The influence of lateral connections on the structure of cortical maps

Miguel Á. Carreira-Perpiñán^{*} & Geoffrey J. Goodhill Department of Neuroscience Georgetown University Medical Center 3900 Reservoir Road NW, Washington, DC 20007, US miguel@cns.georgetown.edu, geoff@georgetown.edu

June 6, 2004

Abstract

Maps of ocular dominance and orientation in primary visual cortex have a highly characteristic structure. The factors that determine this structure are still largely unknown. In particular, it is unclear how short-range excitatory and inhibitory connections between nearby neurons influence structure both within and between maps. Using a generalized version of a well-known computational model of visual cortical map development, we show that the number of excitatory and inhibitory oscillations in this interaction function critically influences map structure. Specifically, we demonstrate that functions that oscillate more than once do not produce maps closely resembling those seen biologically. This strongly suggests that local lateral connections in visual cortex oscillate only once and have the form of a Mexican hat.

1 Introduction

The primary visual cortex of cats, monkeys and ferrets is characterized by a series of columnar maps, each representing particular features of the visual input. The most noticeable and well-characterized of these are visual field position (VF, topography), ocular dominance (OD) and orientation preference (OR). These maps co-exist on the same neural substrate in highly stereotyped patterns. For instance, OD columns tend to intersect OR columns at steep angles, and OR pinwheels tend to lie at the center of OD columns (Bartfeld and Grinvald, 1992; Obermayer and Blasdel, 1993; Hübener et al., 1997; Crair et al., 1997; Kim et al., 1999; Müller et al., 2000). These maps and their interrelationships have provided a paradigm example for the investigation of how cortical structure arises during development.

In particular, theoretical models of columnar development have been remarkably successful at accounting for many of the observed features of cortical maps (reviewed in Erwin et al., 1995; Swindale, 1996; Carreira-Perpiñán and Goodhill, 2002b). Many of these models can be seen, explicitly or implicitly, as trading off coverage versus continuity (Swindale, 1996; Swindale et al., 2000; Carreira-Perpiñán and Goodhill, 2002a). Coverage refers to a desire to represent all input features uniformly, while continuity refers to the desire to represent all input features smoothly. In the cortex, continuity arises from the pattern of short-range lateral connections. By specifying the degree to which neurons at different distances from a particular cortical neuron are excited or inhibited by that neuron, the continuity term in theoretical models effectively controls the type of surface defined by the sheet of cortical receptive fields through the higher-dimensional space of input features, and thus the relations between maps.

A large number of experimental studies have investigated the form these lateral connections take in V1 (reviewed in Callaway, 1998; Lund et al., 2003). Anatomically, a localized injection of an axonal tracer such as biocytin (e.g. Amir et al., 1993; Bosking et al., 1997; Kisvárday et al., 1997) or GFP adenovirus (Stettler et al., 2002) leads to a small, densely labeled central region of diameter roughly 500 μ m, surrounded by a halo of more patchy and diffuse connectivity. Physiological studies investigating the sign of the connections within the central region have generally found a mixture of excitation and inhibition (e.g. Hata et al., 1991; Roerig and Chen, 2002). It has also been suggested that within this region excitation is more spatially restricted than

^{*}Current address: Dept. of Computer Science, University of Toronto. 6 King's College Road. Toronto, ON M5S 3H5, Canada. Email: miguel@cs.toronto.edu.

inhibition (Hata et al., 1991), though this is not generally confirmed (see also Swindale, 1996). While the shortrange connections are largely non-specific, the longer-range connections tend to link cells with similar orientation preferences (e.g. Bosking et al., 1997; Roerig and Chen, 2002). Although the anatomical data does not directly support a lateral-connection pattern with a Mexican-hat form (central excitation sorrounded by inhibition), other sources of evidence exist. For example, in mouse visual cortex, extracellular recordings and optical imaging (see Schuett et al., 2002 and references inside) have shown that stimulation of a limited region of the visual cortex also elicits surround inhibition, thus suggesting a Mexican-hat pattern. The range of inhibition also seems to extend farther than suggested by the purely anatomical evidence, and this has a role on map development (see Hensch and Stryker, 2004; Fagiolini et al., 2004; and references inside).

Models of cortical pattern formation during development have generally considered only the effect of the dense, short range mixture of excitatory and inhibitory connections, and have ignored the patchy longer-range connections. In particular, most models have assumed lateral interactions take the form of a radially symmetric disk of short-range excitation surrounded by an annulus of longer range inhibition (e.g. von der Malsburg, 1973; Willshaw and von der Malsburg, 1976; Miller et al., 1989; Goodhill, 1993; reviewed in Swindale, 1996). Even models which allow plasticity of lateral interactions have imposed the condition of short-range excitation and longer-range inhibition (Sirosh and Miikkulainen, 1997). Although as discussed above there is little direct biological evidence to support this assumption, such "Mexican hat" functions for short-range connectivity are very popular in models for cortical function as well as cortical development (e.g. Somers et al., 1998; Dragoi and Sur, 2000; Ernst et al., 2001; Kang et al., 2003). One very appealing feature of Mexican hat interactions is their powerful pattern-forming properties (Turing, 1952; Meinhardt, 1982; Ball, 1999).

However, despite this popularity there has been very little investigation of how important the precise form of short-range Mexican hat interactions (i.e., the continuity term) is to theoretical models of cortical development. If it were the case that different types of continuity term led to maps with broadly similar properties, then models would not provide strong constraints on what patterns of lateral connections might exist in the real cortex. Here we show that, on the contrary, the precise form of the continuity term provides very sensitive control over the details of the maps formed. This greatly strengthens the prediction that lateral connections in the real cortex do indeed have the form of short-range excitation and longer-range inhibition. Our results demonstrate the fragility of the interrelations between maps with respect to a class of Mexican-hat functions which could be implemented in biological hardware in terms of localized arbors (of different widths and amplitudes) of excitatory and inhibitory connections.

The particular model we investigate is the elastic net (Durbin and Willshaw, 1987; Durbin et al., 1989; Durbin and Mitchison, 1990; Goodhill and Willshaw, 1990; Dayan, 1993). Together with the Kohonen model, the elastic net exceeds all other models in the closeness with which the maps it produces match the details of real maps (Erwin et al., 1995; Swindale, 1996). Unlike the Kohonen model however, the elastic net has an objective function which implements explicitly the trade-off between coverage and continuity. This property allows the effect of the continuity term to be cleanly separated from that of the coverage term. We first present a novel interpretation of the continuity term in terms of a pattern of lateral interactions in the cortex, and show that the standard form of the continuity term, in effect a first-order derivative, represents short-range excitation and longer-range inhibition. We also generalize the continuity term to higher-order derivatives, and show how these correspond to patterns of lateral interactions which oscillate more than once (i.e., the annulus of inhibition is now surrounded at longer range by an annulus of excitation, and so forth). Simulations of the joint development of OD and OR columns are then compared using these differing assumptions regarding lateral interactions. We find that as derivative order increases normal map structure is gradually disrupted: in particular many OR pinwheels move to OD borders, and the histogram of intersection angles between OR and OD columns becomes flat though the maps themselves remain columnar. These results lead to two testable biological predictions: (1) the sign of short-range lateral interactions changes only once in visual cortices that have been well-characterized so far; (2) if cortices are found where columns do not intersect at right angles and OR pinwheels are not constrained to mostly lie on OD centers, these cortices may contain short-range lateral interactions that change sign more than once.

2 Materials and methods

The elastic net (Durbin and Willshaw, 1987; Durbin and Mitchison, 1990; Goodhill and Willshaw, 1990) produces maps that attempt to minimize a tradeoff between coverage C, the desire to represent all stimulus features uniformly in the cortex (Swindale, 1991), and continuity R, the desire to produce a smooth mapping of stimulus features. This trade-off is expressed in terms of minimizing an "energy function" $E = C + \frac{\beta}{2}R$, where β controls the weight of R with respect to C. Let \mathbf{x}_n represent the stimulus space vector n (combination of preferred position, ocular dominance and orientation), and \mathbf{y}_m represent the center in stimulus space of the receptive field of cortical neuron m. (Note: although it is convenient to refer to m as being a single cortical neuron, it should more properly be thought of as representing a small patch of neurons. We use the term "receptive field" to refer to all features jointly, i.e., visual field, ocular dominance and orientation, and not only to the visual field.) In the elastic net the coverage term is then defined as follows:

$$C(\mathbf{y}_1, \dots, \mathbf{y}_M; K) = -K \sum_{n=1}^N \log \sum_{m=1}^M e^{-\frac{1}{2} \left\| \frac{\mathbf{x}_n - \mathbf{y}_m}{K} \right\|^2}.$$

K is a parameter that gives the receptive field size in feature space ($\|\cdot\|$ represents the Euclidean norm, or length, of a vector), such that large values of K mean cortical neurons are quite unselective for specific features, while small values of K mean they are quite selective. We reduce (anneal) K over the course of the simulation, corresponding to a gradual sharpening of the cortical receptive fields. This implements an efficient optimization procedure known as deterministic annealing (Rose, 1998). This iterative minimization procedure can be interpreted in terms of Hebbian learning (Durbin and Mitchison, 1990), although other interpretations are also possible. The stimulus space is densely populated with points along the dimensions of visual field (VFx, VFy), ocular dominance (OD), and orientation preference and selectivity (OR, ORr). In our case, the stimulus values (training set) are a grid of $N_x \times N_y \times N_{\text{OD}} \times N_{\text{OR}} \times 1$ in the rectangle $[0,1] \times [0,1] \times [-l,l] \times [-\frac{\pi}{2}, \frac{\pi}{2}] \times [0,r]$, where OR is conventionally represented by two variables in polar coordinates, with N_{OR} values uniformly arranged on a ring of radius r (see e.g. Swindale, 1996). There are $N = N_x N_y N_{\text{OD}} N_{\text{OR}}$ such 5-dimensional points \mathbf{x}_n . The net consists of a square lattice with M centroids, representing a square array of cortical neurons (the depth dimension of the cortex is not modeled).

The continuity term in the elastic net is defined as

$$R(\mathbf{y}_1,\ldots,\mathbf{y}_M) = \sum_{m=1}^M \|\mathbf{f}(\mathbf{y}_m)\|^2$$

where for each neuron m the term \mathbf{f} is a fixed linear combination of cortical neurons in a neighborhood of m. That is, \mathbf{f} specifies the strength of the interaction between a cortical neuron and its neighbors. We call the coefficients of the linear combination the *stencil*. In the original formulation of the elastic net (Durbin and Willshaw, 1987), R was defined as $\sum_{m=1}^{M} \|\mathbf{y}_{m+1} - \mathbf{y}_m\|^2$, i.e., the sum of squared distances in feature space between adjacent receptive field centers of neighboring cortical neurons. In the generalized elastic net (Carreira-Perpiñán and Goodhill, 2004), each term $\mathbf{f}(\mathbf{y}_m) = (\mathbf{y}_{m+1} - \mathbf{y}_m)$ is viewed as a discretized derivative operator of order p = 1, i.e., an approximate gradient. We can prove (Carreira-Perpiñán and Goodhill, 2004) that the continuity term can be rewritten as a Mexican hat function as follows (see p = 1 in table 1). Define the *cortical interaction function* e_k and the local linear combination \mathbf{g} with coefficients e_k as:

$$e_k = \frac{-4}{\pi(4k^2 - 1)}, \ k = 0, \pm 1, \pm 2, \dots$$
 $\mathbf{g}(\mathbf{y}_m) = \sum_{k = -\infty}^{\infty} e_k \mathbf{y}_{m+k}.$ (1)

Then, for a large number M of net centroids, it holds that

$$R(\mathbf{y}_{1},...,\mathbf{y}_{M}) = \sum_{m=1}^{M} \|\mathbf{y}_{m+1} - \mathbf{y}_{m}\|^{2} \approx \sum_{m=1}^{M} \|\mathbf{g}(\mathbf{y}_{m})\|^{2}$$

(exact expressions can be derived for any finite M, but are complicated and nearly coincide numerically with the limit approximation). In other words, we can represent the continuity term $R(\mathbf{y}_1, \ldots, \mathbf{y}_M)$ either in terms of the original linear combinations $\mathbf{f}(\mathbf{y}_m) = \mathbf{y}_{m+1} - \mathbf{y}_m$ or in terms of the new linear combinations $\mathbf{g}(\mathbf{y}_m)$ given in eq. (1). That is, the continuity term in the original elastic net is mathematically equivalent to a Mexican hat function consisting of short-range excitation and long-range inhibition. This demonstrates the equivalence in the elastic net model of the Mexican-hat cortical interaction and the interpretation of Durbin and Mitchison (1990) in terms of wirelength minimization (in the stimulus space). It also shows that the sum-of-squared-distances equation, which involves only nearest-neighbor connections, can be implemented with a denser net of connections (whose strength decays rapidly with distance).

Instead of choosing, for the function $\mathbf{f}(\mathbf{y}_m)$, the first order derivative of sensory space over cortex, we could explore the effects of choosing higher order derivatives. For this paper we chose derivatives of order 2, 3 and 4. The discrete forms of these higher-order derivatives are shown in table 1; they are obtained by repeatedly



Table 1: Cortical interaction functions for continuity terms of differential order p = 1 to 4 in the generalized elastic net in 1D. The equations for R are the discretized difference formulation (using the **f** linear combinations), while the graphs show the equivalent Mexican-hat representation and corresponding cortical interaction function (using the **g** linear combinations). The vertical axes have been rescaled to aid the visualization. All functions are constrained to integrate to 0 (note the sum of the coefficients of the discretized difference is 0 in each row). For p = 1, the cortical interaction function has only one change of sign (considering only one half of the horizontal axis, by symmetry); as p increases, the function develops progressively more changes of sign (oscillations). For a 2D net, the resulting lateral connections consist of several concentric regions of alternating excitation and inhibition.

applying the first-order derivative $(\mathbf{y}_{m+1} - \mathbf{y}_m)$. In each case, we can show that the sum $\sum_{m=1}^{M} \|\mathbf{f}(\mathbf{y}_m)\|^2$ is equivalent, in the limit of large M, to $\sum_{m=1}^{M} \|\mathbf{g}(\mathbf{y}_m)\|^2$ with an appropriate \mathbf{g} interaction function; and that this \mathbf{g} interaction function has a progressively larger number of distal oscillations, or alternating rings of excitatory and inhibitory connections, with increasing order of the derivative. Since the continuity term R should not depend on the location of the net, but only the relative arrangement of its centroids, the coefficients in the stencil should add to zero, which is the same as requiring that applying the stencil to a constant function give zero (by definition of differential stencil). 2D stencils can be readily obtained by passing a 1D stencil horizontally and vertically, resulting in two terms $R_{\rm h}$ and $R_{\rm v}$ that are then added. This is the natural extension of the first-order 2D stencil used in the original elastic net, though it leads to a Fourier spectrum that is increasingly anisotropic as p increases. There are other ways to construct 2D stencils, e.g. the Laplacian stencil $\nabla^2_+ = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \end{pmatrix}$, of order p = 2, which has a more isotropic spectrum. Note however that in a discrete grid there is not a unique definition of the concept of isotropy. Although in reality it is likely that the lateral connections are maturing as the maps develop (Sur and Leamey, 2001), the elastic net model (in common with most other models) simplifies this situation by assuming a fixed pattern of lateral connections.

Simulation setup We trained nets for derivative orders p = 1, 2, 3 and 4 with the following configuration: training set: $N_x = N_y = 20$, $N_{OD} = 2$, $N_{OR} = 12$, l = 0.09 and r = 0.16 (totalling N = 9600 points in 5 dimensions); net: 128×128 centroids; $\beta = 10$; nonperiodic boundary conditions. We reduced K with an annealing rate of 0.992 from 0.1 to the point at which the OD and OR maps have just arisen ($K \approx 0.05$ for all values of p). We achieved the minimization of E at each value of K by Cholesky factorization (Carreira-Perpiñán and Goodhill, 2004). This method is both more efficient and more robust than the numerical procedures used previously (e.g. Durbin and Willshaw, 1987; Durbin and Mitchison, 1990; Goodhill and Willshaw, 1990), which were limited to very small β values. We ran 5 different simulations for each value of p in order to obtain error bars in figs. 3–4, adding a small amount of random noise to the initial net and the training set in each case. The simulations were performed using custom software written in Matlab.

Several of the statistics reported are based on histograms; in all cases we used n = 10 bins and took each bin's center as class representative. Using a different number of bins (100 instead of 10) did not significantly alter any of the results. In order to quantitatively characterize the histograms (average and variability) as a function of p in a compact way in fig. 3, we summarized each histogram by 3 statistics. (1) Its Kullback-Leibler divergence (Cover and Thomas, 1991) to the uniform histogram in the same domain, defined as $KL(u) = \sum_{i=1}^{n} p_i \log \frac{p_i}{u_i}$, where p_i and u_i are the values (normalised to sum 1) at bin i of the histogram and the uniform histogram to $\log n \approx 2.303$ when all p_i are zero except at a single bin. (2) Its Fisher skewness $\gamma_1 = \mu_3 \mu_2^{-3/2}$ (where μ_i is the *i*th central moment), which is 0 for a histogram symmetric around its mean and negative (positive) for tails on the left (right) of the mean. (3) Its mean μ . Both γ_1 and μ are independent of the number of bins n; all three are independent of translations or scalings of the data. Other summary statistics are possible, but these three sufficed to characterize the histograms we obtained.

Stripe width The power of the discrete Fourier transform of OD and OR angle maps was roughly isotropic and concentrated around a ring. We summarized it by the mean wavelength (mean stripe width), averaged over all directions.

Crossing angles We disregarded the points lying 5 pixels or less from the boundaries of the net (a thin inner stripe framing the maps in fig. 1) to eliminate border effects. At each pixel, we computed the angle between the gradient vectors of OD and OR (see below) and mapped it to $[0^{\circ}, 90^{\circ}]$, thus obtaining the angle between contours of OD and OR. Each such angle counted with a weight proportional to the product of the gradient moduli of OD and OR, $\|\nabla OD\| \times \|\nabla OR\|$, so that pixels lying in areas of either constant OD or constant OR (i.e., away from borders of OD or contours of OR) were effectively removed from the histogram. This is because, in an area where either map is nearly constant, the gradient vector is negligible in modulus and its direction basically arbitrary; unlike along borders, where the gradient vector is large and its direction well defined.

Pinwheel distributions An OR singularity, or pinwheel, is defined positive if the orientation angle increases in a clockwise direction around the pinwheel and negative if anticlockwise. Pinwheels were automatically located in the OR maps as follows. First, the winding number of each pixel in the OR map was estimated by summing the increments of OR angle (in $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$) along a closed path (a square of radius 1 pixel centered in the pixel considered) in a clockwise direction and dividing the result by 2π ; this results in 0 for non-pinwheel points and $+\frac{1}{2}$ $\left(-\frac{1}{2}\right)$ for pixels at or adjacent to a positive (negative) pinwheel, respectively. The exact pinwheel location was obtained by grouping clusters of nonzero winding number and computing their centers of mass. To quantify the layout of pinwheels with respect to OD stripes, we computed, for every pinwheel, the distribution of distances of a given pinwheel to its closest OD border. We prefer this to calculating distances to OD centers since the "center of an OD stripe" is generally not well defined, except for stripes which are translations of each other, and becomes difficult to use with stripes with irregular boundaries, forks, islands or other complex shapes. In contrast, the OD borders are well defined in all these cases.

To compute the distances between a pinwheel and its closest OD border, we represented the OD borders by a finely spaced collection of points (the contours generated by Matlab for the OD map); the said distance was then given by the point in this collection closest to the given pinwheel. We considered that a pinwheel lay on an OD border if the distance between the two was 1 pixel or less. The distributions of random pinwheels were built by generating, uniformly over the cortex, the same number of pinwheels as in the actual map produced by the simulation, and then computing the distance histogram; this process was repeated 50 times to obtain an average histogram with error bars that is the one plotted along with the actual ones for comparison. Note that theoretical distributions (such as the Rayleigh distribution for nearest-neighboring pinwheels; Müller et al., 2000) differ from the true one because they do not take into account the border effects of a finite cortex.

Gradient correlations The gradient of a variable f (one of VFx, VFy, OD, OR) with respect to the cortical location (x, y), $\nabla f(x, y) = \left(\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right)$, was approximated at each pixel (i, j) by a central finite difference: $\nabla f_{ij} = \left(\frac{f_{i+1,j}-f_{i-1,j}}{2h}, \frac{f_{i,j+1}-f_{i,j-1}}{2h}\right)$, where h = 1 pixel is the grid size. For points along the cortex boundaries, a forward difference $\frac{f_{i+1,j}-f_{i,j}}{h}$ was used at a small accuracy loss. In both cases, for OR the gradient values were corrected

to have a maximal absolute value of $\frac{\pi}{2}$. The 2D gradient vector ∇f was converted to polar coordinates (modulus, angle) for each variable separately, with the angle being used in the histogram of crossing angles. The moduli were used for the correlation plots (the modulus of VF gradient was defined operationally as the sum of the gradient moduli of VFx and VFy). The units of the gradient modulus of variable f are those of variable f per pixel. The linear correlation was quantified by Pearson's correlation coefficient r. The customary regression lines $X \to Y$ and $Y \to X$ minimize the vertical and horizontal errors, while the line X, Y minimizes the orthogonal distances to the line and is given by the principal component of the data covariance matrix. All 3 lines pass through the data mean but with different slopes in general. For highly correlated data ($r^2 \approx 1$), all 3 lines are very close, while for highly uncorrelated data ($r^2 \approx 0$) $X \to Y$ and $Y \to X$ become horizontal and vertical, respectively.

Gradient correlations have also been reported in the experimental literature (e.g. Das and Gilbert, 1997; Hetherington and Swindale, 1999), where the preferences of VF and OR were determined electrophysiologically at several cortical sites and then the difference in preference between neighboring sites divided by their cortical distance. However, for practical reasons "neighboring" sites are much further apart over the cortical surface in these studies than neighboring points in our simulated cortex.

3 Results

We present the results obtained with a generalized elastic net model of the maps of visual field position (VF), ocular dominance (OD) and orientation preference (OR). The model attempts to represent the 5 dimensions of VF (x and y directions), OD and OR (preferred angle and selectivity) using a 2D array of cortical neurons (the elastic net), thus achieving a reduction of dimension. The maps are the result of maximizing the coverage of these stimuli while minimizing a continuity constraint on the shape of net (the geometry of the maps). The continuity constraint penalizes the squared magnitude of the derivative of order p integrated over the whole net, thus forcing a representation of the maps that is as smooth as possible. Biologically, the continuity constraint represents a cortical interaction function with the shape of a Mexican hat, in which the number of excitatory-inhibitory regions increases with the derivative order p (table 1). We considered p = 1 to 4. Table 1 shows, for each of p = 1 to 4, the corresponding form of the continuity term in the elastic net, and the equivalent (short-range) cortical interaction function (for details see Methods). For p = 1 (the standard elastic net) the cortical interaction function has the form of short-range excitation surrounded by an annulus of longer-range inhibition which decreases smoothly to zero. For p = 2, the magnitude of the longer-range inhibition decreases quickly to zero. For p = 3, the magnitude of the longer-range inhibition decreases quickly to zero, and then there is a small annulus of excitation surrounding this which decreases smoothly to zero. For p = 4, the annulus of excitation is stronger but falls rapidly to zero. It should be emphasized that in each case these functions represent only radially-symmetric short-range cortical interactions, i.e., within a radius of about 500 microns, and longer-range patchy connections are not considered.

By simulating the development of VF, OD and OR maps using the elastic net with each of these continuity terms, we demonstrate that these very subtle differences in the form of the Mexican hat cortical interactions have profound effects on the detailed structure of the maps obtained. Fig. 1 shows the maps obtained from one simulation for derivative orders p = 1 to 4. The maps obtained in simulations using different random numbers differed in the detailed layout but were qualitatively similar. Figs. 2 and 5 quantitatively assess the specific maps of fig. 1, while figs. 3–4 summarize the various statistics as a function of the derivative order; the average value and error bars (standard deviations) result from 5 simulations per order.

Figure 1 (following page): Simulated maps of ocular dominance and orientation for a generalized elastic net of derivative orders p = 1 to 4. Left column: ocular dominance map. Center column: orientation polar map. Right column: contours of ocular dominance and orientation. For each order p, the continuity terms and cortical interaction functions are given in table 1. The grids contain 128×128 cortical points. This is one of 5 simulations run with different noise in the initial net and training set. The training set consists of a uniform, dense grid in the 5-dimensional stimulus space of visual field (VFx, VFy), ocular dominance (OD) and orientation (selectivity ORr and preferred angle ORt). The maps are the result of minimizing the generalized elastic net energy $E = C + \frac{\beta}{2}R$ with $\beta = 10$. Note how as p increases the pinwheels increasingly co-locate with OD borders, and the contours increasingly align along the diagonals. For p = 1, the pinwheels are preferentially away from OD borders, near OD centers, and the contours of OD and OR intersect preferentially at right angles. As p increases, more and more pinwheels lie right on OD borders, like beads on a string, and the contours tend to intersect at any angle.



p = 1

p = 2

p = 3

p = 4

Crossing angles The distribution of crossing angles of the contours of OD and OR for the maps of fig. 1 are shown in the left column of fig. 2. For p = 1, this distribution is strongly skewed towards orthogonality, in agreement with experimental data (macaque: Bartfeld and Grinvald, 1992; Obermayer and Blasdel, 1993; cat: Hübener et al., 1997; Crair et al., 1997; Kim et al., 1999; Müller et al., 2000). However, as p increases, the distribution quickly loses its bias, and becomes practically uniform for $p \geq 3$. The loss of orthogonality is apparent by close inspection of the contours in fig. 1 (right column). The transformation of the distribution from being sharply skewed towards orthogonality (mean 65°) to being practically uniform (mean $\approx 45^{\circ}$) is clearly shown in fig. 3 (first row).

Stripe widths Fig. 4 (left) plots the mean wavelength of OD and OR in pixels, showing that the stripes become narrower as the derivative order increases. The stripe width depends for any derivative order on the variance along the OD and OR dimensions in the training set, and is related to the order in which both maps arise (Goodhill and Cimponeriu, 2000). For the results shown in this paper we chose these variance parameters such that OD columns are slightly wider than OR columns, as is the case in macaques (e.g. Obermayer and Blasdel, 1993). However, we also investigated the case where OR columns are wider than OD columns, corresponding to the case in cats (e.g. Löwel et al., 1988). No important differences were found of increasing derivative order in the effects we have quantified here (data not shown). Specifically, all histograms and curves in figures 2–5 remained basically the same (same increasing or decreasing pattern, similar location of peaks, etc.), and in figure 4 (left) the two curves cross over between p = 1 and p = 2, since the OR wavelength decreases faster than the OD one.

It is noticeable that the stripes become increasingly more aligned with the diagonals of the cortical region considered as p increases, as seen from fig. 1. This is due to the fact that the 2D stencils (pattern of lateral interactions) we use are increasingly anisotropic in Fourier space as p increases (Carreira-Perpiñán and Goodhill, 2004). In order to eliminate the possibility that some of the other effects on map structure we report might not appear using more isotropic 2D stencils representing equal derivative order, we also investigated a 2D Laplacian stencil (order p = 2), which has a more isotropic spectrum. The matching of pinwheels and OD borders and the uniformity of crossing angles still appeared with this stencil (results not shown), so we conclude that the anisotropy is not responsible for the effects we report.

Orientation pinwheels Fig. 2 (right column) shows the distribution of *distances between nearest-neighboring pinwheels*, normalized by the wavelength of OR. The histograms are very similar for all orders *p*. When pinwheels are considered independently of their sign, the histogram peaks quite sharply at half the OR wavelength and has a slight tail on the right towards higher distances, while quickly dropping to 0 for small distances. If the same number of pinwheels are uniformly distributed at random, then the histogram peaks at a smaller distance and has a much longer right tail. When only pinwheels of the same sign are considered, the histogram contains half

Figure 2 (following page): Distribution of crossing angles of OD and OR, distance pinwheel-to-closest-OD-border and distance pinwheel-to-closest-pinwheel for derivative orders p = 1 to 4, for the maps of fig. 1. The number of bins is 10 for all histograms; the value at each bin is plotted against the center of the bin and the resulting points are connected by lines. Left column: histogram of crossing angles in $[0^{\circ}, 90^{\circ}]$ of the contours of OD and OR. The horizontal dashed line represents the uniform histogram. Note the bias for orthogonality for p = 1 and the progressive flattening as p increases. Center column: histogram of the distance from a pinwheel to its closest OD border. The axes represent the number of pinwheels as a function of the distance in fractions of the OD wavelength. The two vertical dashed lines mark $\frac{1}{4}$ of the wavelengths of OR and OD. "Random (all)" means pinwheels distributed uniformly at random in the map (the standard deviation bars resulting from 50 such random sets); "actual (all)" means the actual pinwheels shown in fig. 1 irrespectively of their sign, with "actual (+)" and "actual (-)" meaning the histogram considering only positive or negative pinwheels, respectively. Note that for p = 1 there is a single peak (at around $\lambda(OD)/4$, the center of an OD stripe); but for p > 1 there is also a second peak (at 0, the OD border) whose height progressively increases with p at the expense of the OD-center peak. In all cases the histogram is very different from that resulting from random pinwheels. *Right column*: histogram of the distance between nearest-neighboring pinwheels. The axes represent the number of pinwheels as a function of the distance in fractions of the OR wavelength. The two vertical dashed lines mark $\frac{1}{2}$ of the wavelengths of OR and OD. "Random (all/all)" means pinwheels distributed uniformly at random in the map (the standard deviation bars resulting from 50 such random sets); "actual (all/all)" means the actual pinwheels shown in fig. 1 irrespectively of their sign, with "actual (+/+)" and "actual (-/-)" meaning the histogram considering only positive or negative pinwheels, respectively. This histogram remains essentially the same for all values of p.





Figure 3: Summary statistics for the histograms of fig. 2 as a function of the derivative order p = 1 to 4 (explanations of what these statistics mean are in the Methods). The graphs show the average over 5 different runs of a histogram statistic and the error bars represent standard deviations. The statistics shown are the Kullback-Leibler divergence to the uniform histogram KL(u) (left column), the Fisher skewness γ_1 (center column) and the mean μ (right column). The vertical axis for the KL(u) statistic is the same for all rows, as is for the γ_1 . Row 1: statistics for the histogram of crossing angles. All three statistics decrease with increasing p, indicating a left-skewed histogram for p = 1 that quickly flattens for p > 1. Row 2: statistics for the histogram of distance pinwheel-to-closest-OD-border. This histogram remains nonuniform for all p (high KL(u)) but changes the skewness sign and decreases the mean for p > 1, indicating the growth of the peak at distance 0 (OD border). Row 3: statistics for the histogram of distance pinwheel-to-closest-ode pinwheel-to-closest-ode pinwheel-to-closest-pinwheel. This histogram remains essentially unchanged, with a tiny decrease of its mean for increasing p due to the increase of the number of pinwheels per OR module (fig. 4, right column).



Figure 4: Summary statistics for map wavelengths and pinwheels on OD borders as a function of the derivative order p = 1 to 4. The graphs show the average over 5 different runs of a statistic and the error bars represent standard deviations. *Left graph*: mean wavelengths of OD and OR in pixels (note the cortical grid is 128×128 pixels). *Right graph*: percentage of pinwheels per module of OR lying right on OD borders, where a module of OR is defined as a square of side equal to the OR wavelength (the plot using the module of OD looked exactly the same). Note how both the OR and OD wavelengths decrease with p and the percentage of pinwheels on OD borders per module of OR increases dramatically from a minority (10%) for p = 1 to half for p = 4.

the number of pinwheels and has the same shape, but peaks at about $\frac{3}{4}$ the OR wavelength. This is confirmed by the summary statistics of fig. 3 (third row); the large error bars for p = 1 are due to the very small number of pinwheels present in the map (around 50), which results in a more variable histogram. A slight tendency for pinwheels to be more closely packed as p increases is visible in fig. 3 (μ graph), consistent with the slight increase of pinwheels per OR module mentioned above.

This reflects the fact that pinwheels are laid out in a noisy grid-like arrangement (of grid length half the OR wavelength), with pinwheels repelling each other, so that same-sign pinwheels alternate in the nodes of the grid. Thus, while the numbers of pinwheels and stripe widths vary with p, the arrangement of pinwheels relative to each other remains the same for all p, in agreement with optical imaging data for squirrel and macaque monkeys (Obermayer and Blasdel, 1997).

The effect on the distribution of the distance from a pinwheel to its closest OD border is more striking. In fig. 2 (center column), for p = 1 the histogram shows a quite sharp peak around $\frac{1}{4}$ of the OD wavelength, with a long tail on the left. Thus, pinwheels tend to locate at the centers of OD columns, almost always avoiding OD borders, as apparent in fig. 1 and in agreement with optical imaging data (macaque: Bartfeld and Grinvald, 1992; Obermayer and Blasdel, 1993; cat: Hübener et al., 1997). However, for p = 2 the histogram becomes bimodal, preserving a now less sharp peak at $\frac{1}{4}$ of the OD wavelength and developing a very sharp peak at distance 0, indicating that pinwheels generally lie either away from the OD borders (at a distance around 0.17 of the OD wavelength) or right on top of an OD border. As p increases, the peak at distance 0, which remains

Figure 5 (following page): Pairwise scatterplots of the gradient moduli of visual field, ocular dominance and orientation angle for derivative orders p = 1 to 4, for the maps of fig. 1. The total number of points in each cloud is $128 \times 128 = 16\,384$, equal to the number M of points in the net. The pinwheels are marked by red (positive) and blue (negative) circles. The units of the axes are # per pixel, where # is degrees for the orientation angle and arbitrary for the visual field and ocular dominance (i.e., as given by the training set). Note the different scales in the axes. The lines shown are the regression of Y on X ($X \to Y$), of X on Y ($Y \to X$) and the principal component (X, Y). The clouds of points are diffuse, indicating that all combinations of gradients are abundantly represented, also evidenced by the low correlation coefficients. In general, the clouds appear compressed because of a rescaling of the axes needed to accommodate a few points of very high gradient (such as pinwheels). The apparent trend in the slope of the regression lines as p increases is caused by the increase in the number of pinwheels and their separation into two groups: right at OD borders (high OD gradient) and away from them (near-zero OD gradient). These two groups correspond to the two peaks at 0 and $\approx \lambda(OD)/4$, respectively, in the histogram of distance pinwheel-to-closest-OD-border of fig. 2 (center column).



Gradient modulus of OD vs

sharp, increases at the expense of the peak at distance 0.17 of the OD wavelength, which becomes broader. This remarkable effect is apparent in fig. 1: note for $p \ge 2$ how the OD borders meet runs of pinwheels like beads on a string. As fig. 2 shows, this distribution is very different from that resulting from the same number of pinwheels but uniformly distributed at random, which would have a roughly flat value for distances up to 0.15 of the OD wavelength to then taper off at larger distances. The histograms also show that, for all orders, the distance to OD borders is equally distributed with respect to the sign of the pinwheels (naturally, the histograms of positive and negative pinwheels are half the height of the histogram for all signs). The summary statistics in fig. 3 (second row) confirm that the histogram remains nonuniform for all p, its mean μ (normalized by the OD wavelength) monotonically decreases and its skewness changes from negative for p = 1 (tail on the left) to positive for high p(tail on the right). Fig. 4 (right) shows how the proportion of pinwheels that lie right on OD borders grows from barely 10% for p = 1 to almost 50% for p = 4. Note that the flattening of the angle histogram is consistent with the increasing proportion of pinwheels on OD borders, since an OD border crossing a pinwheel must be parallel to some OR contours. In other words, an orthogonality bias requires that pinwheels avoid OD borders.

Gradient correlations Fig. 5 shows the pairwise correlation plots between the gradient moduli of VF, OD and OR (cf. Das and Gilbert, 1997; Hetherington and Swindale, 1999; White et al., 2001; Bosking et al., 2002; Buzás et al., 2003) for orders p = 1 to 4. In general, although there is an absence of pixels where the gradient of two or more variables are simultaneously high, there is practically no correlation between gradients, as shown by the low r value and the horizontal and vertical regression lines $X \to Y$ and $Y \to X$, respectively. As p increases, the maximal gradient values increase (note the increasing ranges of the axes) due to a small proportion of high-gradient points, but by and large the bulk of the cloud changes little. In particular, there is neither a positive correlation between gradients nor a negative one, or very small if at all (as in Hetherington and Swindale, 1999; White et al., 2001; Bosking et al., 2002; Buzás et al., 2003; and in contrast with Das and Gilbert, 1997). What is remarkable, and confirms what we said above, is the correlation between the gradients of OD and OR at points at or near pinwheels (which is responsible for the apparent trend in the slope of the regression lines). For p = 1, pinwheels (by definition characterized by a high OR gradient) lie almost exclusively in areas of zero OD gradient, i.e., away from borders. As p increases, the pinwheel population separates in two groups, one at near-zero OD gradient and another at high OD gradient, i.e., the OD borders, with very few pinwheels lying in between.

4 Discussion

Although a dense network of short-range cortical connections is known to exist in visual cortex (e.g. Amir et al., 1993; Bosking et al., 1997; Kisvárday et al., 1997; Callaway, 1998; Stettler et al., 2002; Lund et al., 2003), the precise form of the resulting intracortical interaction function is unclear. We have examined the effect of varying the form of the continuity term representing these short-range interactions in the elastic net model of visual cortical map development. In particular, we have investigated continuity terms corresponding to 1st, 2nd, 3rd and 4th order derivative constraints on the rate of change of cortical receptive field location in a feature space consisting of visual field position, ocular dominance and orientation preference. We have shown that this progression of derivatives corresponds to a progression of patterns of short-range intracortical interactions with increasing numbers of oscillations in the sign of the interactions. Our simulations lead us to hypothesize that the greater the number of oscillations in the continuity term the stronger the co-location between OR pinwheels and OD borders. In particular, we predict that lateral connections oscillate only once in the regions of cat and monkey cortex that have been well-characterized by optical imaging so far. These results provide a way to infer the form of the lateral connections from the observed map structure. Thus, if in other species, visual cortical areas, or at other developmental stages a strong co-location between OR pinwheels and OD borders is observed, this can be interpreted in our model in terms of lateral interactions that oscillate more than once.

In fact, a recent experimental report might be consistent with this. Kisvárday et al. (2001) optically imaged the maps of OD and OR in adult cats (> 4 months) and kittens (< 4 months). They found that the tendency of pinwheels to lie near OD centers and of OD and OR maps to intersect at steep angles was hardly present in the region between areas 17 and 18, and very weak in area 18. Further, the tendency was significantly weaker for adult cats than for kittens in all three regions. These results suggest that the arrangement of pinwheels with respect to OD domains may change with age and cortical region. It is also worth noting that published experimental data do show a proportion, though small, of pinwheels right on OD borders (e.g. fig. 4A of Hübener et al., 1997).

The growing co-location of discontinuities of OR (pinwheels) and OD (borders) and the flattening of the angle histogram with increasing p seems to be an essential feature of the elastic net. However, although the dependence of these statistics on p is generally gradual and monotonic, there seem to be two qualitatively different classes

of lateral interactions: one for p = 1 and another for $p \ge 2$, with p = 2 often being a transition case, as evidenced by the sharp migration of pinwheels to OD borders for p > 1 (figs. 1–2). Another characteristic feature of the generalized elastic net is the narrowing of the stripes and the fact that the ratio of OD to OR wavelength is smaller for the first-order derivative than for the higher-order ones, given by fig. 4 (left). In principle, this might suggest cats have cortical interaction functions of higher order than monkeys (consistent with there being less tendency towards pinwheels lying in OD centers in cats compared to monkeys). Also, this could suggest that any changes in the ratio of wavelengths in an individual between the fovea and periphery might be indicative of a change in lateral interactions over the cortex. However, the stripe widths are also strongly determined by the correlations in the input space (Goodhill and Cimponeriu, 2000) and the strength of the continuity term β , which makes it difficult to separate the individual contributions of each of these factors.

Our results show that interrelations between OD and OR maps are caused by the specific form of lateral interactions used by the elastic net. From a structural point of view, what is special about the first-order derivative, i.e., lateral interactions that oscillate only once? This form was originally used in the elastic net since it is relevant to solving the Traveling Salesman Problem, i.e., minimizing (squared) distances between the receptive fields in feature space of neurons that are neighboring in the cortex (Durbin and Willshaw, 1987; Durbin and Mitchison, 1990). This also had a flavor of wirelength minimization though in the input feature space (Durbin and Mitchison, 1990) rather than the cortical space (Koulakov and Chklovskii, 2001). Durbin and Mitchison (1990) explored the effect of varying the exponent in the sum-of-distances term, but did not explore higher-order derivatives as we have done. The motivation for exploring higher-order derivatives comes also from regularization theory (Tikhonov and Arsenin, 1977). Derivative-based penalties can be used to represent biological or physical constraints and so obtain good solutions to inverse problems in areas as diverse as vision (Poggio et al., 1985), geophysics (Tarantola, 1987) or statistical smoothing and spline approximation (Wahba, 1990; Green and Silverman, 1994). The effect of the derivative order on the stripe width is theoretically understood in some of these problems (Wahba, 1990; Carreira-Perpiñán and Goodhill, 2004). However, at present such a theoretical understanding is lacking for the other aspects of the 2D structure of the maps that are more subtle but occur reliably in our simulations, such as crossing angles and pinwheel locations. Note also that the different maps are independent in the continuity term, so that these effects are due to the combined action with the coverage term. However, the uncertainty about the definition of coverage is smaller than that of continuity (Swindale et al., 2000; Carreira-Perpiñán and Goodhill, 2002a) and we may expect similar effects with slightly different definitions of coverage. From a functional point of view, the consequences of any of these different forms for the lateral interactions remain unclear.

It should be noted that different models have somewhat different ways of mapping between the form of cortical continuity they implicitly or explicitly employ and a pattern of lateral interactions. In, for instance, correlational models such as that of Miller et al. (1989), the effective cortical interaction **I** maps to a pattern of lateral interactions **B** by the relation $\mathbf{I} = (\mathbf{1} - \mathbf{B})^{-1}$ (where **I** and **B** are matrices and **1** represents the identity matrix; Miller et al., 1989; Miller and Stryker, 1990). In this model, **I**'s of Mexican hat form generally thus correspond to patterns of lateral interactions with many oscillations, and when the joint development of OR and OD is simulated with this model, neither the orthogonality between the two columnar systems nor the positioning of OR pinwheels at the center of OD columns are robustly reproduced (Erwin and Miller, 1998; see also Piepenbrock et al., 1997). Mathematical relationships between correlational and elastic-net-type models are discussed by Dayan (1993) and Yuille et al. (1996).

In general, we find a substantially weaker relationship in our simulations between the rate of change of different input features with cortical position than has generally been thought to occur in the elastic net (Durbin and Mitchison, 1990). Also, by simulating tangential penetrations across pinwheels (data not shown) we do not generally find a jump in the VF position (as confirmed by the fact that pinwheels lie in areas of low to middle VF gradient in fig. 5). The diffuse clouds of points in fig. 5 demonstrate that many types of relationships between gradients of variables coexist in the simulated cortex. The elastic net (or dimension reduction models in general) is not best characterized as producing purely negatively correlated representations where, in general, when one feature varies much, the remaining ones vary little (Durbin and Mitchison, 1990). Das and Gilbert (1997) have claimed that the rate of change of orientation preference is positively correlated with that of visual field position in the cat, though this has not been reproduced by other investigators (cat: Hetherington and Swindale, 1999; Buzás et al., 2003; ferret: White et al., 2001; tree shrew: Bosking et al., 2002). It is interesting to note that a model specifically formulated to obtain this result (Mitchison and Swindale, 1999) yielded as a by-product pinwheels on OD borders (corresponding to a positive correlation between the rate of change of OR and OD), which is contrary to what is observed in optical imaging images. Our results show that, as the derivative order increases, the discontinuities of OD and OR increasingly tend to co-locate, but the gradients of OR and VF do not correlate. However, it is important to note that the methodology of experimental studies is practically constrained to be quite different from the complete analysis of rates of changes between feature variables that is possible in theoretical models (fig. 5). Our gradients are, up to the small discretization error induced by the grid, "true" gradients. Gradients computed from a set of electrode penetrations cannot be considered true gradients for three reasons: first, they are really directional derivatives along the line joining two sites (the gradient equation takes into account all directions); second, the separation between sites is usually so large (over 100 μ m, or about 10% of the OR wavelength) that the discretization error becomes very large and can corrupt any local correlation or lack of it; third, the number of sites is very small (a few dozens compared to M = 16384 in our simulations). The significance of the location of pinwheels in areas of low to medium VF gradients with respect to coverage uniformity (Das and Gilbert, 1997; Bosking et al., 2002) remains unclear.

5 Acknowledgements

We are very grateful to Peter Dayan for helpful discussions.

References

- Y. Amir, M. Harel, and R. Malach. Cortical hierarchy reflected in the organization of intrinsic connections in macaque monkey visual cortex. J. Comp. Neurol., 334(1):19–46, Aug. 1 1993.
- P. Ball. The Self-Made Tapestry. Pattern Formation in Nature. Oxford University Press, 1999.
- E. Bartfeld and A. Grinvald. Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular-dominance columns in primate striate cortex. *Proc. Natl. Acad. Sci. USA*, 89(24):11905–11909, Dec. 15 1992.
- W. H. Bosking, J. C. Crowley, and D. Fitzpatrick. Spatial coding of position and orientation in primary visual cortex. Nat. Neurosci., 5(9):874–882, Sept. 2002.
- W. H. Bosking, Y. Zhang, B. Schofield, and D. Fitzpatrick. Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. J. Neurosci., 17(6):2112–2127, Mar. 15 1997.
- P. Buzás, M. Volgushev, U. T. Eysel, and Z. F. Kisvárday. Independence of visuotopic representation and orientation map in the visual cortex of the cat. *Eur. J. Neurosci.*, 18(4):957–968, Aug. 2003.
- E. M. Callaway. Local circuits in primary visual cortex of the macaque monkey. Annu. Rev. Neurosci., 21:47–74, 1998.
- M. Á. Carreira-Perpiñán and G. J. Goodhill. Are visual cortex maps optimized for coverage? Neural Computation, 14(7):1545–1560, July 2002a.
- M. Á. Carreira-Perpiñán and G. J. Goodhill. Development of columnar structures in visual cortex. In G. A. Ascoli, editor, *Computational Neuroanatomy: Principles and Methods*, chapter 15, pages 337–357. Humana Press Inc., Totowa, NJ, 2002b.
- M. Á. Carreira-Perpiñán and G. J. Goodhill. Generalized elastic nets. Submitted to Neural Computation, currently in revision., 2004.
- T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, New York, London, Sydney, 1991.
- M. C. Crair, E. S. Ruthazer, D. C. Gillespie, and M. P. Stryker. Ocular dominance peaks at pinwheel center singularities of the orientation map in cat visual cortex. *J. Neurophysiol.*, 77(6):3381–3385, June 1997.
- A. Das and C. D. Gilbert. Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature*, 387(6633):594–598, June 5 1997.
- P. Dayan. Arbitrary elastic topologies and ocular dominance. Neural Computation, 5(3):392-401, 1993.
- V. Dragoi and M. Sur. Dynamic properties of recurrent inhibition in primary visual cortex: Contrast and orientation dependence of contextual effects. J. Neurophysiol., 83(2):1019–1030, Feb. 2000.

- R. Durbin and G. Mitchison. A dimension reduction framework for understanding cortical maps. *Nature*, 343 (6259):644–647, Feb. 15 1990.
- R. Durbin, R. Szeliski, and A. Yuille. An analysis of the elastic net approach to the traveling salesman problem. *Neural Computation*, 1(3):348–358, Fall 1989.
- R. Durbin and D. Willshaw. An analogue approach to the traveling salesman problem using an elastic net method. *Nature*, 326(6114):689–691, Apr. 16 1987.
- U. A. Ernst, K. R. Pawelzik, C. Sahar-Pikielny, and M. V. Tsodyks. Intracortical origin of visual maps. Nat. Neurosci., 4(4):431–436, Apr. 2001.
- E. Erwin and K. D. Miller. Correlation-based development of ocularly matched orientation and ocular dominance maps: Determination of required input activities. J. Neurosci., 18(23):9870–9895, Dec. 1 1998.
- E. Erwin, K. Obermayer, and K. Schulten. Models of orientation and ocular dominance columns in the visual cortex: A critical comparison. *Neural Computation*, 7(3):425–468, May 1995.
- M. Fagiolini, J.-M. Fritschy, K. Löw, H. Möhler, U. Rudolph, and T. K. Hensch. Specific GABA_A circuits for visual cortical plasticity. *Science*, 303(5664):1681–1683, Mar. 12 2004.
- G. J. Goodhill. Topography and ocular dominance: A model exploring positive correlations. *Biol. Cybern.*, 69 (2):109–118, June 1993.
- G. J. Goodhill and A. Cimponeriu. Analysis of the elastic net applied to the formation of ocular dominance and orientation columns. *Network: Computation in Neural Systems*, 11(2):153–168, May 2000.
- G. J. Goodhill and D. J. Willshaw. Application of the elastic net algorithm to the formation of ocular dominance stripes. *Network: Computation in Neural Systems*, 1(1):41–59, Jan. 1990.
- P. J. Green and B. W. Silverman. Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach. Chapman & Hall, London, New York, 1994.
- Y. Hata, T. Tsumoto, H. Sato, and H. Tamura. Horizontal interactions between visual cortical neurones studied by cross-correlation analysis in the cat. J. Physiol., 441(1):593–614, Sept. 1991.
- T. K. Hensch and M. P. Stryker. Columnar architecture sculpted by GABA circuits in developing cat visual cortex. Science, 303(5664):1678–1681, Mar. 12 2004.
- P. A. Hetherington and N. V. Swindale. Receptive field and orientation scatter studied by tetrode recordings in cat area 17. Vis. Neurosci., 16(4):637–652, July 1999.
- M. Hübener, D. Shoham, A. Grinvald, and T. Bonhoeffer. Spatial relationships among three columnar systems in cat area 17. J. Neurosci., 17(23):9270–9284, Dec. 1997.
- K. Kang, M. Shelley, and H. Sompolinsky. Mexican hats and pinwheels in visual cortex. Proc. Natl. Acad. Sci. USA, 100(5):2848–2853, Mar. 4 2003.
- D.-S. Kim, Y. Matsuda, K. Ohki, A. Ajima, and S. Tanaka. Geometrical and topological relationships between multiple functional maps in cat primary visual cortex. *Neuroreport*, 10(12):2515–2522, Aug. 20 1999.
- Z. F. Kisvárday, É. Tóth, M. Rausch, and U. T. Eysel. Orientation-specific relationship between populations of excitatory and inhibitory lateral connections in the visual cortex of the cat. *Cerebral Cortex*, 7(7):605–618, Oct.–Nov. 1997.
- Z. F. Kisvárday, T. Womelsdorf, P. Buzás, and U. T. Eysel. Variability of pinwheel locations in ocular dominance maps according to area and age. Society for Neuroscience Abstracts, 27:285.4, 2001.
- A. A. Koulakov and D. B. Chklovskii. Orientation preference patterns in mammalian visual cortex: A wire length minimization approach. *Neuron*, 29(2):519–527, Feb. 2001.
- S. Löwel, H. J. Bischof, B. Leutenecker, and W. Singer. Topographic relations between ocular dominance and orientation columns in the cat striate cortex. *Experimental Brain Research*, 71(1):33–46, 1988.

- J. S. Lund, A. Angelucci, and P. C. Bressloff. Anatomical substrates for functional columns in macaque monkey primary visual cortex. *Cerebral Cortex*, 13(1):15–24, Jan. 2003.
- H. Meinhardt. Models of Biological Pattern Formation. Academic Press, New York, 1982.
- K. D. Miller, J. B. Keller, and M. P. Stryker. Ocular dominance column development: Analysis and simulation. Science, 245(4918):605–615, Aug. 11 1989.
- K. D. Miller and M. P. Stryker. The development of ocular dominance columns: Mechanisms and models. In S. J. Hanson and C. R. Olson, editors, *Connectionist Modeling and Brain Function: The Developing Interface*, chapter 9, pages 255–350. MIT Press, Cambridge, MA, 1990.
- G. J. Mitchison and N. V. Swindale. Can Hebbian volume learning explain discontinuities in cortical maps? *Neural Computation*, 11(7):1519–1526, Oct. 1999.
- T. Müller, M. Stetter, M. Hübener, F. Sengpiel, T. Bonhoeffer, I. Gödecke, B. Chapman, S. Löwel, and K. Obermayer. An analysis of orientation and ocular dominance patterns in the visual cortex of cats and ferrets. *Neural Computation*, 12(11):2573–2595, Nov. 2000.
- K. Obermayer and G. G. Blasdel. Geometry of orientation and ocular dominance columns in monkey striate cortex. J. Neurosci., 13(10):4114–4129, Oct. 1993.
- K. Obermayer and G. G. Blasdel. Singularities in primate orientation maps. Neural Computation, 9(3):555–575, Apr. 1997.
- C. Piepenbrock, H. Ritter, and K. Obermayer. The joint development of orientation and ocular dominance: Role of constraints. *Neural Computation*, 9(5):959–970, July 1997.
- T. Poggio, V. Torre, and C. Koch. Computational vision and regularization theory. Nature, 317:314–319, 1985.
- B. Roerig and B. Chen. Relationships of local inhibitory and excitatory circuits to orientation preference maps in ferret visual cortex. *Cerebral Cortex*, 12(2):187–198, Feb. 2002.
- K. Rose. Deterministic annealing for clustering, compression, classification, regression, and related optimization problems. Proc. IEEE, 86(11):2210–2239, Nov. 1998.
- S. Schuett, T. Bonhoeffer, and M. Hübener. Mapping retinotopic structure in mouse visual cortex with optical imaging. J. Neurosci., 22(15):6549–6559, Aug. 1 2002.
- J. Sirosh and R. Miikkulainen. Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. *Neural Computation*, 9(3):577–594, Apr. 1997.
- D. C. Somers, E. V. Todorov, A. G. Siapas, L. J. Toth, D.-S. Kim, and M. Sur. A local circuit approach to understanding integration of long-range inputs in primary visual cortex. *Cerebral Cortex*, 8(3):204–217, Apr. 1998.
- D. D. Stettler, A. Das, J. Bennett, and C. D. Gilbert. Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron*, 36(4):739–750, Nov. 14 2002.
- M. Sur and C. A. Leamey. Development and plasticity of cortical areas and networks. Nature Reviews Neuroscience, 2(4):251–262, Apr. 2001.
- N. V. Swindale. Coverage and the design of striate cortex. Biol. Cybern., 65(6):415-424, 1991.
- N. V. Swindale. The development of topography in the visual cortex: A review of models. Network: Computation in Neural Systems, 7(2):161–247, May 1996.
- N. V. Swindale, D. Shoham, A. Grinvald, T. Bonhoeffer, and M. Hübener. Visual cortex maps are optimised for uniform coverage. *Nat. Neurosci.*, 3(8):822–826, Aug. 2000.
- A. Tarantola. Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation. Elsevier Science Publishers B.V., Amsterdam, The Netherlands, 1987.
- A. N. Tikhonov and V. Y. Arsenin. Solutions of Ill-Posed Problems. John Wiley & Sons, New York, London, Sydney, 1977.

- A. M. Turing. The chemical basis of morphogenesis. Proc. R. Soc. Lond. B, 237(641):37–72, Aug. 14 1952.
- C. von der Malsburg. Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, 14:85–100, 1973.
- G. Wahba. Spline Models for Observational Data. SIAM Publ., Philadelphia, PA, 1990.
- L. E. White, W. H. Bosking, and D. Fitzpatrick. Consistent mapping of orientation preference across irregular functional domains in ferret visual cortex. *Vis. Neurosci.*, 18(1):65–76, Jan. 2001.
- D. J. Willshaw and C. von der Malsburg. How patterned neural connections can be set up by self-organization. *Proc. R. Soc. Lond. B*, 194(1117):431–445, Nov. 12 1976.
- A. L. Yuille, J. A. Kolodny, and C. W. Lee. Dimension reduction, generalized deformable models and the development of ocularity and orientation. *Neural Networks*, 9(2):309–319, Mar. 1996.