Network: Comput. Neural Syst. **11** (2000) 153–168. Printed in the UK PII: S0954-898X(00)13504-3

Analysis of the elastic net model applied to the formation of ocular dominance and orientation columns

Geoffrey J Goodhill and Andrei Cimponeriu†

Department of Neuroscience and Georgetown Institute for Cognitive and Computational Sciences, Georgetown University Medical Center, 3900 Reservoir Road NW, Washington DC 20007, USA

E-mail: geoff@giccs.georgetown.edu and andrei@camelot.mssm.edu

Received 6 April 2000

Abstract. The development and structure of orientation (OR) and ocular dominance (OD) maps in the primary visual cortex of cats and monkeys can be modelled using the elastic net algorithm, which attempts to find an 'optimal' cortical representation of the input features. Here we analyse this behaviour in terms of parameters of the feature space. We derive expressions for the OR periodicity, and the first bifurcation point as a function of the annealing parameter using the methods of Durbin *et al* (Durbin R, Szeliski R and Yuille A 1989 *Neural Computation* **1** 348–58). We also investigate the effect of the relative order of OR and OD development on overall map structure. This analysis suggests that developmental order can be predicted from the final OR and OD periodicities. In conjunction with experimentally measured values for these periodicities, the model predicts that (i) in normal macaques OD develops first, (ii) in normal cats OR develops first and (iii) in strabismic cats OD develops first.

1. Introduction

Feature-based algorithms have been popular tools for modelling the formation of maps in primary visual cortex such as those for orientation and ocular dominance (OD) (reviewed in Erwin *et al* (1995), Swindale (1996)). These models represent the problem as a mapping between a low-dimensional feature space representing the visual input and a two-dimensional sheet of neurons representing the cortex. Different dimensions of the feature space correspond to different variables, e.g. two for spatial position, one for OD and two for orientation (OR). The feature space is populated with points representing combinations of variables, with a distribution that attempts to capture the structure of the visual environment. One model in this class is the elastic net (Durbin and Willshaw 1987, Durbin and Mitchison 1990, Goodhill and Willshaw 1990, 1994, Goodhill *et al* 1997). This trades off a matching constraint, which matches cortical cells to particular features, with a regularization constraint, which minimizes the sum of (squared) differences in feature selectivity between neighbouring cells in the cortex. The structure of the maps of OD and OR provided by the elastic net gives an extremely good match to some basic data regarding the structure of experimentally determined maps (Erwin *et al* 1995, Swindale 1996).

Two versions of the elastic net that can be applied to cortical mapping problems are 'annealed' and 'non-annealed'. In the annealed version a parameter k is gradually reduced

† Present address: The Mount Sinai School of Medicine, One Gustave L Levy Place, Department of Ophthalmology, Box 1183, New York 10029, USA.

0954-898X/00/020153+16$30.00

over time, where k is interpreted as the broadness of tuning of cortical cells for particular features. Initially cortical cells are broadly tuned, i.e. respond to many different orientations and spatial positions in both eyes, but as $k \to 0$ cortical cells become specialized for just a single feature combination. This implements a form of deterministic annealing, an efficient strategy for solving combinatorial optimization problems such as the one of minimizing the differences in feature selectivity between neighbours (Rose 1998). As the net develops with decreasing k, it passes through a set of bifurcation points at certain critical values of k (Durbin *et al* 1989), each corresponding to a point at which the cortex makes a strong commitment to selectivity along a particular feature dimension (e.g. space, OD, OR). By varying the parameter values it is possible to make the net pass through these bifurcation points in different sequences, and this influences the final structure of the maps produced (Hoffsümmer *et al* 1995, 1996). In particular, the OD map is more similar to that seen in monkeys if it develops before the orientation map, and more similar to that seen in cats if it develops after the orientation map (Hoffsümmer *et al* 1995, 1996). In the alternative non-annealed version of the elastic net, k is held fixed at a low value, below any of the critical values, so that expansion commences along all feature dimensions immediately. A much stronger initial topographic bias is required to produce a final topographic map in this version compared with the annealed version. Also, since k does not tend to zero thus freezing the maps, they can continue to evolve for a very long time: in this regime Wolf and Geisel (1998) identified the phenomenon of 'pinwheel annihilation' in OR maps where pairs of pinwheels can collide and disappear. However, current experimental data suggest that once OR and OD maps have initially appeared their structure is stable over subsequent development (Chapman *et al* 1996).

In this paper we analyse the annealed version of the elastic net as applied to the formation of spatial position, OD and OR maps to obtain a deeper understanding of how the sequence of bifurcations affects map structure. (i) We extend a simple optimization argument previously used for the OD map to predict the periodicity of the OR map, and show that it provides a good match with simulation results. (ii) We calculate the first bifurcation point as k decreases in terms of parameters of the feature space when feature points are arranged on a regular grid, using the methods of Durbin *et al* (1989). Further bifurcation points are determined by simulation. (iii) We look more closely at the differences in map structure as a function of developmental order. The periodicity of the OR and OD maps is not strongly affected. However, the distribution of the intersection angles between the OR and OD columns does depend on developmental order. We investigate this by making quantitative statistical comparisons between these distributions as the parameters of the feature space are varied. (iv) We show that developmental order can be predicted in the elastic net model based on the ratio of final periodicities in the map, and make predictions about the relative order of development in different species under normal and strabismic rearing conditions.

2. The elastic net model of cortical mapping

The elastic net was first proposed as a biologically motivated method for solving combinatorial optimization problems such as the travelling salesman problem (TSP) (Durbin and Willshaw 1987). As applied to the development of cortical maps, it consists of a set of $\boldsymbol{y}_j$, $j = 1 \ldots N$ vectors, modelling the receptive fields of cells in the primary visual cortex (V1). The net learns a set of $\boldsymbol{x}_i$, $i = 1 \ldots M$ prototypes, which represent the visual input. Both the prototypes and the cortical cell vectors consist of several features, specifically the x and y retinal position, the ocularity and the OR preferred angle and 'strength' as introduced by Swindale (1982).

Learning is performed by minimizing an energy function E given by

$$E = -\alpha k \sum_i \log \sum_j \Phi(\boldsymbol{x}_i, \boldsymbol{y}_j, k) + \frac{\beta}{2} \sum_{j'} \|\boldsymbol{y}_{j'} - \boldsymbol{y}_j\|^2 \tag{1}$$

where α and β are constants, the j' are the neighbours of j in the cortex, and Φ is given by

$$\Phi(\boldsymbol{x}_i, \boldsymbol{y}_j, k) = \exp\left(\frac{-\|\boldsymbol{x}_i - \boldsymbol{y}_j\|^2}{2k^2}\right). \tag{2}$$

The parameter k gives the extent of the neighbourhood in the input space over which each $\boldsymbol{y}_j$ is excited, i.e. the extent of its receptive field. In the 'annealed' version of the elastic net model, k is gradually reduced; since it acts as a scale for the difference between $\boldsymbol{x}_i$ and $\boldsymbol{y}_j$ features, this can be thought of as increasing the power of a magnifying glass through which the net's vectors can be distinguished from the prototypes; the smaller k, the greater the magnification, and the smaller the distance between them, minimized by the relaxation of E. To solve the TSP, $M \leqslant N$; in this case, when k becomes very small, each $\boldsymbol{y}_j$ is matched to an $\boldsymbol{x}_i$ (Durbin *et al* 1989, Simmen 1991). Meanwhile, the 'length' of the cortical map, given by the second term in the energy function, is kept small, requiring that neighbouring cortical cells learn similar features. The elastic net is closely related to other models based on local learning rules (Simic 1990, Yuille 1990, Dayan 1993, Yuille *et al* 1996), and implements a form of deterministic annealing (Rose 1998).

Two possible ways in which the prototypes can be distributed is randomly, sampling from a continous distribution of feature values (e.g. Wolf and Geisel 1998), or regularly, where feature values are evenly spaced (e.g. Goodhill and Willshaw 1990, 1994). Since the interpretation of the input space is somewhat abstract, it is hard to give biological justifications for choosing one type of distribution over the other. Here we use a regular distribution, and derive expressions for both the periodicity of the maps and how the order of development depends on the feature space parameters. The prototypes are placed on a five-dimensional regular grid. The five dimensions correspond to the x and y positions, evenly spaced with a step d to a total of unit length (i.e. n = $1 + 1/d$ positions in each direction), OD $\in \{-l, +l\}$, and $r\sin(\theta)$, $r\cos(\theta)$, where θ takes m evenly spaced values on $[0, 2\pi)$, and r is the 'strength' of OR tuning (since orientation preference is π periodic, we use the standard device of doubling this angle to give an appropriate 2π periodicity). All these components are mutually orthogonal (for a discussion see Goodhill *et al* (1996)).

3. Periodicity of the orientation map

3.1. Analysis

We previously derived a simple expression for the expected periodicity of OD columns given by the elastic net, by minimizing the sum of squared lengths of the distances in feature space between neighbouring cortical cells (Goodhill and Willshaw 1990, Goodhill 1992). A similarly simple calculation can be provided for the expected periodicity of orientation columns. Consider a mapping between a one-dimensional cortex and a cylinder representing m orientations separated by an angle of $2\pi/m$ at each position (see figure 1). A class of mappings can be defined where the cortex progresses through n spatial positions for one orientation, then doubles back to pick up the next orientation for the same n spatial positions, and so on until it has wrapped once round the cylinder. It then moves to the $(n + 1)$th spatial position and progresses around the same cycle again, and so on until it reaches the end of the cylinder (see figure 1). If there are a total of N spatial positions, then there are N/n complete

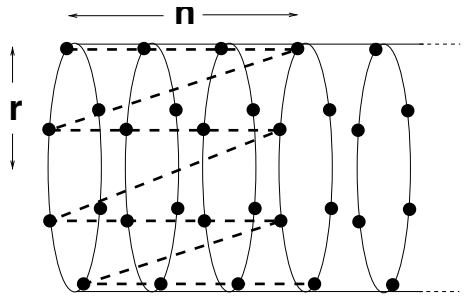

Figure 1. One possible type of mapping between a cylindrical feature space representing orientation and one spatial dimension, and a one-dimensional elastic net (solid circles, prototypes; dashed line, elastic net; solid lines, to aid visualization). Here the net forms isorientation domains of smoothly varying orientation, each of which traverses n units in the spatial dimension.

cycles through all orientations. Assuming for simplicity this is integer, the total sum-of-squares length L of a map in this class is given by

$$L = \frac{N}{n}\left[m(n-1)d^2 + (m-1)\left([(n-1)d]^2 + \left[2r\sin\left(\frac{\pi}{m}\right)\right]^2\right) + d^2\right]$$

which simplifies to

$$\frac{L}{N} = d^2\left[n(m-1) + 2 - m + \frac{4(m-1)r^2}{nd^2}\sin^2\left(\frac{\pi}{m}\right)\right]. \tag{3}$$

Note that the 'spiral' map, where all orientations are traversed at one spatial position before moving to the next spatial position, is a special case of this when $n = 1$. Since the elastic net was conceived as an algorithm for minimizing tour length in the TSP it should find the minimum of L. L is minimized with respect to n when

$$n = \frac{2r}{d}\sin\left(\frac{\pi}{m}\right).$$

If m is sufficiently large so that $\sin(\frac{\pi}{m}) \approx \frac{\pi}{m}$ ($m \geqslant 6$) then

$$n = \frac{2\pi r}{md}.$$

This gives the size of one iso-orientation patch. Multiplying by m gives the overall wavelength for orientation:

$$\lambda_{\mathrm{or}} = \frac{2\pi r}{d} \tag{4}$$

which is, appropriately, independent of the quantization factor m. This has an analogous form to the expression for the optimal OD wavelength (Goodhill 1992):

$$\lambda_{\mathrm{od}} = \frac{8l}{d}. \tag{5}$$

(Note that the expression given by Goodhill (1992) was $\frac{2l}{d}$; this was for one OD patch rather than a complete cycle, and there OD $\in \{0, l\}$ rather than $\{-l, l\}$ as here.) Since in the model the optimal OR and OD wavelengths exist in different dimensions they are completely decoupled, and in principle any pairing of different orientation and OD wavelengths can be obtained by suitable choice of l and r.

The above analysis considers only one class of mapping. At least one other type of map exists which is of shorter 'length'. In the OD case this is the map that traverses one eye in the forward direction followed by the other eye in the reverse direction, which is shorter than the periodic map (Goodhill and Willshaw 1990, Goodhill 1992). The analogous map in the OR case is where the cortex traverses all spatial positions for one orientation, moves to the next

orientation, then travels back through all spatial positions in the reverse direction, and so on. If m is odd the two ends of the cortex represent opposite ends of visual space, and if m is even they represent the same end. The total sum-of-squares length of this path L^* is approximately (for large N)

$$L^* = Nmd^2.$$

Compare this to the expression for L in equation (3) (approximated for $m \geqslant 6$) when the optimal $n = 2\pi r/md$ is substituted:

$$L = Nmd^2(2n - 1).$$

Since in this class of mappings $n \geqslant 1$ then $L^* < L$ always. However, such a solution is at odds with experimental data, and is generally not found by the elastic net. We return to this issue in the discussion.

3.2. Simulations

Figure 2 shows the periodicities of the OR and OD maps produced by elastic net simulations as a function of r and l compared with the predictions above (typical OD and OR maps can be seen in figure 4). As expected, λ_{or} is independent of l and λ_{od} is independent of r. Although the expected linear relationships between λ_{or} and r and λ_{od} and l are not exactly followed, the predictions still work surprisingly well given the simplicity of the one-dimensional optimization argument that produced them. In both cases the actual curve crosses the predicted curve at the point at which the OD and OR maps form simultaneously: for small l orientation forms first, and for small r OD forms first. We now discuss the bifurcation sequence in more detail.

4. Bifurcation points

4.1. Analysis

Durbin *et al* (1989) analysed the dynamics of the elastic net as applied to the TSP. They showed that for a sufficiently large value of k the energy function has one minimum, with all points in the elastic net at the origin. As k is reduced the energy function bifurcates. These bifurcation points occur when the Hessian matrix of the energy function becomes singular, i.e. has a zero eigenvalue. At the first bifurcation the net starts to expand along the direction of the principal eigenvector of the covariance matrix of the set of prototypes; a first map has emerged. As k is reduced further the energy function bifurcates again and a second map emerges, and so on. The full Hessian is complicated; however, Durbin *et al* (1989) showed that when the net lies entirely at the origin, $\{\boldsymbol{y}_j = 0\}_{j=1\ldots N}$, the Hessian has a much simpler form and the prototype distribution enters only via its covariance matrix. The value of k for which the net breaks out from the origin can be shown to be approximately $k = \sqrt{\lambda_{\max}}$, where $\lambda_{\max}$ is the largest eigenvalue of the city/prototype convariance matrix. This relation also follows from a more general consideration of deterministic annealing (Rose 1998).

For the regular distribution of the prototypes the covariance matrix and its eigenvalues can be calculated analytically in terms of d, l and r, and thus values derived for the critical k value at which the first map starts to expand. Because of the orthogonality of the components

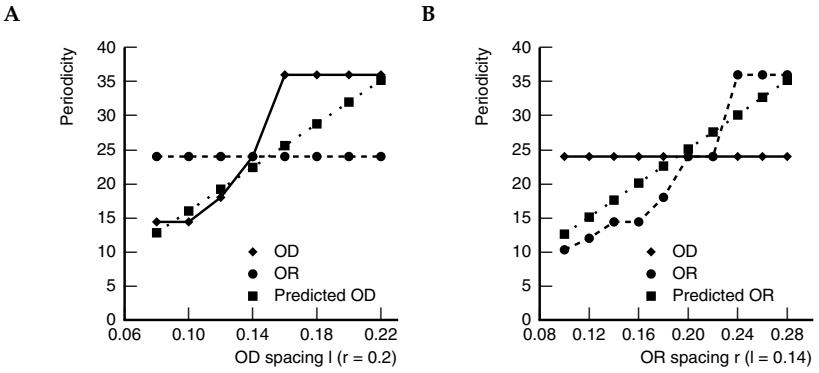

Figure 2. (*a*) Periodicity of the OD and OR maps as a function of l, the spacing in the ocularity dimension ($r = 0.20$). (*b*) Periodicity of the OR and OD maps as a function of r, the spacing in the orientation dimension ($l = 0.14$). Periodicity was taken to be the position of peak power in the Fourier power spectra of the patterns averaged over all directions; the actual periodicities show discontinuous jumps due to the limited number of integer wavelengths available in the finite cortical grid. Three simulations with different random seeds were performed for each parameter set, and these gave the same result in each case. Simulation parameters: $\alpha = 0.2$, $\beta = 2.0$, k was started at an initial value of 0.5 and was subsequently reduced by multiplying it by 0.99 after each iteration, $d = 0.05$, $m = 6$, $r \in [0.08, 0.28]$, $l \in [0.08, 0.2]$. The cortical sheet was a square array of size 72 by 72 cells, so that there were roughly the same number of cortical cells (5184) as prototype vectors (5292). The neighbourhood of a cortical cell was taken to be its four nearest neighbours in the cortical array. The net was started with random OR and OD components and a crude retinotopy, and the simulation was terminated at 400 iterations ($k \approx 0.01$), by which time the maps were stable. Open boundary conditions were used for the cortical sheet.

the covariance matrix is diagonal, and can be calculated to be

$$C = \begin{pmatrix} \frac{n^2 d_x^2}{12} & 0 & 0 & 0 & 0 \\ 0 & \frac{n^2 d_y^2}{12} & 0 & 0 & 0 \\ 0 & 0 & \frac{r^2}{2} & 0 & 0 \\ 0 & 0 & 0 & \frac{r^2}{2} & 0 \\ 0 & 0 & 0 & 0 & l^2 \end{pmatrix} \tag{6}$$

where the first two rows/columns give the variance in the x, y directions, the next two the variance along the polar coordinates of the (r, θ) plane, and the last the variance along the OD direction. For generality we have included the possibility that the spacing of feature points along the x and y components may be different (cf Goodhill *et al* 1997). The critical values k_{xy}, k_{or} and k_{od} at which initial movement from the origin occurs along each of these dimensions are thus

$$\begin{aligned} k_{xy} &= \frac{nd}{\sqrt{12}} \\ k_{\mathrm{or}} &= \frac{r}{\sqrt{2}} \\ k_{\mathrm{od}} &= l. \end{aligned} \tag{7}$$

Note that if the spatial dimensions are of unit length then $nd \approx 1$ and k_{xy} becomes $\sqrt{\frac{1}{12}} = 0.29$. Our previous applications of the elastic net to the formation of OD columns have usually started at $k_{\mathrm{init}} = 0.2$ with $l < k_{\mathrm{init}}$ (Goodhill and Willshaw 1990, 1994, Goodhill *et al* 1997). Thus, as expected, in those cases we saw an immediate expansion along the spatial dimensions without

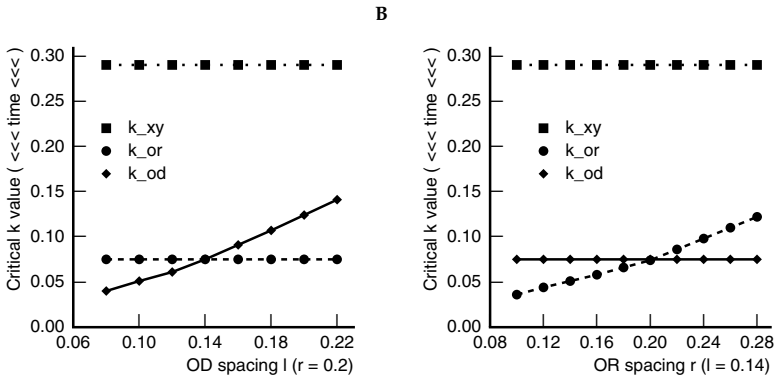

Figure 3. The dependence of the critical values of k on feature space parameters. (*a*) Varying OD parameter l with $r = 0.20$. (*b*) Varying OR parameter r with $l = 0.14$. Simulation details as in figure 2.

expansion along the OD dimension. In addition, these values allow conditions to be derived on d, r and l so as to ensure that one particular feature map develops first. Biologically it is clear that the topographic map develops first, thus the biologically appropriate parameter regime (in the limit of small d) is $r < 0.41$ and $l < 0.29$. Once it is true that $\{\boldsymbol{y}_j \neq 0\}_{j=1\ldots N}$ the subsequent null eigenvalues of the Hessian are no longer easy to compute analytically. However, the bifurcation points can be studied by simulation.

4.2. Simulations

As predicted analytically, during a simulation with k initially large enough, the net initially collapsed to a point at the centre of the feature space. We examined the development of the subsequent maps by observing each feature separately. As k decreased to its first critical value the retinotopic map started to emerge at $k = 0.29$, as predicted by equation (7). We also ran simulations in the non-biological regime where the OR or OD maps form before the retinotopic map. As predicted by our calculations, whereas if $r < 0.41$ and $l < 0.29$ the retinotopic map emerged first, if $r > 0.41$ and $l < r/\sqrt{2}$ the OR map emerged first (i.e. each cortical receptive field had a different orientation but the same retinotopic location), and if $l > 0.29$ and $r < l\sqrt{2}$ the OD map emerged first.

We investigated the dependence of the critical values of k on the parameters of the feature space. Figure 3 shows the dependence k_{or} and k_{od} on l and on r after the retinotopic map has formed. Since on the ordinates ks are represented as increasing, while in simulations k decreases, these plots should be read from top to bottom: for example on the left figure, for a given value of l, e.g. 0.10, the first map to emerge is the retinotopic one, at $k_{xy} = 0.29$, the second is the OR map, at $k_{\mathrm{or}} = 0.075$, and the last is the OD map, at $k_{\mathrm{od}} = 0.051$. These simulation values for k_{or} and k_{od} do not match the numerical values predicted by equation (7), since the predictions are valid only if OR and OD respectively is the first map to emerge. However, surprisingly the analysis still correctly predicts several more general aspects of these graphs. Firstly, in both analysis and simulations k_{or} depends only on r and k_{od} depends only on l. Secondly, in both analysis and simulations this is a linear dependence, though with a change in the slope in the simulations at (OR) or near (OD) the cross-over point. Thirdly, in the simulations the slope of the k_{od} versus l line is $\approx$0.5 before cross-over and $\approx$0.8 after cross-over, while the slope of the k_{or} versus r line is $\approx$0.35 before cross-over and $\approx$0.6 after

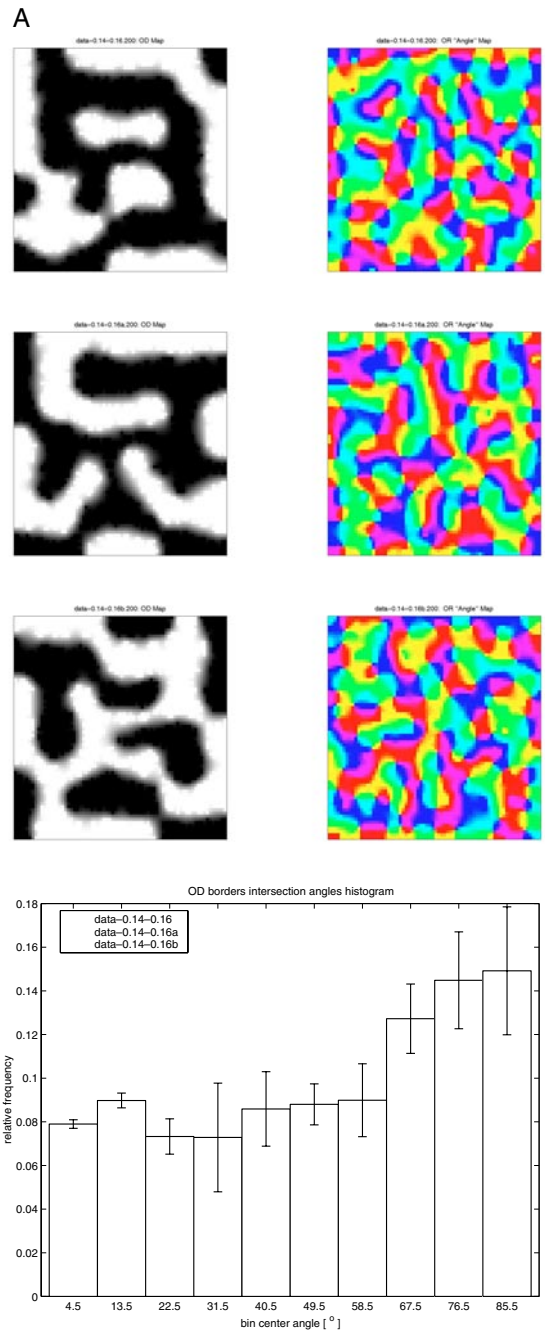

Figure 4. Map structure depends on developmental order. (*a*) OD develops before OR, (*b*) OR develops before OD. Above are the final OD and OR maps for three simulations in each case, each starting from different random initial conditions. Below are histograms of intersection angles between OD and OR columns. A clear difference can be seen between the two sets of simulations. Parameters as in figures 2, 3 except: annealing rate = 0.98, $l = 0.14$, $r = 0.16$ (*a*), $r = 0.24$ (*b*). Note that OR columns are wider in (*b*) than (*a*) because r is larger.

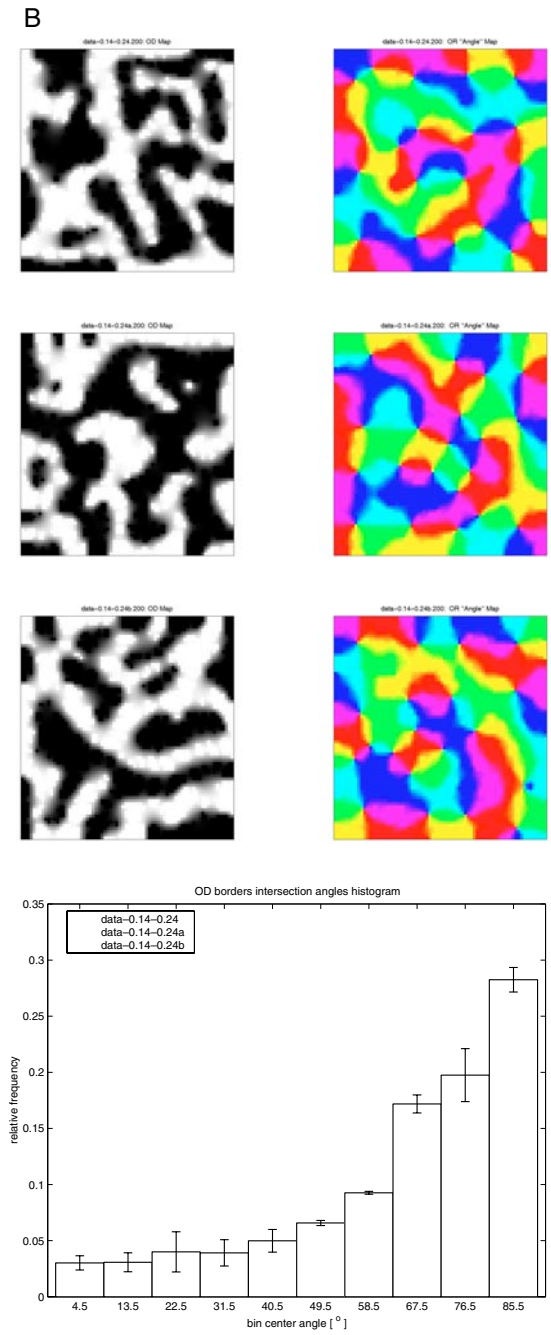

Figure 4. (Continued)

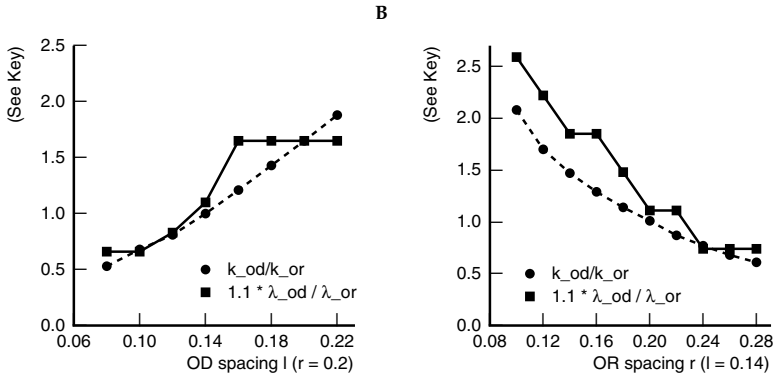

Figure 5. Assessment of the quality of the approximation $\frac{k_{od}}{k_{or}} = \frac{\pi}{\sqrt{8}}\frac{\lambda_{od}}{\lambda_{or}}$. (*a*) Varying OD parameter l with $r = 0.20$. (*b*) Varying OR parameter r with $l = 0.14$ (values of k_{od}, k_{or}, λ_{od}, λ_{or} taken from figures 2 and 3). Note that, although the curves sometimes diverge, in general if one curve is greater than unity then so is the other, and *vice versa*.

cross-over. Thus, the ratio of the k_{od} and k_{or} slopes both before and after cross-over is $\approx\sqrt{2}$, as in equation (7).

5. Map structure as a function of developmental order

5.1. Simulations

Figure 4 compares results for two sets of simulations: one set where OD develops before OR (*a*) and one set where OR develops before OD (*b*). In each case three simulations are shown starting from different initial conditions, with all other parameters unchanged. An immediate qualitative impression, first noted by Hoffsümmer *et al* (1995, 1996), is that when OD develops first (*a*), OD columns are regular and stripelike, whereas when OR develops first (*b*), OD columns are less regular and patchy. Qualitatively, it appears that OD column periodicity is also less sharply defined in the latter compared with former case: this shows up quantitatively as a broader peak in the power spectrum of the Fourier transform (data not shown).

5.2. Periodicity as a predictor of developmental order

Comparing equations (4), (5) and (7), the following relationship can be derived:

$$\frac{k_{od}}{k_{or}} = \frac{\pi}{\sqrt{8}}\frac{\lambda_{od}}{\lambda_{or}} \approx 1.1\frac{\lambda_{od}}{\lambda_{or}}. \tag{8}$$

The expressions for k_{od} and k_{or} are only valid when these are the first bifurcation points. However, we observed by simulation earlier that the ratio k_{od}/k_{or} remains approximately as given by equation (7), even when OD and OR are not the first bifurcations. Figure 5 plots the empirically determined values of $\frac{k_{od}}{k_{or}}$ and $\frac{\pi}{\sqrt{8}}\frac{\lambda_{od}}{\lambda_{or}}$ as a function of r and l, and demonstrates that in practice the relation between $\frac{k_{od}}{k_{or}}$ and $\frac{\lambda_{od}}{\lambda_{or}}$ of equation (8) is roughly valid even when OR or OD is not the first bifurcation. Thus the order of bifurcations in a simulation can be roughly predicted from whether $\frac{\pi}{\sqrt{8}}\frac{\lambda_{od}}{\lambda_{or}}$ is greater than or less than one. We return to the biological consequences of this result in the discussion. A link between bifurcation points and resulting

wavelengths in the elastic net model was first noted by Hoffsümmer *et al* (1995, 1996) (see the discussion).

5.3. Intersection angles

A useful way to quantify the overall joint structure of OD and OR maps is to calculate the distribution of intersection angles (DIA) between the two sets of columns. The intersection angle ϕ between OD and OR columns at a particular point is given by the angle between the local gradient vectors of the two maps at that point (Hübener *et al* 1997, Löwel *et al* 1998):

$$\phi = \cos^{-1}\left(\frac{\nabla(\mathrm{OD}) \cdot \nabla(\mathrm{OR})}{\|\nabla(\mathrm{OD})\|\|\nabla(\mathrm{OR})\|}\right).$$

This equation works for continuous data; however in our case OR and OD are defined on a lattice, and the partial derivatives that form the gradient vector do not exist. They can be replaced with finite differences, but since OR is represented by its principal value while being a periodic function, discontinuities appear at transitions from $-\pi$ to π. To avoid the effect of this discontinuity, a smooth periodic function of OR can be taken, e.g. sin(OR). This gives the formula we actually used in computing intersection angles:

$$\phi = \cos^{-1}\left(\frac{\nabla \mathrm{OD} \cdot \nabla \sin(\mathrm{OR})}{\|\nabla \mathrm{OD}\|\|\nabla \sin(\mathrm{OR})\|}\right)$$

The effect of the sine disappears when normalizing $\nabla \sin(\mathrm{OR})$ by $\|\nabla \sin(\mathrm{OR})\|$. Intersection angles were calculated only for pixels on the borders of OD columns: in the simulations $\|\nabla \mathrm{OD}\|$ is very small away from the borders.

Figure 4 shows (averaged) DIAs for each set of simulations. They look different, and perhaps surprisingly there appears to be a greater tendency towards orthogonality in the case where OR developed first and OD columns are more patchy than the case where OD developed first and OD columns are more stripelike. In order to more rigorously test the hypothesis that a change in developmental order causes a significant change in the DIA, we used the Kolmogorov–Smirnov (KS) test to compare DIAs arising from different parameters. Given two random samples, the KS test gives a confidence level for the hypothesis that the two samples were drawn from the same distribution, so that a result of less than 0.05 means the two samples probably originated from different distributions. Fixing the OD parameter $l = 0.14$, we ran three simulations with different starting conditions for each of the r values $\{0.12, 0.14, \ldots, 0.26\}$ (parameters as in figure 4). DIAs were calculated for each simulation and compared between values of r (note that the complete DIA was used, not the binned version shown in figure 4). For the same r value there were three comparisons, whereas for different r values there were $3 \times 3 = 9$ comparisons. The result of the KS test was averaged over each set of comparisons, and the resulting mean and standard deviation are plotted in figure 6. Note that there is a large variability, and occasionally two simulations varying only in initial conditions can be significantly different. However, despite this large variability, there is a clear discontinuity at the point at which the order of map development reverses ($r = 0.2$): the DIAs for the case where the OR map forms first are quite different from the DIAs where the OD map forms first, confirming the qualitative impression of figure 4.

A possible alternative hypothesis is that it is the change in the ratio of column periodicities that causes this qualitative change in the relationship between the two maps. If this were true then the natural transition point would be when the two periodicities were equal. Comparing equations (4) and (5) and setting $l = 0.14$ as in figure 4 then predicts that this transition should occur at $r = 0.18$. Figure 6 shows rather that the transition occurs at $r = 0.20$, when developmental order changes.

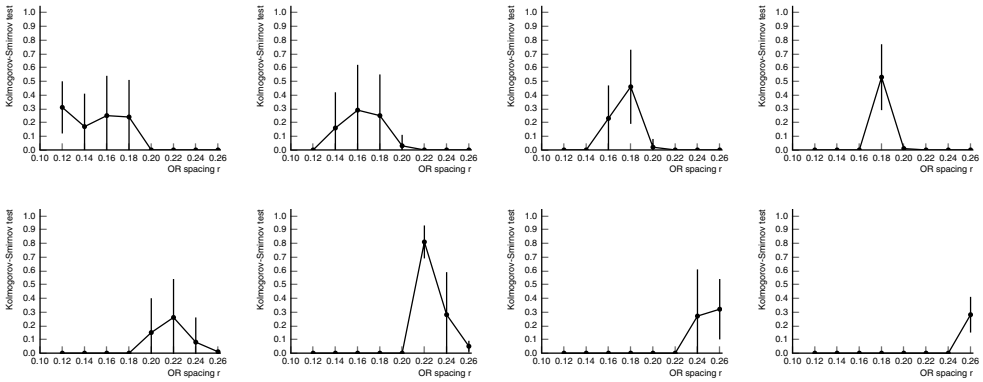

Figure 6. Statistical analysis of map structure as developmental order changes. The mean and standard deviation of the KS test between DIAs are plotted against r for eight different sets of three simulations. Top left: the case $r = 0.12$ is compared with itself and with different r values. Next on the right $r = 0.14$ is compared, and so on. A clear transition point is evident at $r = 0.20$, where developmental order changes (see figure 3): DIAs for pairs of simulations with different r both above or below $r = 0.20$ tend not to be significantly different, while pairs of simulations with r on opposite sides of 0.20 do tend to be significantly different.

6. Discussion

Relation to previous work. In two brief earlier contributions, Hoffsümmer *et al* (1995, 1996) studied a continuous version of the elastic net from the perspective of dynamics rather than optimization. By performing a linear stability analysis from a homogeneous state where the topographic but not OR and OD maps have formed, they calculated the following expressions for the critical values k_i and resulting wavelengths λ_i (where i indexes a dimension such as OR or OR):

$$k_i = \sqrt{\langle S_i^2 \rangle - \eta - \eta \log(\langle S_i^2 \rangle / \eta)}$$

$$\lambda_i = \frac{2\pi k_i}{\sqrt{\log(\langle S_i^2 \rangle / \eta)}}$$

where $\langle S_i^2 \rangle$ is the variance of the stimulus distribution along dimension i, and η is analogous (though not equal) to our parameter β. Note that for η small, $k_i \approx \sqrt{\langle S_i^2 \rangle}$, which is the same as equation (7) for the particular stimulus distribution we considered. However, the equation for λ_i is rather different from our equations (4) and (5). From these expressions it follows that

$$\frac{k_i}{k_j} = \frac{\log(\langle S_i^2 \rangle / \eta)}{\log(\langle S_j^2 \rangle / \eta)} \frac{\lambda_i}{\lambda_j}$$

which can be compared with our equation (8). However, the approach of the current paper has been different. Firstly, we have calculated periodicities from an optimization rather than dynamical perspective, and thus equations (4) and (5) depend only on the geometry of the feature space and not on the dynamical parameter β. Secondly we have used arguments based on the eigenvalues of the Hessian of the energy function to calculate the critical points. (Note that although there is a slight dependence on β in the expression of Durbin *et al* (1989) for the smallest eigenvalue of the Hessian, this is negligible and has been ignored here.) It is intriguing that these two quite different perspectives yield somewhat similar results.

Optimal periodicities. To calculate 'optimal' periodicities we have optimized over only a small class of easily parametrized mappings. Clearly many other types of map are possible, and we have shown that at least one of these is shorter. There are two principal reasons why we have not considered other classes of map here. Firstly, our approach still yields a reasonable approximation to simulation results, and the shorter mapping discussed above is generally not found by the elastic net. When expansion along the dimensions representing spatial position forms first, the net is then unable to completely double back on itself. It is likely that in reality earlier, activity-independent, mechanisms (reviewed by Goodhill and Richards (1999)) prevent such solutions from occurring (for further discussion see Goodhill (1998)). Secondly, we have not so far been able to conceive of any other classes of periodic mappings which satisfy the constraint that topography never systematically reverses: although experimentally it is clear that topography makes reverse jumps, so far it has not been shown to vary *continuously* in the reverse direction over several neighbouring cortical neurons in V1 (though it may in V2 Roe and T'so (1995)). Das and Gilbert (1997) plotted the magnitude of the rate of change of topography against that of orientation and found a positive correlation. The equivalent plot for the class of mappings shown in figure 1 consists of just two points: one at the position $(d, 0)$ corresponding to the horizontal segments of the cortex in figure 1, and one at the position $((n-1)d, 2\pi/m)$, corresponding to the diagonal segments (assuming cortical cells are one unit apart). Although this does give a positive correlation, in two-dimensional elastic net simulations there is a wider distribution of points, and these show a negative correlation (S-M Wu and G J Goodhill, unpublished data). Mitchison and Swindale (1999) have recently proposed a modification of Kohonen's (1982) feature-based learning rule which can reproduce the positive correlation of Das and Gilbert. However, this algorithm also places pinwheel centres at the borders of OD columns rather than at the centres, unlike experimental data and more standard feature-based algorithms (Erwin *et al* 1995). The fact that in reality the rates of change of some variables are positively correlated whereas others are negatively correlated suggests that models such as the elastic net, which treat all variables uniformly, may need to be modified.

Relation to experimental data. The elastic net is an abstract algorithm and its elements do not map 1–1 to biology. However, despite its abstract form it does have explanatory power. For instance, it predicted the increase in width of OD columns when between-eye correlations are reduced (Goodhill and Willshaw 1990, Goodhill 1992) before this had been observed experimentally (Löwel 1994, Tieman and Tumosa 1997, Horton *et al* 1999, though see also Crawford 1998, Murphy *et al* 1998).

In ferrets optical imaging studies suggest that OR develops slightly before OD (Chapman *et al* 1996). However in cats and macaques experimental data have not definitively established the precise ordering of OR and OD map development. In monkeys the OD map emerges before birth (Horton and Hocking 1996a) and the OR map is present very soon after birth (Wiesel and Hubel 1974, Blasdel *et al* 1995), though it may emerge earlier. In cat, although it was originally thought that the OD map emerges starting at about 3 weeks after birth (LeVay *et al* 1978), more recent optical imaging data suggests that both OR and OD maps are present as early as 2 weeks after birth (Crair *et al* 1998).

Table 1 summarizes values from the experimental literature regarding OR and OD column periodicities. Unfortunately, no study has yet systematically examined both OR and OD periodicity in both normal and strabismic animals. A notable feature of the table is the large variability in wavelengths between individuals of the same species, and the rough correlation between OR and OD wavelengths within each individual. From the results of Obermayer and Blasdel (1993) and Löwel *et al* (1988), predictions of the elastic net model are that in

Table 1. Experimentally measured mean OR and OD column wavelengths (in micrometres), and their ratios multiplied by $\frac{\pi}{\sqrt{8}}$ (quoted only when the comparison can be made within a single animal). The superscript N refers to normal animals, S to strabismics. Standard deviations are for wavelengths averaged over several different animals, except for those marked *, which give the variability in wavelength within a single animal. The model predicts that when $\frac{\pi}{\sqrt{8}}\frac{\lambda_{od}}{\lambda_{or}}$ is greater than one OD developed first, while if is less than one then OR developed first. The blank columns are included to emphasize where additional data are needed. Note that the study of Murphy *et al* (1998) used the spacing of cytochrome oxidase blobs as an indicator of OD periodicity.

Study	Species	λ_{or}^N	λ_{or}^S	λ_{od}^N	λ_{od}^S	$\frac{\pi}{\sqrt{8}}\frac{\lambda_{od}^N}{\lambda_{or}^N}$	$\frac{\pi}{\sqrt{8}}\frac{\lambda_{od}^S}{\lambda_{or}^S}$
Obermayer and Blasdel (1993)	Macaque	733		903		1.37	
		615		773		1.40	
		588		796		1.50	
		670		829		1.37	
		623		829		1.48	
		611		800		1.45	
Horton and Hocking (1996b)	Macaque			536 ± 82			
Murphy *et al* (1998)	Macaque			590 ± 22	598 ± 11		
Crawford (1998)	Macaque			762 ± 90	770 ± 51		
Löwel *et al* (1987)	Cat	1060 ± 13*					
		1000 ± 14*					
Löwel *et al* (1988)	Cat	915		660		0.80	
		920		725		0.87	
		1100		820		0.83	
		1190		870		0.81	
Löwel (1994)	Cat			900 ± 85	1199 ± 48		
Rao *et al* (1997)	Cat	1070 ± 270*					
	Ferret	720 ± 180*					
Crowley and Katz (1999)	Ferret			612 ± 42			

normal macaques OD develops before OR, whereas in normal cats OR develops before OD. Predictions about how strabismus affects developmental order are difficult due to a lack of data on the wavelength of OR columns in strabismic animals. Löwel *et al* (1998) found no change in pinwheel density in strabismic cats compared with normal ones, possibly implying no change in overall wavelength. Assuming this to be the case, dividing the OD data for strabismics from Löwel (1994) by the OR data for normals from Löwel *et al* (1988) (and multiplying by $\pi/\sqrt{8}$) gives a value of 1.29 ± 0.18, predicting that the order of development in strabismic cats is reversed compared with normal cats. Intuitively one can imagine that the reduced correlation between inputs in the two eyes might drive OD segregation faster than the normally positively correlated inputs. Some further data that support this are those of Löwel (1994), where λ_{od} is both larger *and more sharply defined* (narrower peak in the Fourier power spectrum) than in the strabismic compared with the normal, as in our simulations. However Löwel *et al* (1998) found no change in the DIA for strabismic compared with normal cats, whereas the model predicts a reduced tendency to orthogonality in strabismics compared with normal ones. The prediction of a reversal of order in strabismics could be tested using optical imaging by comparing the time of emergence of OR and OD columns in normal compared with strabismic kittens.

The reason λ_{od} does not also increase in the strabismic macaque is probably because OD segregation in the macaque is already well established at birth (Horton and Hocking 1996a), and overall periodicity is then already too well determined to be affected by subsequent strabismus.

In macaque OD periodicity varies across the global extent of V1, being wider in the foveal representation than in the periphery (Horton and Hocking 1996a). Unless there is an equivalent change in the periodicity of OR columns, which has not been determined experimentally, the model thus predicts that although OR and OD may develop at a similar time in most of V1, OD may develop first in the foveal representation. This would also predict a reduced tendency towards orthogonality in the DIA in this part of V1 compared with elsewhere.

Acknowledgments

This work was supported by NSF grant IBN-9808364 and NIH grant 1R01EY12544. We thank Nick Swindale for many helpful suggestions, and Shau-Ming Wu for coding an earlier version of the simulator.

References

Blasdel G G, Obermayer K and Kiorpes L 1995 Organization of ocular dominance and orientation columns in the striate cortex of neonatal macaque monkeys *Visual Neurosci.* **12** 589–603

Chapman B, Stryker M P and Bonhoeffer T 1996 Development of orientation preference maps in ferret primary visual cortex *J. Neurosci.* **16** 6443–53

Crair M C, Gillespie D C and Stryker M P 1998 The role of visual experience in the development of columns in cat visual cortex *Science* **279** 566–70

Crawford M L 1998 Column spacing in normal and visually deprived monkeys *Exp. Brain Res.* **123** 282–8

Crowley J C and Katz L C 1999 Development of ocular dominance columns in the absence of retinal input *Nature Neurosci.* **2** 1125–30

Das A and Gilbert C D 1997 Distortions of visuotopic map match orientation singularities in primary visual cortex *Nature* **387** 594–8

Dayan P S 1993 Arbitrary elastic topologies and ocular dominance *Neural Comput.* **5** 392–401

Durbin R and Mitchison G 1990 A dimension reduction framework for understanding cortical maps *Nature* **343** 644–7

Durbin R, Szeliski R and Yuille A 1989 An analysis of the elastic net approach to the travelling salesman problem *Neural Comput.* **1** 348–58

Durbin R and Willshaw D J 1987 An analogue approach to the travelling salesman problem using an elastic net method *Nature* **326** 689–91

Erwin E, Obermayer K and Schulten K 1995 Models of orientation and ocular dominance columns in the visual cortex: a critical comparison *Neural Comput.* **7** 425–68

Goodhill G J 1992 Correlations, competition and optimality: modelling the development of topography and ocular dominance *University of Sussex Cognitive Science Research Paper* CSRP 226

——1998 The influence of neural activity and intracortical interactions on the periodicity of ocular dominance stripes *Network: Comput. Neural Syst.* **9** 419–32

Goodhill G J, Bates K R and Montague P R 1997 Influences on the global structure of cortical maps *Proc. R. Soc.* B **264** 649–55

Goodhill G J, Finch S and Sejnowski T J 1996 Optimizing cortical mappings *Advances in Neural Information Processing Systems* vol 8, ed D S Touretzky, M C Mozer and M E Hasselmo (Cambridge, MA: MIT Press) pp 330–6

Goodhill G J and Richards L J 1999 Retinotectal maps: molecules, models, and misplaced data *Trends Neurosci.* **22** 529–34

Goodhill G J and Willshaw D J 1990 Application of the elastic net algorithm to the formation of ocular dominance stripes *Network: Comput. Neural Syst.* **1** 41–59

——1994 Elastic net model of ocular dominance: overall stripe pattern and monocular deprivation *Neural Comput.* **6** 615–21

Hoffsümmer F, Wolf F, Geisel T, Löwel S and Schmidt K 1995 Sequential bifurcation of orientation—and ocular dominance maps *ICANN95: Proc. Int. Conf. On Artificial Neural Networks (Paris, 1995)* vol 1, p 535

——1996 Sequential bifurcation and dynamic rearrangement of columnar patterns during cortical development *Computational Neuroscience: Trends in Research 1995* ed J Bower (New York: Academic) pp 197–202

Horton J C and Hocking D R 1996a An adult-like pattern of ocular dominance columns in striate cortex of newborn monkeys prior to visual experience *J. Neurosci.* **16** 1791–807

——1996b Intrinsic variability of ocular dominance column periodicity in normal macaque monkeys *J. Neurosci.* **16** 7228–39

Horton J C, Hocking D R and Adams D L 1999 Comparison of ocular dominance columns in normal and strabismic squirrel monkeys *Soc. Neurosci. Abstracts* **25** 1808

Hübener M, Shoham D, Grinvald A and Bonhoeffer T 1997 Spatial relationships among three columnar systems in cat area 17 *J. Neurosci.* **17** 9270–84

Kohonen T 1982 Self-organized formation of topologically correct feature maps *Biol. Cybernet.* **43** 59–69

LeVay S, Stryker M P and Shatz C J 1978 Ocular dominance columns and their development in layer IV of the cat's visual cortex: a quantitative study *J. Comput. Neurol.* **179** 223–44

Löwel S 1994 Ocular dominance column development: strabismus changes the spacing of adjacent columns in cat visual cortex *J. Neurosci.* **14** 7451–68

Löwel S, Bischof H J, Leutenecker B and Singer W 1988 Topographic relations between ocular dominance and orientation columns in the cat striate cortex *Exp. Brain Res.* **71** 33–46

Löwel S, Freeman B and Singer W 1987 Topographic organization of the orientation column system in the large flat-mounts of the cat visual cortex: a 2-deoxyglucose study *J. Comput. Neurol.* **255** 401–15

Löwel S, Schmidt K E, Kim D-S, Wolf F, Hoffsümmer F, Singer W and Bonhoeffer T 1998 The layout of orientation and ocular dominance domains in area 17 of strabismic cats *Eur. J. Neurosci.* **10** 2629–43

Mitchison G and Swindale N V 1999 Can Hebbian volume learning explain discontinuities in cortical maps? *Neural Comput.* **11** 1519–26

Murphy K M, Jones D G, Fenstemaker S B, Pegado V D, Kiorpes L and Movshon J A 1998 Spacing of cytochrome oxidase blobs in visual cortex of normal and strabismic monkeys *Cerebral Cortex* **8** 237–44

Obermayer K and Blasdel G G 1993 Geometry of orientation and ocular dominance columns in monkey striate cortex *J. Neurosci.* **13** 4114–29

Obermayer K, Blasdel G G and Schulten K 1992 Statistical–mechanical analysis of self-organization and pattern formation during the development of visual maps *Phys. Rev.* **45** 7568–89

Rao S C, Toth L J and Sur M 1997 Optically imaged maps of orientation preference in primary visual cortex of cats and ferrets *J. Comput. Neurol.* **387** 358–70

Roe A W and Ts'o D Y 1995 Visual topography in primate V2: multiple representations across functional stripes *J. Neurosci.* **15** 3689–715

Rose K 1998 Deterministic annealing for clustering, compression, classification, regression and related optimization problems *Proc. IEEE* **86** 2210–39

Simic P D 1990 Statistical mechanics as the underlying theory of 'elastic' and 'neural' optimizations *Network: Comput. Neural Syst.* **1** 89–103

Simmen M W 1991 Parameter sensitivity of the elastic net approach to the travelling salesman problem *Neural Comput.* **3** 363–74

Swindale N V 1982 A model for the formation of orientation columns *Proc. R. Soc.* **215** 211–30

——1996 The development of topography in the visual cortex: a review of models *Network: Comput. Neural Syst.* **7** 161–24

Tieman S B and Tumosa N 1997 Alternating monocular exposure increases the spacing of ocularity domains in area 17 of cats *Visual Neurosci.* **14** 929–38

Wiesel T N and Hubel D H 1974 Ordered arrangement of orientation columns in monkeys lacking visual experience *J. Comput. Neurol.* **158** 307–18

Wolf F and Geisel T 1998 Spontaneous pinwheel annihilation during visual development *Nature* **395** 73–8

Yuille A L 1990 Generalized deformable models, statistical physics, and matching problems *Neural Comput.* **2** 1–24

Yuille A L, Kolodny J A and Lee C W 1996 Dimension reduction, generalized deformable models and the development of ocularity and orientation *Neural Networks* **9** 309–19