation* **9**:1291-1303.

Goodhill, GJ & Willshaw, DJ (1990) Application of the elastic net algorithm to the formation of ocular dominance stripes. *Network: Computation in Neural Systems* **1**:41-61.

Goodhill, GJ & Willshaw, DJ (1994) Elastic net model of ocular dominance: Overall stripe pattern and monocular deprivation. *Neural Computation* **6**:615-621.

Grossberg, S & Olson, SJ (1994) Rules for the cortical map of ocular dominance and orientation columns. *Neural Networks* **7**:883-894.

Hubel, DH & Wiesel, TN (1977) Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society of London, Series B: Biological Sciences* **B198**:1-59.

Hubener, M, Shoham, D, Grinvald, A & Bonhoeffer, T (1997). Spatial relationships among three columnar systems in cat area 17. *Journal of Neuroscience* **17**:9270-9284.

Kohonen, T (1982) Self-organised formation of topologically correct feature maps. *Biological Cybernetics* **43**:59-69.

Kohonen, T (1995) *Self-Organizing Maps*. Berlin, New York:Springer-Verlag.

LeVay, S, Connolly, M, Houde, J & Van Essen DC (1985) The complete pattern of ocular dominance stripes in the striate cortex and visual field of the macaque monkey. *Journal of Neuroscience* **5**:486-501.

von der Malsburg, C & Willshaw DJ (1976) A mechanism for producing continuous neural mappings: Ocular dominance stripes and ordered retino-tectal projections. *Experimental Brain Research* **1**:463-469.

von der Malsburg, C & Willshaw DJ (1977) How to label nerve cells so that they can interconnect in an ordered fashion. *Proceedings of the National Academy of Sciences of the United States of America* **74**:5176-5178.

Miller, KD (1992) Development of orientation columns via competition between on- and off-center inputs. *Neuroreport* **3**:73-76.

Miller, KD, Keller, JB & Stryker, MP (1989) Ocular dominance column development: Analysis and simulation. *Science* **245**:605-615.

Mitchison, G (1991) Neuronal branching patterns and the economy of cortical wiring. *Proceedings of the Royal Society of London, Series B: Biological Sciences* **245**:151-158.

Mitchison, G (1995) A type of duality between self-organizing maps and minimal wiring. *Neural Computation* **7**:25-35.

Montague, PR, Gally, JA & Edelman GM (1991) Spatial signaling in the development and function of neural connections. *Cerebral Cortex* **1**:199-220.

Niebur, E & Wörgötter, F (1993) Orientation column structures from first principles. In JA Bower, editor, *Computation and Neural Systems, 92*. Norwell, MA:Kluwer, chapter 62.

Obermayer, K & Blasdel, GG (1993) Geometry of orientation and ocular dominance columns in monkey striate cortex. *Journal of Neuroscience* **13**:4114-4129.

Obermayer, K, Blasdel, GG & Schulten, K (1992) Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Physical Review A* **45**:7568-7589.

Obermayer, K, Ritter, H & Schulten K (1990) A principle for the formation of the spatial structure of cortical feature maps. *Proceedings of the National Academy of Sciences of the United States of America* **87**:8345-8349.

Piepenbrock, C, Ritter, H & Obermayer, K (1997) The joint development of orientation and ocular dominance: Role of constraints. *Neural Computation* **9**:959-970.

Riesenhuber, M, Bauer, HU & Geisel, T (1996) Analyzing phase transitions in high-dimensional self-organizing maps. *Biological Cybernetics* **75**:397-407.

Rojer, AS & Schwartz, EL (1990) Cat and monkey cortical columnar patterns modeled by bandpass-filtered 2D white noise. *Biological Cybernetics* **62**:381-391.

Sirosh, J & Miikkulainen, R (1994) Cooperative self-organization of afferent and lateral connections in cortical maps. *Biological Cybernetics* **71**:65-78.

Sperry, RW (1963) Chemoaffinity in the orderly growth of nerve fiber patterns and connections. *Proceedings of the National Academy of Sciences of the United States of America* **50**:703-710.

Swindale, NV (1980) A model for the formation of ocular dominance stripes. *Proceedings of the Royal Society of London, Series B: Biological Sciences* **208**:243-264.

Swindale, NV (1982) A model for the formation of orientation columns. *Proceedings of the Royal Society of London, Series B: Biological Sciences* **215**:211-230.

Swindale, NV (1992) A model for the coordinated development of columnar systems in primate striate cortex. *Biological Cybernetics* **66**:217-230.

Swindale, NV (1996) Visual cortex: Looking into a Klein bottle. *Current Biology* **6**:776-779.

Swindale, NV, Matsubara, JA & Cynader MS (1987) Surface organization of orientation and direction selectivity in cat area 18. *Journal of Neuroscience* **7**:1414-1427.

Takeuchi, A & Amari, S (1979) Formation of topographic maps and columnar microstructures in nerve fields. *Biological Cybernetics* **35**:63-72.

Tal, D & Schwartz, EL (1997) Topological singularities in cortical orientation maps: The sign theorem correctly predicts orientation column patterns in primate striate cortex. *Network* **8**:229-238.

Tanaka, S (1990) Theory of self-organization of cortical maps: Mathematical framework. *Neural Networks* **3**:625-640.

Whitelaw, VA & Cowan, JD (1981) Specificity and plasticity of retinotectal connections: A computational model. *Journal of Neuroscience* **1**:1369-1387.

Willshaw, DJ & von der Malsburg, C (1976) How patterned neural connections can be set up by self-organization. *Proceedings of the Royal Society of London, Series B: Biological Sciences* **194**:431-445.

Willshaw, DJ & von der Malsburg, C (1979) A marker induction mechanism for the establishment of ordered neural mappings: Its application to the retinotectal problem. *Philosophical Transactions of the Royal Society* **B287**:203-243.

Wiskott, L & Sejnowski, TJ (1998) Constrained optimization for neural map formation: A unifying framework for weight growth. *Neural Computation* **10**671-716.

Wolf, F & Geisel, T (1998) Spontaneous pinwheel annihilation during visual development. *Nature* **395**:73-78.

Wolf, F, Pawelzik, K, Geisel, T, Kim, D-S & Bonhoeffer, T (1993) Optimal smoothness of orientation preference maps. In E Marder, J Rinzel & J Bower, editors, *Computation and Neural Systems, II* Boston, MA:Kluwer.

9.7 Exercises

1. Replicate figure 9.6 using (subtractively normalised) Hebbian learning as in chapter 8 to refine a one-dimensional topographic map. Study empirically the requirement for a topographic arbor function and/or starting condition for the weights, and also the condition that the Mexican hat (equation 9.7) should not be too narrow compared with the breadth of the correlations.

2. Extend this model to the case of ocular dominance by including input from a second projection with weak correlations (or anti-correlations) with the first. How common is the orderly progression of topography amidst the ocular dominance stripes?

3. Implement the weight-based competitive learning model of figure 9.9. Show that this model still generates monocular cells even if there are strong positive correlations between the eyes, and study the relationship between the between-eye correlations and the width of the stripes. How sensitive are the results to the rate at which the width σ of the cortical activity patterns is annealed. Repeat these experiments for a fixed σ – does the relationship between correlations and stripe width change?

4. Implement the feature-based competitive learning model (equations 9.11 and 9.12) and apply it to the full case of two-dimensions of topography and ocular dominance (using inputs $\tilde{\mathbf{u}} = (x, y, z)$ with $z = \pm Z$) for the case of two square input regions with wrap-around connections (*ie* two toroidal input regions). Compare the resulting ocular dominance stripes with parts of the full ocular dominance pattern shown in figure 9.15A). Study the effect of changing Z on the existence and width of the stripes.

5. Extend the model of exercise 4 to the case of orientation, using inputs $\tilde{\mathbf{u}} = (x, y, z, R\cos 2\theta, R\sin 2\theta)$, where R is fixed and θ has a uniform distribution in $[0, \pi)$. Show topographically ordered maps of ocular dominance and orientation preferences. Do orientation pinwheels occur at any time during the simulation as in figure 9.16? If the width σ of cortical activity patterns is kept fixed rather than being decreased, does this affect the result? What relationship is there between the maps of ocular and orientation preference?