1

Receptive Fields and Maps in the Visual Cortex: Models of Ocular Dominance and Orientation Columns

Kenneth D. Miller[1]

Published in: *Models of Neural Networks III*, E. Domany, J.L. van Hemmen, and K. Schulten, Eds. (Springer-Verlag, NY), 1995, pp. 55–78.

An earlier and briefer version of this article appeared in The Handbook of Neural Networks (M.A.Arbib, Ed.), The MIT Press,1995, under the title "Models of Ocular Dominance and Orientation Columns." Reused by permission.

ABSTRACT The formation of ocular dominance and orientation columns in the mammalian visual cortex is briefly reviewed. Correlation-based models for their development are then discussed, beginning with the models of Von der Malsburg. For the case of semi-linear models, model behavior is well understood: correlations determine receptive field structure, intracortical interactions determine projective field structure, and the "knitting together" of the two determines the cortical map. This provides a basis for simple but powerful models of ocular dominance and orientation column formation: ocular dominance columns form through a correlation-based competition between left-eye and right-eye inputs, while orientation columns can form through a competition between ON-center and OFF-center inputs. These models account well for receptive field structure, but are not completely adequate to account for the details of cortical map structure. Alternative approaches to map structure, including the self-organizing feature map of Kohonen, are discussed. Finally, theories of the computational function of correlation-based and self-organizing rules are discussed.

[1]Depts. of Physiology and Otolaryngology, W.M. Keck Center for Integrative Neuroscience, and Sloan Center for Theoretical Neurobiology, University of California, San Francisco, CA 94143-0444

Postscript version of this figure
not available; hence figure is
omitted from preprint of paper.

FIGURE 1.1. Schematic of the mature visual system.
Retinal ganglion cells from the two eyes project to separate layers of the lateral geniculate nucleus (LGN). Neurons from these two layers project to separate patches or stripes within layer 4 of visual cortex (V1). Binocular regions (receiving input from both eyes) are depicted at the borders between the eye-specific patches. The cortex is depicted in cross-section, so that layers 1-3 are above and layers 5-6 below the LGN-recipient layer 4. Reprinted by permission from [42]. © 1989 by the AAAS.

1.1 INTRODUCTION

The brain is a learning machine. An animal's experience shapes the neural activity of its brain; this activity in turn modifies the brain, so that the animal learns from its experience. This self-organization, the brain's reshaping of itself through its own activity (reviewed in [7, 14, 39, 51]), has long fascinated neuroscientists and modelers.

The classic example of activity-dependent neural development is the formation of ocular dominance columns in the cat or monkey primary visual cortex (reviewed in [44]). The cerebral cortex is the uniquely mammalian part of the brain. It is thought to form the complex, associative representations that characterize mammalian and human intelligence. The primary visual cortex (V1) is the first cortical area to receive visual information. It receives signals from the lateral geniculate nucleus of the thalamus (LGN), which in turn receives input from the retinae of the two eyes (Fig. 1.1).

To describe ocular dominance columns, several terms must be defined. First, the *receptive field* of a cortical cell refers to the area on the retinae in which appropriate light stimulation evokes a response in the cell, and also to the pattern of light stimulation that evokes such a response. Second, a

Postscript version of this figure
not available; hence figure is
omitted from preprint of paper.

FIGURE 1.2. Ocular dominance columns from cat V1.
A horizontal cut through the layer 4 of V1 is shown. Terminals serving a single eye are labeled white. Dark regions at edges are out of plane containing LGN terminals. Region shown is 5.3 x 7.9 mm. Photograph generously supplied by Dr. Y. Hata.

column is defined as follows. V1 extends many millimeters in each of two, "horizontal" dimensions. Receptive field positions vary continuously along these dimensions, forming a *retinotopic* map, a continuous map of the visual world. In the third, "vertical" dimension, the cortex is about 2 mm in depth, and consists of six layers. Receptive field positions do not significantly vary through this depth. Such organization, in which cortical properties are invariant through the vertical depth of cortex but vary horizontally, is called *columnar* organization and is a basic feature of cerebral cortex.

Third, *ocular dominance* must be defined. Cells in the LGN are *monocular*, responding exclusively to stimulation of a single eye (Fig. 1.1). LGN cells project to layer 4 of V1, where they terminate in alternating stripes or patches of terminals representing a single eye (Figs. 1.1, 1.2). Most or, in some species, all layer 4 V1 cells are monocular. Cells in other layers of V1 respond best to the eye that dominates layer 4 responses at that horizontal location. Thus, V1 cells can be characterized by their *ocular dominance*, or eye preference. The stripes or patches of cortex that are dominated throughout the cortical depth by a single eye are known as *ocular dominance columns.*

The segregated pattern of termination of the LGN inputs to V1 arises early in development. Initially, LGN inputs project to layer 4 of V1 in an overlapping manner, without apparent distinction by eye represented. The terminal arbors of individual LGN inputs extend horizontally in layer 4 for distances as large as 2 mm (for comparison, a typical spacing between cortical cells is perhaps 20 μm). Subsequently, beginning either prenatally or shortly after birth depending on the species, the inputs representing each eye become horizontally confined to the alternating, approximately 1/2-mm wide ocular dominance patches.

This segregation results from an activity-dependent competition between the geniculate terminals serving the two eyes (see discussion in [44]). The signal indicating that different terminals represent the same eye appears to be the correlations in their neural activities [54]. These correlations exist due both to spontaneous activity, which is locally correlated within each retina [36, 37, 38, 64], and to visually-induced activity, which correlates the activities of retinotopically nearby neurons within each eye and, to a lesser extent, between the eyes [26]. The segregation process is competitive. If one eye is caused to have less activity than the other during a critical period in which the columns are forming, the more active eye takes over most of the cortical territory [25, 52, 60]; but the eye with reduced activity suffers no loss of projection strength in retinotopic regions in which it lacks competition from the other eye [15, 16]. In summary, ocular dominance column formation is a simple system in which correlated patterns of neural activity sculpt the patterns of neural connectivity.

Orientation columns are another striking feature of visual cortical organization. Most V1 cells are orientation selective, responding selectively to light/dark edges over a narrow range of orientations. The preferred orientation of cortical cells varies regularly and periodically across the horizontal dimension of cortex, and is invariant in the vertical dimension. The maturation of orientation selectivity is activity-dependent (e.g., [6, 11]). However, it has not yet been possible to test whether the initial development of orientation selectivity is activity-dependent. This is because some orientation selectivity already exists at the earliest developmental times at which visual cortical responses can be recorded [1, 4, 6, 20, 61], and it has not been possible to block visual system activity immediately before this time. Nonetheless, it has long been a popular notion that the initial development of orientation selectivity, like that of ocular dominance, may occur through a process of activity-dependent synaptic competition.

The inputs from LGN to V1 serving each eye are of two types: ON-center and OFF-center. Both kinds of cells have circularly symmetric, orientation-insensitive receptive fields, and respond to contrast rather than uniform luminance. ON-center cells respond to light against a dark background, or to light onset; OFF-center cells respond to dark against a light background, or to light offset. In the cat, the orientation-selective V1 cells in layer 4 are *simple cells*: cells with receptive fields consisting of alternating oriented subregions that receive exclusively ON-center or exclusively OFF-center input (Fig. 1.3). As shall be discussed, one theory for the development of orientation selectivity is that, like ocular dominance, it develops through a competition between two input populations: in this case, a competition

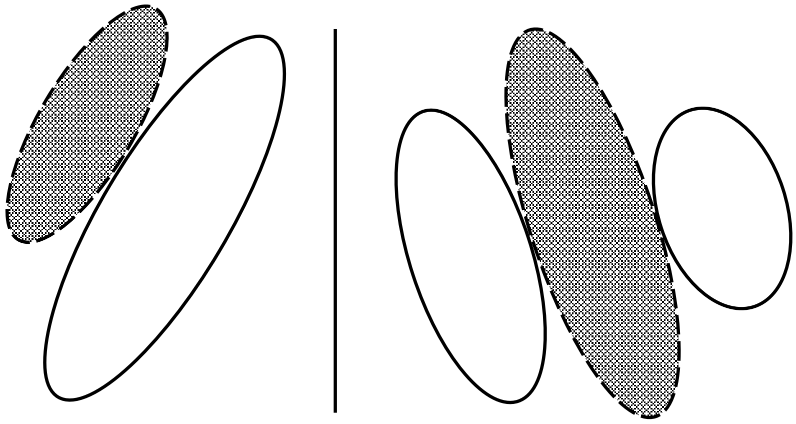

FIGURE 1.3. Two examples of simple cell receptive fields (RFs). Regions of the visual field from which a simple cell receives ON-center (white) or OFF-center (dark) input are shown. Note: Ocular dominance columns (Fig. 1.2) represent an alternation, across cortex, in the type of input (left- or right-eye) received by different cortical cells; while a simple cell RF (this figure) represents an alternation across visual space in the type of input (ON- or OFF-center) received by a *single* cortical cell.

between the ON-center and the OFF-center inputs [41].

1.2 CORRELATION-BASED MODELS

To understand ocular dominance and orientation column formation, two processes must be understood: (1) the development of *receptive field structure*: under what conditions do receptive fields become monocular (drivable only by a single eye) or orientation selective? (2) the development of *periodic cortical maps* of receptive field properties: what leads ocular dominance or preferred orientation to vary periodically across the horizontal dimensions of cortex, and what determines the periodic length scales of these maps? Typically, the problem is simplified by consideration of a two-dimensional model cortex, ignoring the third dimension in which properties such as ocular dominance and orientation are invariant.

One approach to addressing these problems is to begin with a hypothesized mechanism of synaptic plasticity, and to study the outcome of cortical development under such a mechanism. Most commonly, theorists have considered a “Hebbian synapse”: a synapse whose strength is increased when pre- and post-synaptic firing are correlated, and possibly decreased when they are anticorrelated. Other mechanisms, such as activity-dependent release and uptake of a diffusible modification factor, can lead to similar dynamics [42], in which synaptic plasticity depends on the correlations among the activities of competing inputs. Models based on such mechanisms are referred to as "correlation-based models" [39].

1.2.1 The Von Der Malsburg Model Of V1 Development

Von der Malsburg [57, 59] first formulated a correlation-based model for the development of visual cortical receptive fields and maps. His model had two basic elements. First, synapses of LGN inputs onto cortical neurons were modified by a Hebbian rule that is *competitive*, so that some synapses were strengthened only at the expense of others. He enforced the competition by holding constant the total strength of synapses converging on each cortical cell (conservation rule). Second, cortical cells tended to be activated in *clusters*, due to intrinsic cortical connectivity, *e.g.* short-range horizontal excitatory connections and longer-range horizontal inhibitory connections.

The conservation rule leads to competition among the inputs to a single target cell. Inputs that tend to be coactivated — that is, that have correlated activities — are mutually reinforcing, working together to activate the postsynaptic cells and thus to strengthen their own synapses. Different patterns that are mutually un- or anti-correlated compete, since strengthening of some synapses means weakening of others. Cortical cells eventually develop receptive fields responsive to a correlated pattern of inputs.

The clustered cortical activity patterns lead to competition between different groups of cortical cells. Each input pattern comes to be associated with a cortical cluster of activity. Overlapping cortical clusters contain many coactivated cortical cells, and thus become responsive to overlapping, correlated input patterns. Adjacent, non-overlapping clusters contain many anti-correlated cortical cells, and thus become responsive to un- or anti-correlated input patterns. Thus, over distances on the scale of an activity cluster, cortical cells will have similar response properties; while, on the scale of the distance between non-overlapping clusters, cortical cells will prefer un- or anti-correlated input patterns. This combination of local continuity and larger-scale heterogeneity leads to continuous, periodic cortical maps of receptive field properties.

In computer simulations, this model was applied to the development of orientation columns [57] and ocular dominance columns [59]. For orientation columns, inputs were activated in oriented patterns of all possible orientations. Individual cortical cells then developed selective responses, preferring one such oriented pattern, with nearby cortical cells preferring nearby orientations. For ocular dominance columns, inputs were activated in monocular patterns consisting of a localized set of inputs from a single eye. Individual cortical cells came to be driven exclusively by a single eye, and clusters of cortical cells came to be driven by the same eye. The final cortical pattern consisted of alternating stripes of cortical cells preferring a single eye, with the width of a stripe approximately set by the diameter of an intrinsic cluster of cortical activity.

In summary, a competitive Hebbian rule leads individual receptive fields to become selective for a correlated pattern of inputs. Combined with the idea that the cortex is activated in intrinsic clusters, this suggests an origin for cortical maps: coactivated cells in a cortical cluster tend to become selective for similar, coactivated patterns of inputs. These basic ideas are used in most subsequent models.

1.2.2 Mathematical Formulation

A typical correlation-based model is mathematically formulated as follows [57, 27, 40, 42]. Let $x, y, \ldots$ represent retinotopic positions in V1, and let $\alpha, \beta, \ldots$ represent retinotopic positions in the LGN. Let $S^\mu(x,\alpha)$ be the synaptic strength of the connection from α to x of the LGN projection of type μ, where μ may signify left-eye, right-eye, ON-center, OFF-center, etc. Let $B(x,y)$ represent the synaptic strength and sign of connection from the cortical cell at y to that at x. For simplicity, $B(x,y)$ is assumed to take different signs for a fixed y as x varies, but alternatively, separate excitatory-projecting and inhibitory-projecting cortical neurons may be used. Let $a(x)$ and $a^\mu(\alpha)$ represent the activity of a cortical or LGN cell, respectively.

The activity $a(x)$ of a cortical neuron is assumed to depend on a linear combination of its inputs:

$$a(x) = f_1\left(\sum_{\mu,\alpha} S^\mu(x,\alpha)a^\mu(\alpha) + \sum_y B(x,y)a(y)\right) \tag{1.1}$$

Here, f_1 is some monotonic function such as a sigmoid or linear threshold.

A Hebbian rule for the change in feedforward synapses can be expressed

$$\Delta S^\mu(x,\alpha) = A^\mu(x,\alpha) f_2\left[a(x)\right] f_3\left[a^\mu(\alpha)\right] \tag{1.2}$$

Here, $A(x,\alpha)$ is an "arbor function", expressing the number of synapses of each type from α to x; a minimal form is $A(x,\alpha)=1$ if there is a connection from α to x, $A(x,\alpha)=0$ otherwise. A typical form for the functions f_2 and f_3 is $f(a)=(a-\langle a\rangle)$, where $\langle a\rangle$ indicates an average of a over input patterns. This yields a *covariance rule*: synaptic change depends on the covariance of postsynaptic and presynaptic activity.

Next, the Hebbian rule must be made *competitive*. This can be accomplished by conserving total synaptic strength over the postsynaptic cell [57], which in turn may be done either subtractively or multiplicatively

[43]. The corresponding equations are

$$\frac{d}{dt}S^{\mu}(x,\alpha) = \Delta S^{\mu}(x,\alpha) - \epsilon(x)A(x,\alpha) \qquad \text{(Subtractive)} \tag{1.3}$$

$$\frac{d}{dt}S^{\mu}(x,\alpha) = \Delta S^{\mu}(x,\alpha) - \gamma(x)S^{\mu}(x,\alpha) \qquad \text{(Multiplicative)} \tag{1.4}$$

where $\epsilon(x) = \frac{\sum_{\kappa,\alpha}\Delta S^{\kappa}(x,\alpha)}{\sum_{\kappa,\alpha}A(x,\alpha)}$, and $\gamma(x) = \frac{\sum_{\kappa,\alpha}\Delta S^{\kappa}(x,\alpha)}{\sum_{\kappa,\alpha}S^{\kappa}(x,\alpha)}$. Either form of constraint ensures that $\sum_{\mu,\alpha}\frac{d}{dt}S^{\mu}(x,\alpha) = 0$. Alternative methods have been developed to force Hebbian rules to be competitive (see Miller and MacKay, 1994).

Finally, synaptic weights may be limited to a finite range, $s_{\min}A(x,\alpha) \leq S^{\mu}(x,\alpha) \leq s_{\max}A(x,\alpha)$. Typically, $s_{\min} = 0$ and $s_{\max}$ is some positive constant.

1.2.3 Semi-Linear Models

In semi-linear models, the f's in (1.1-1.2) are chosen to be linear. Then, after substituting for $a(x)$ from (1.1) and averaging over input patterns (assuming that all inputs have identical mean activity, and that changes in synaptic weights are negligibly small over the averaging time), (1.2) becomes

$$\Delta S^{\mu}(x,\alpha) = \lambda A(x,\alpha)\left[\sum_{y,\beta,\kappa} I(x-y)\left[C^{\mu\kappa}(\alpha-\beta) - k_2\right]S^{\kappa}(y,\beta) + k_1\right] \tag{1.5}$$

Here, $I(x-y)$ is an element of the intracortical interaction matrix

$$\mathbf{I} \equiv (\mathbf{1}-\mathbf{B})^{-1} = \mathbf{1} + \mathbf{B} + \mathbf{B}^2 + \ldots,$$

where the matrix $\mathbf{B}$ is defined in (1.1). This summarizes intracortical synaptic influences including contributions via $0,1,2,\ldots$ synapses. The covariance matrix

$$C^{\mu\kappa}(\alpha-\beta) = \langle (a^{\mu}(\alpha) - \bar{a})\,(a^{\kappa}(\beta) - \bar{a})\rangle$$

expresses the covariation of input activities. The factors λ, k_1 and k_2 are constants. Translation invariance has been assumed in both cortex and LGN.

When there are two competing input populations, (1.5) can be further simplified by transforming to sum and difference variables: $S^S \equiv S^1 + S^2$, $S^D \equiv S^1 - S^2$. Assuming equivalence of the two populations (so that $C^{11} = C^{22}$, $C^{12} = C^{21}$), (1.5) becomes

$$\Delta S^S(x,\alpha) = \lambda A(x,\alpha)\left\{\sum_{y,\beta} I(x-y)\left[C^S(\alpha-\beta) - 2k_2\right]S^S(y,\beta) + 2k_1\right\} \tag{1.6}$$

$$\Delta S^D(x,\alpha) = \lambda A(x,\alpha)\sum_{y,\beta} I(x-y)C^D(\alpha-\beta)S^D(y,\beta) \tag{1.7}$$

Here, $C^S \equiv C^{11} + C^{12}$, $C^D \equiv C^{11} - C^{12}$. Subtractive renormalization (1.3) alters only (1.6) for S^S, by subtraction of $2\epsilon(x)A(x-\alpha)$, while leaving (1.7) for S^D unaltered. Multiplicative renormalization (1.4) alters both (1.6) and (1.7), by subtraction of $\gamma(x)S^S(x,\alpha)$ and $\gamma(x)S^D(x,\alpha)$ respectively.

1.2.4 How Semi-Linear Models Behave

Linear equations like (1.6) and (1.7) can be understood by finding the eigenvectors or "modes" of the operators on the right side of the equation. The eigenvectors are the synaptic weight patterns that grow independently and exponentially, each at its own rate. The fastest-growing eigenvectors typically dominate development and determine basic features of the final pattern, although the final pattern ultimately is stabilized by nonlinearities such as the limits on the range of synaptic weights or the nonlinearity involved in multiplicative renormalization (1.4).

I will focus on the behavior of (1.7) for S^D (for analysis of (1.6), see [34, 35]). S^D describes the difference in the strength of two competing input populations. Thus, it is the key variable describing the development of ocular dominance segregation, or development under an ON-center/OFF-center competition. In many circumstances, (1.7) can be derived directly from (1.1-1.2) by linearization about $S^D \equiv 0$ [40], without need to assume a semi-linear model. The condition $S^D \approx 0$ corresponds to an initial condition in which the projections of the two input types are approximately equal. Thus, study of (1.7) can lend insight into early pattern formation in more general, nonlinear correlation-based models.

Equation 1.7 can be simply solved in the case of full connectivity from the LGN to the cortex, when $A(x,\alpha) \equiv 1$ for all x and α. Then modes of $S^D(x,\alpha)$ of the form $e^{ikx}e^{il\alpha}$ grow exponentially and independently, with rate proportional to $\tilde{I}(k)\tilde{C}^D(l)$, where $\tilde{I}$ and $\tilde{C}^D$ denote the Fourier transforms of I and C^D, respectively (for a description of the modes as real rather than complex functions, see [44]). The wavenumber k determines the wavelength $2\pi/|k|$ of an oscillation of S^D across cortical cells, while the wavenumber l determines the wavelength $2\pi/|l|$ of an oscillation of S^D across geniculate cells. The fastest-growing modes, which will dominate early development, are determined by the k and l that maximize $\tilde{I}(k)$ and $\tilde{C}^D(l)$ respectively. The peak of a function's Fourier transform corresponds to the cosine wave that best matches the function, and thus represents the "principal oscillation" in the function.

To understand these modes (Fig. 1.4), consider first the set of inputs received by a single cortical cell, that is, the shape of the mode for a fixed cortical position x. This can be regarded as the "receptive field" of the cortical cell. Each receptive field oscillates with wavenumber l. This oscillation, of $S^D \equiv S^1 - S^2$, is an oscillation between receptive field subregions dominated by S^1 inputs and subregions dominated by S^2 inputs. Thus, in ocular dominance competition, monocular cells (cells whose entire receptive fields are dominated by a single eye) are formed only by modes with $l = 0$ (no oscillation). Monocular cells thus dominate development if the peak of the Fourier transform of the C^D governing left/right competition is at $l = 0$. Now instead consider an ON/OFF competition: S^1 and S^2 represent ON- and OFF-center inputs from a single eye. Then the receptive fields of modes with nonzero l resemble simple cells: they receive predominantly ON-center and predominantly OFF-center inputs from successive, alternating subregions of the visual world. Thus, simple cells can form if the C^D governing ON/OFF competition has its peak at a nonzero l.

Now consider the arborizations or "projective fields" projecting from a single geniculate point, that is, the shape of the mode for a fixed geniculate position α. These oscillate with wavenumber k. In ocular dominance competition, this means that left- and right-eye cells from α project to alternating patches of cortex. When monocular cells form ($l = 0$), these alternating patches of cortex are the ocular dominance columns: alternating patches of cortex receiving exclusively left-eye or exclusively right-eye input respectively. Thus, the width of ocular dominance columns — the wavelength of alternation between right-eye and left-eye dominated cortical cells — is determined by the peak of the Fourier transform of the intracortical interaction function I. In ON/OFF competition, with $l \neq 0$, the identity of the cortical cells receiving the ON-center or OFF-center part of the projection varies as α varies, so individual cortical cells receive both ON- and OFF-center input, but from distinct subregions of the receptive field.

In summary, there is an oscillation within receptive fields, with wavenumber l determined by the peak of $\tilde{C}^D$; and an oscillation within arbors, with wavenumber k determined by the peak of $\tilde{I}$ (Fig. 1.4). These two oscillations are "knit together" to determine the overall pattern of synaptic connectivity. The receptive field oscillation, which matches the receptive field to the correlations, quantitatively describes von der Malsburg's finding that individual receptive fields become selective for a correlated pattern of inputs. Similarly, the arbor oscillation matches projective fields to the intracortical interactions, and thus to the patterns of cortical activity clusters. This quantitatively describes the relationship between activity clusters and maps. Note that the factor e^{ikx} can be regarded as inducing a phase shift, for varying x, in the structure of receptive fields. Thus, cortical cells that are nearby on the scale of the arbor oscillation have similar receptive fields, while cells 1/2 wavelength apart have opposite receptive fields.

An alternative viewpoint on the same pattern is obtained by rewriting the modes as $e^{i(k+l)x}e^{-il(x-\alpha)}$. The argument $l(x - \alpha)$ represents the oscillation with wavenumber l within the receptive field, now expressed in coordinates relative to the center of the receptive field rather than in absolute position across the geniculate. The argument $(k+l)x$ represents a shift, for varying x, in the phase of the receptive field relative to the receptive field center. For the case of ocular dominance, with $l = 0$, this is just the

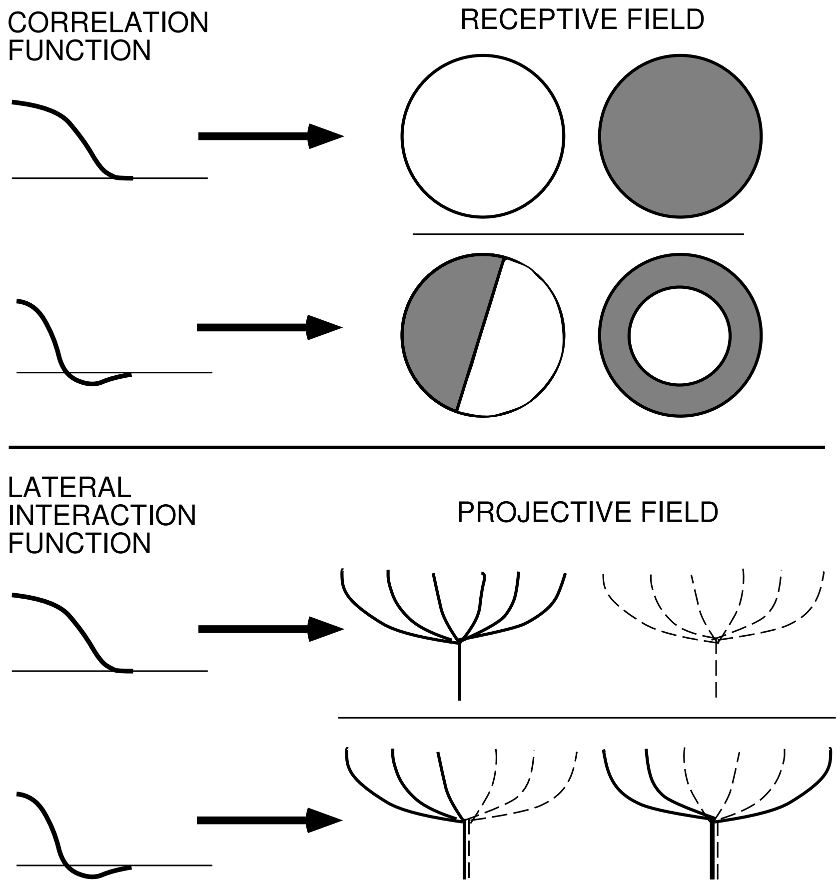

FIGURE 1.4. Schematic of the outcome of semi-linear correlation-based development.
TOP: The correlation function (C^D) determines the structure of receptive fields (RFs). White RF subregions indicate positive values of S^D; dark subregions, negative values. When C^D does not oscillate, individual cortical cells receive only a single type of input, as in ocular dominance segregation. If C^D oscillates, there is a corresponding oscillation in the type of input received by individual cortical cells, as in simple cell RFs. Alternative RF structures could form, as in the center-surround structure shown; but oriented simple-cell-like outcomes predominate for reasonable parameters [41]. Simple cells then develop with various numbers of subregions and various spatial phases; only a single example, of a cell with two subregions and odd spatial symmetry, is pictured. BOTTOM: The intracortical interactions (I) similarly determine the structure of projective fields. Here, solid lines indicate positive values of S^D, dotted lines indicate negative values. Adapted from [43].

shift, with wavenumber k, between left-eye dominance and right-eye dominance of cortical cells. For ON/OFF competition with $l \neq 0$, this represents a periodic shifting, with movement across cortex, as to which subregions of the receptive field are dominated by ON-center inputs, and which subregions are dominated by OFF-center inputs. Thus, we can view the results as an oscillation within receptive fields, with wavenumber l; combined with a shift with cortical position in the spatial phase of receptive fields, this shift occurring with wavenumber $k + l$, the vector sum of the projective field or arbor oscillation and the receptive field oscillation.

The competitive, renormalizing terms (1.3–1.4) do not substantially alter these pictures, except that multiplicative renormalization can suppress ocular dominance development in some circumstances [43].[2] These results hold also for localized connectivity (finite arbors), and thus generally characterize the behavior of semi-linear models [39, 44]. The major difference in the case of localized connectivity is that, if k or l corresponds to a wavelength larger than the diameter of connectivity from or to a single cell, then it is equivalent to $k = 0$ or $l = 0$ respectively. A good approximation to the leading eigenvectors in the case of finite connectivity is given simply by $A(x - \alpha)e^{ikx}e^{il\alpha}$, where k and l are determined as above by the peaks of $\tilde{I}(k)$ and $\tilde{C}^D(l)$ (unpublished results).

1.2.5 Understanding Ocular Dominance And Orientation Columns With Semilinear Models

This understanding of semilinear models leads to simple models for the development of both ocular dominance columns [42] and orientation columns [41], as follows (Fig. 1.4).

Monocular cells develop through a competition of left- and right-eye inputs in a regime in which $\tilde{C}^D(l)$ is peaked at $l = 0$. The wavelength of ocular dominance column alternation is then determined by the peak of $\tilde{I}(k)$.

Orientation-selective simple cells develop through a competition of ON-center and OFF-center inputs in a regime in which $\tilde{C}^D(l)$ is peaked at

[2]Subtractive renormalization (1.3) has no effect on the development of S^D. Multiplicative renormalization (1.4) lowers the growth rates of all modes of both S^D and S^S by the factor $\gamma(x)$, which depends only on S^S. The result is that, in order for S^D to grow at all, its modes must have larger unconstrained growth rates than those of S^S; that is, the peak of the Fourier transform of C^D must be larger than that of C^S. In practice, this condition is met only if there are anticorrelations between S^1 and S^2, that is, if C^{12} is significantly negative. When this condition is met, then the modes that dominate S^D are just as described above; they are not altered by the constraint term in 1.4. These and other effects of renormalizing terms are discussed in detail in [43].

$l \neq 0$. The mean wavelength of alternation of ON-center and OFF-center subregions in the simple cells' receptive fields is determined by the peak of $\tilde{C}^D(l)$. This wavelength corresponds to a cell's preferred spatial frequency under stimulation by sinusoidal luminance gratings. In individual modes, all cortical cells have the same preferred orientation, but their spatial phase varies periodically with cortical position. The mixing of such modes of all orientations leads to a periodic variation of preferred orientation across cortex. The period with which preferred orientations change across cortex is more complex to determine [41].

This model of ocular dominance column formation is similar to that of von der Malsburg [59]. The latter model assumed anti-correlation between the two eyes; this was required due to the use of multiplicative renormalization (1.4). With subtractive renormalization (1.3), ocular dominance column formation can occur even with partial correlation of the two eyes [43]. The model can be compared to experiment, particularly through the prediction of the relation between intracortical connectivity and ocular dominance column width.

The model of orientation-selective cell development is quite different from that of von der Malsburg [57]. Von der Malsburg postulated that oriented input patterns lead to the development of orientation-selective cells. The ON/OFF model instead postulates that ON/OFF competition results in oriented receptive fields, in the absence of oriented input patterns; the circular symmetry of the input patterns is spontaneously broken. This symmetry-breaking potential of Hebbian development was first discovered by Linsker [28]. In all of these models, the continuity and periodic alternation of preferred orientation is due to the intracortical connectivity. The ON/OFF model can be compared to experiment most simply by the measurement of C^D, to determine whether it has the predicted oscillation.

1.2.6 Related Semi-Linear Models

Linsker [27, 28, 29] proposed a model that was highly influential in two respects. First, he pointed out the potential of Hebbian rules to spontaneously break symmetry, yielding orientation-selective cells given approximately circularly symmetric input patterns. Second, he demonstrated that Hebbian rules could lead to segregation *within* receptive fields, so that a cell came to receive purely excitatory or purely inhibitory input in alternating subregions of the receptive field. This model was thoroughly analyzed in [34, 35].

Linsker used a semi-linear model, with a single input type that could have positive or negative synaptic strengths ($s_{\min} = -s_{\max}$). He largely restricted study to the case of a single postsynaptic cell. Because the equation for a single input type and a single postsynaptic cell (Eq. 1.5, with

$I(x-y) = \delta(x-y)$), is circularly symmetric,[3] its eigenfunctions are also eigenfunctions of the rotation operator. Thus, the eigenfunctions can be written in polar coordinates (r, θ) as $cos(n\theta)f_{nj}(r)$, $sin(n\theta)f_{nj}(r)$, where $f_{nj}(r)$ is a radial function and n and j are integers indexing the eigenfunctions. In quantum mechanics, atomic orbitals are named Nx, where N is a number representing one plus the total number of angular and radial nodes, and x is a letter denoting the number of angular nodes (s,p,d,f,g,... corresponding to n=0,1,2,3,4,... angular nodes). Thus, 1s is a function with zero nodes, 2s has one node which is radial, 2p has one node which is angular, 3p has two nodes (one radial, one angular), etc. This naming scheme can be applied to any rotationally symmetric system, and in particular can be applied to the eigenfunctions of Linsker's system [34, 35], a fact which physicists have found amusing. The nature of these eigenfunctions, their dependence on parameters, and their role in determining the outcomes Linsker observed in simulations are described in [34, 35].

For our present purposes, the essential results of this analysis are as follows. Two factors underlay Linsker's results. One factor was that oscillations in a correlation function can induce oscillations in a receptive field, as described above. The other factor was a constraint in the model fixing the percentage of positive or negative synapses received by a cell; this forced an alternation of positive and negative subregions even when the correlation function did not oscillate. These two causes were not disentangled in Linsker's simulations, but only the first appears likely to be of biological relevance.

Tanaka [45, 56] has independently formulated models of ocular dominance and orientation columns that are similar to those described in section 1.2.5. The major difference is that he works in a regime in which each cortical cell comes to receive only a single LGN input. Tanaka defines cortical receptive fields as the convolution of the input arrangement with the intracortical interaction function. This means that a cortical cell's receptive field is due to its single input from the LGN, plus its input from all other cortical cells within reach of the intracortical interaction function. Thus, orientation selectivity in this model arises from the breaking of circular symmetry in the pattern of inputs to different cortical cells, rather than to individual cortical cells.

1.3 THE PROBLEM OF MAP STRUCTURE

The above models account well for basic features of primary visual cortex. However, many details of real cortical maps are not replicated by these

[3]The assumption is made that the arbor and correlation functions depend only on distance.

models [9, 12, 63]. One reason may be the simplicity of the model of cortex: the real cortex is 3-dimensional, rather than 2; has cell-specific connectivity, rather than connectivity that depends only on distance; and has plastic rather than fixed intracortical connections. Another reason is that the details of map structure inherently involve nonlinearities, by which the fastest-growing modes interact and compete; whereas the semi-linear framework only focuses on early pattern formation, in which the fastest-growing modes emerge and mix randomly without interacting.

Some simple models that focus on map development rather than receptive field development strikingly match the map structures observed in monkeys [9]. One such model [46] uses the self-organizing feature map (SOFM) of Kohonen [24, 48]. In the SOFM, only a single cluster of cortical cells is activated in response to a given input pattern. This is an abstraction of the idea that the cortex responds in localized activity clusters. The single activated cluster is centered on the cell whose weight vector best matches the direction of the input activation vector. Hebbian learning then takes place on the activated cells, bringing their weight vector closer to the input activation vector. The size of an activity cluster is gradually decreased as the mapping develops; this is akin to annealing, helping to ensure a final mapping that is optimal on both coarse and fine scales.

Except for the restriction to a single activity cluster and the gradual decrease in cluster size, the SOFM is much like the correlation-based models. However, an abstract representation of the input is generally used. In correlation-based models, the input space may have thousands of dimensions, one for each input cell. In the SOFM model of visual cortex, the input space instead has five dimensions: two represent retinotopic position, and one represents each of ocular dominance, orientation selectivity, and preferred orientation. Each cortical cell receives five "synapses", corresponding to these five "inputs". Assumptions are made as to the relative "size" of, or variance of the input ensemble along, each dimension. There is no obvious biological interpretation for this comparison between dimensions. Under the assumptions that the ocular dominance and orientation dimensions are "short" compared to the retinotopic dimensions, and that only one input point is activated at a time, Hebbian learning can lead to maps of orientation and ocular dominance that are, in detail, remarkably like those seen in macaque monkeys [9, 46].

The SOFM, and other models based on the "elastic net" algorithm [8, 13], lead to locally continuous mappings in which a constant distance across the cortex corresponds to a roughly constant distance in the reduced "input space". This means that, when one input feature is changing rapidly across cortex, the others are changing slowly. Thus, the models predict that orientation changes rapidly where ocular dominance changes slowly, and vice versa. It may be this feature that is key to replicating the details of macaque orientation and ocular dominance maps. A model that forces such a relationship to develop between ocular dominance and orientation, while assuring periodic representation of each, also gives a good match to primate visual maps [55].

The SOFM also replicates aspects of the retinotopic maps seen in higher areas of cat visual cortex [62]. For these studies, the input and output space are each taken to be two-dimensional, representing retinotopic position. The input space is taken to be a half-circle, representing a hemi-retina, and the shape of the output space is varied. When this shape is long and narrow, as in cat cortical areas 18 and 19, the retinotopic map developed by the SOFM has a characteristic pattern of discontinuities closely resembling those observed experimentally in those areas [62]. The SOFM achieves maps in which nearby points in the output space correspond to nearby points in the input space, while each area of the input space receives approximately equal representation provided each is equally activated ([48]; see further discussion of the SOFM below). The success of the SOFM models of retinotopic maps suggests that these are constraints that should be satisfied by any model of cortical maps. One would like to determine more precisely the constraints on a retinotopic mapping, embodied by the SOFM, that are sufficient to replicate these results.

It has recently been reported that input correlations can alter the spacing of ocular dominance columns in cat visual cortex by perhaps 20-30% [32]. A smaller ocular dominance column spacing develops when the activities of the two eyes are correlated by normal vision than when the two eyes' activities are decorrelated (decorrelation is achieved by inducing divergent strabismus, which causes the two eyes to see different parts of the visual world).[4] This effect was anticipated theoretically by Goodhill [12], who argued essentially that correlation of the activities of the two eyes brings them "closer together", and so the two eyes should be brought closer together in their cortical representation by reduction of column size. This effect could also have been anticipated by the SOFM models of ocular dominance, because decorrelation corresponds to an increase in the variance of ocular dominance and thus an increase in the "size" of the ocular dominance dimension, which results in increased column size [48]. In semi-linear models, in contrast, the column width does not appear to be significantly affected by between-eye correlations. Rather, as the degree of between-eye correlation is increased, more binocular cells form at the column borders, until at some critical level of correlation ocular dominance

[4]This story is still somewhat tentative. Another worker (K. Murphy, private communication) reports no difference between the column spacings found in normal and strabismic cats. She reports that in both cases, column periods range from 800-1300 μm, whereas Löwel [32] reports that normal cats have periods of 800-1000 μm and strabismic cats have periods of 1100-1300 μm. Neither study used litter-matched controls; given natural variation in ocular dominance column width, such studies may be needed to determine the strength and robustness of a strabismic effect.

segregation no longer occurs (unpublished results). That is, the two eyes are brought "closer together" through alteration of receptive fields rather than through alteration of the map. One can anticipate several biological mechanisms that might be added to instead yield reduction in column size, such as nonlinearities that discourage formation of binocular cells, or nonlinearities in cortical activation that cause the size of activity clusters to depend on the correlations of the inputs.[5]

Finally, it has recently been noted that cat orientation maps are significantly smoother than could be achieved by simple linear considerations [63]. The analysis in [63] suggests that these maps could result, mathematically, from a local "diffusion" of preferred orientations. It will be interesting to develop a biologically interpretable model of such a process.

1.4 THE COMPUTATIONAL SIGNIFICANCE OF CORRELATION-BASED RULES

1.4.1 Efficient Representation of Information

A simple correlation-based rule for a single postsynaptic cell can, if properly designed, lead to development of a receptive field that corresponds to the principal component of the input data (that is, to the principal eigenvector of the covariance matrix of the inputs to the cell) [30, 43, 47]. This receptive field in turn maximizes the variance of the postsynaptic cell's activity, given the ensemble of input patterns. It has been argued that correlation-based rules thus maximize the information carried in the postsynaptic cell's activity about the input patterns [30]. Intuitively, by varying as much as possible in its response to different inputs, the postsynaptic cell draws the greatest possible distinction between the different input patterns.

More generally, a number of closely related (and in many circumstances identical) computational functions have been proposed for brain areas near the sensory periphery. These include maximization of information about the inputs [30], minimization of redundancy or correlation in the activities of output cells [3], statistical independence of the output activities [3], or encoding of the input information as compactly as possible (for example, requiring as little dynamic range as possible per neuron) [2]. These func-

[5]Goodhill's theory [12] is identical to the model of [42], except that the translation from input activations to cortical activations has been altered. Miller *et al.* [42] used linear activation rules, whereas Goodhill [12] uses the SOFM rule [24] in which a single activity bubble of cells surrounds the best-activated or "winning" cortical cell; also, Goodhill modifies the rule so that the more often a cell has won, the less likely it is to win again. Thus, use of nonlinear intracortical interactions can transform a semi-linear model [42] into one in which inter-ocular correlation modulates column width.

tions all involve representing the input information in an efficient way, in the sense of information theory. These measures of efficiency take into account the statistics of the input ensemble, but disregard the "semantics", the meaning or survival value to the animal, of the inputs.

The interpretation that the function of a correlation-based rule is to yield such an efficient representation is inviting, but it carries two major problems. First, the principal component representation achieved by correlation-based rules is optimally efficient only for a Gaussian distribution of input patterns, or in other words, it reflects only the second order or two-point statistics (the covariance) of the input data. It is possible that a great deal of information may reside in higher-order statistics, but a correlation-based rule as conceived above will ignore this information. Intrator has suggested that a variant of standard Hebbian rules can instead maximize a third-order statistic of the output activity, and argues that this may be a better statistic for distinguishing among the elements of real-world ensembles [22, 23]. While one statistic or the other may be best for characterizing a given set of data, both approaches can suffer from the limitation that they are maximizing one particular statistic rather than maximizing some measure of efficiency.

Second, this interpretation applies only to a single, isolated postsynaptic cell. Multiple cells viewing the same input ensemble will extract the same information from it under a given correlation-based rule. This does not add new information about the input, but only redundantly repeats the same information. Thus, although a single cell may have a receptive field that maximizes the information it could carry about the input ensemble, a group of such cells generally will not much improve on the performance of a single cell and will not carry the maximal possible information about the input ensemble.[6]

One way out of this dilemna is to introduce couplings between the postsynaptic cells that force them to learn independent parts of the input ensemble. Unfortunately, excitatory couplings tend to produce correlated cells, while inhibitory couplings produce anti-correlated cells. The ostensible goal, however, is to produce uncorrelated cells, cells whose activities carry independent information about the input ensemble. Thus, biological couplings will not work. A theoretical way out involves using connections between the postsynaptic cells that are modified by anti-Hebbian rules: if two cells have correlated activity, the connection between them becomes more negative; if two cells have anti-correlated activity, the connection between them

[6]For simplicity, in this discussion I am ignoring noise. Depending on the signal to noise ratio, one will wish to strike a particular balance between variety (carrying more independent components of the input ensemble) and redundancy (e.g., see [2, 30]). However, except in the extreme case of high noise, where complete redundancy is called for, multiple components will always be needed to maximize the information, given multiple output cells.

becomes more positive. The result is that the cells become uncorrelated. Many authors have independently proposed rules that involve such anti-Hebbian learning on lateral connections (e.g., [10, 31, 49]) or related ideas [50]. However, no biological sign of anti-Hebbian synaptic modification has thus far been observed.

An alternative way out of this dilemna stems from the observation that biological receptive fields are localized. Thus, nearby cells see overlapping but not identical sets of inputs. Consider two extreme cases. First, when each input cell is connected to a single output cell, receptive fields are completely localized. In the limit of low noise, the output layer replicates the activity of the input layer, so all information is preserved. However, when noise is significant, some information is lost by this identity mapping, and alternative connectivity schemes may yield greater information about the inputs. Second, when there is global connectivity, so that all input cells are connected to all output cells, receptive fields are completely delocalized. Under a correlation-based rule, each output cell learns the same receptive field. Then, in the low noise limit, most information is being thrown away — only one dimension of the input pattern is being distinguished. However, suppose that this dimension is the most informative dimension about the input ensemble. Then, in the high noise limit, this redundant representation of the most information-rich dimension will maximize the information carried about the input ensemble.

Thus, given a correlation-based learning rule, a completely localized representation can maximize information in the low-noise limit, while a completely delocalized representation can maximize information in the high-noise limit. Intermediate levels of localization should be appropriate for intermediate signal-to-noise ratios (this has recently been demonstrated quantitatively [21]). It seems likely that biology, rather than designing an anti-Hebbian learning rule, has used its own correlation-based rules and has made use of its natural tendency to form partially localized receptive fields in order to ensure efficiency of representation.

1.4.2 Self-Organizing Maps and Associative Memories

The above ideas about efficiency consider only the summed information in the responses of the postsynaptic cells, without regard for location or connectivity. Alternative ideas about the computational significance of correlation-based rules focus on the spatial arrangement of postsynaptic response features and the connectivity between the postsynaptic cells.

One set of such ideas stem from the study of the self-organizing feature map (SOFM) of Kohonen [24, 48] and of related "dimension-reducing mappings" [8]. As previously described, the SOFM corresponds to a Hebbian rule with a nonlinear lateral intracortical interaction, such that each input pattern leads to a single, localized cluster of cortical activity. The SOFM and related algorithms lead to a mapping that matches the topology and geometry of the output space to that of the input space, despite a possible dimensional and/or shape mismatch between the two [8, 24, 48]. That is, nearby points in the output space correspond via the mapping to nearby points in the input space, and input patterns that occur more often develop a larger representation than those that occur less often.

A number of possible functions have been assigned to such mappings. One is the minimization of wiring length, assuming that cortical points representing "nearby" input patterns need to be connected to one another [8]. Another is to represent the input data in a compressed form while minimizing reconstruction error [33, 48]. A specific form of the latter idea is as follows. Suppose that there is noise in the output layer that is distance-dependent, so that the probability of response being centered at a given output point falls off with its distance from the point that is "correct" for that input. Suppose also that there is a metric on the input space, and the error in mistaking one input pattern for another is assigned as the distance between the two patterns. Then the SOFM can be interpreted, approximately, as achieving the input-output mapping that minimizes the average error in reconstructing the input pattern from the output responses [33].

The major problem in applying these ideas to biology is the difficulty in assigning biological meaning to the topology and geometry of the non-retinotopic dimensions of the input space. Given an ensemble of visual input patterns on the retina, for example, how large is the corresponding ocular dominance or orientation dimension relative to the retinotopic dimensions? Without a clear prescription for answering this question, it is difficult to make biological predictions from these ideas. Nonetheless, the computational functions of self-organizing maps, their close connection to correlation-based models, and their ability to replicate many features of cortical maps are intriguing.

Another well-known set of ideas concern the role of correlation-based rules in establishing an associative memory. Suppose one wishes to learn a set of N input-output pairs, $(\mathbf{u}^a, \mathbf{v}^a)$, where $\mathbf{u}^a$ and $\mathbf{v}^a$ are the a^{th} input and output vectors respectively. Let $\mathbf{v}^a = \mathbf{M}\mathbf{u}^a$ for some synaptic matrix $\mathbf{M}$. If the input patterns are orthonormal, $\mathbf{u}^a \cdot \mathbf{u}^b = \delta_{ab}$, then the input-output association is achieved by setting $\mathbf{M} = \sum_a \mathbf{v}^a (\mathbf{u}^a)^{\mathrm{T}}$ (*e.g.* [24]). This relation will be learned by a Hebbian rule, $\frac{d}{dt} M_{ij} = -M_{ij}/N + v_i u_j$, provided there is a "teacher" to clamp the output to $\mathbf{v}^a$ whenever $\mathbf{u}^a$ is presented. A fully connected network with activity states $\mathbf{v}$ will similarly develop the activity states, or "memories", $\mathbf{v}^a$, as stable attracting states if the connection matrix between the cells is determined by the Hebbian prescription $\mathbf{M} = \sum_a \mathbf{v}^a (\mathbf{v}^a)^{\mathrm{T}}$ (*e.g.* [18, 19]). Again, to learn a specific set of memories a "teacher" is required to clamp the network into the appropriate activity states during learning. Given simple nonlinearities in neuronal activation, the stored memories need not be orthogonal to one another, provided the memories are randomly chosen (uncorrelated) and their number is sufficiently small relative to the number of cells (*e.g.* [17]). It is of biological interest to explore how associative properties can develop through correlation-based rules in the absence of a teacher, as well as in the presence of correlated input patterns (for which, see [17]).

1.5 OPEN QUESTIONS

This brief review can only point to a small sample of the rich literature on this topic. Among the many open questions in the field are: How can biologically interpretable models replicate the details of cortical maps? Might orientation selectivity arise from early oriented wave patterns of retinal activity [38, 64] or other mechanisms, rather than through ON/OFF competition? Might the initial development of orientation selectivity occur through patterning of intracortical connections, rather than through patterning of LGN connections to cortex?[7] How might intracortical plasticity affect receptive field and map development [53]? How might input correlations affect column size [12]? How will development be altered by incorporation of more realistic cortical connectivity, and more realistic, nonlinear learning rules? For example, might input correlations help determine the self-organization of plastic intracortical connections or the size of nonlinearly-determined cortical activity clusters, each of which in turn would shape the pattern of input synapses including column size? How can we characterize the computational function of the correlation-based rules used biologically? These and other questions are likely to be answered in the coming years.

1.6 Acknowledgments

K.D. Miller is supported by grants from the National Eye Institute, the Searle Scholars' Program, and the Lucille P. Markey Charitable Trust. I thank Ed Erwin, Sergei Rebrik and Todd Troyer for helpful comments on the manuscript.

[7] See [41] for arguments that the early oriented waves of retinal activity are too large to drive the development of simple cells, *i.e.* their wavelength is much wider than the set of LGN inputs to a single simple cell; but see [58] for an argument that the waves might nonetheless drive development of orientation-selectivity by determining the patterning of intracortical connections rather than of connections from LGN to cortex. The patterning of horizontal connections may take place slightly later than the development of orientation selectivity [1, 5], but both occur sufficiently early that their order remains unclear.

1.7 References

[1] K. Albus and W. Wolf. Early post-natal development of neuronal function in the kitten's visual cortex: A laminar analysis. *J. Physiol.*, 348:153–185, 1984.

[2] J.J. Atick. Could information theory provide an ecological theory of sensory processing? In W. Bialek, editor, *Princeton Lectures on Biophysics*, pages 223–289. World Scientific, Singapore, 1992.

[3] H.B. Barlow. Unsupervised learning. *Neural Comp.*, 1:295–311, 1989.

[4] B.O. Braastad and P. Heggelund. Development of spatial receptive-field organization and orientation selectivity in kitten striate cortex. *J. Neurophysiol.*, 53:1158–1178, 1985.

[5] E.M. Callaway and L.C. Katz. Emergence and refinement of clustered horizontal connections in cat striate cortex. *J. Neurosci.*, 10:1134–1153, 1990.

[6] B. Chapman and M.P. Stryker. Development of orientation selectivity in ferret visual cortex and effects of deprivation. *J. Neurosci.*, 13:5251–5262, 1993.

[7] M. Constantine-Paton, H.T. Cline, and E. Debski. Patterned activity, synaptic convergence and the NMDA receptor in developing visual pathways. *Ann. Rev. Neurosci.*, 13:129–154, 1990.

[8] R. Durbin and G. Mitchison. A dimension reduction framework for understanding cortical maps. *Nature*, 343:644–647, 1990.

[9] E. Erwin, K. Obermayer, and K. Schulten. Models of orientation and ocular dominance columns in the visual cortex: A critical comparison. *Neural Comp.*, 7:425-468, 1995.

[10] P. Foldiak. Adaptive network for optimal linear feature extraction. In *Proceedings, IEEE/INNS International Joint Conference on Neural Networks*, volume 1, pages 401–405. IEEE Press, New York, 1989.

[11] Y. Fregnac and M. Imbert. Development of neuronal selectivity in the primary visual cortex of the cat. *Physiol. Rev.*, 64:325–434, 1984.

[12] G.J. Goodhill. Topography and ocular dominance: A model exploring positive correlations. *Biol. Cybern.*, 69:109–118, 1993.

[13] G.J. Goodhill and D.J. Willshaw. Application of the elastic net algorithm to the formation of ocular dominance stripes. *Network*, 1:41–59, 1990.

[14] C.S. Goodman and C.J. Shatz. Developmental mechanisms that generate precise patterns of neuronal connectivity. *Cell*, 72 Suppl:77–98, 1993.

[15] R.W. Guillery. Binocular competition in the control of geniculate cell growth. *J. Comp. Neurol.*, 144:117–130, 1972.

[16] R.W. Guillery and D.J. Stelzner. The differential effects of unilateral lid closure upon the monocular and binocular segments of the dorsal lateral geniculate nucleus in the cat. *J. Comp. Neurol.*, 139:413–422, 1970.

[17] J.A. Hertz, A.S. Krogh, and R.G. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, 1991.

[18] J.J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA*, 79, 1982.

[19] J.J. Hopfield. Neurons with graded responses have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. USA*, 81, 1984.

[20] D.H. Hubel and T.N. Wiesel. Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *J. Neurophysiol.*, 26:994–1002, 1963.

[21] M. Idiart, B. Berk, and L.F. Abbott. Reduced representation by neural networks with restricted receptive fields. *Neural Comp.*, 7:507-517, 1995.

[22] N. Intrator. Feature extraction using an unsupervised neural network. *Neural Computation*, 4:98–107, 1992.

[23] N. Intrator and L.N. Cooper. Objective function formulation of the BCM theory of visual cortical plasticity: Statistical connections, stability conditions. *Neural Networks*, 5:3–17, 1992.

[24] T. Kohonen. *Self-Organization and Associative Memory.* Springer-Verlag, Berlin, 3rd edition, 1989.

[25] S. LeVay, T.N. Wiesel, and D.H. Hubel. The development of ocular dominance columns in normal and visually deprived monkeys. *J. Comp. Neurol.*, 191:1–51, 1980.

[26] Z. Li and J.J. Atick. Efficient stereo coding in the multiscale representation. *Network*, 5:157–174, 1994.

[27] R. Linsker. From basic network principles to neural architecture: Emergence of spatial-opponent cells. *Proc. Natl. Acad. Sci. USA*, 83:7508–7512, 1986.

[28] R. Linsker. From basic network principles to neural architecture: Emergence of orientation-selective cells. *Proc. Natl. Acad. Sci. USA*, 83:8390–8394, 1986.

[29] R. Linsker. From basic network principles to neural architecture: Emergence of orientation columns. *Proc. Natl. Acad. Sci. USA*, 83:8779–8783, 1986.

[30] R. Linsker. Self-organization in a perceptual network. *Computer*, 21:105–117, March 1988.

[31] R. Linsker. Local synaptic learning rules suffice to maximize mutual information in a linear network. *Neural Comput.*, 4:691–702, 1992.

[32] S. Löwel. Ocular dominance column development: Strabismus changes the spacing of adjacent columns in cat visual cortex. *J. Neurosci.*, 14:7451-7468, 1994.

[33] S. Luttrell. A Bayesian analysis of self-organizing maps. *Neural Comp.*, 6:767–794, 1994.

[34] D.J.C. MacKay and K.D. Miller. Analysis of Linsker's applications of Hebbian rules to linear networks. *Network*, 1:257–298, 1990.

[35] D.J.C. MacKay and K.D. Miller. Analysis of Linsker's simulations of Hebbian rules. *Neural Comput.*, 2:173–187, 1990.

[36] L. Maffei and L. Galli-Resta. Correlation in the discharges of neighboring rat retinal ganglion cells during prenatal life. *Proc. Nat. Acad. Sci. USA*, 87:2861–2864, 1990.

[37] D.N. Mastronarde. Correlated firing of retinal ganglion cells. *Trends Neurosci.*, 12:75–80, 1989.

[38] M. Meister, R.O.L. Wong, D.A. Baylor, and C.J. Shatz. Synchronous bursts of action-potentials in ganglion cells of the developing mammalian retina. *Science*, 252:939–943, 1991.

[39] K.D. Miller. Correlation-based models of neural development. In M.A. Gluck and D.E. Rumelhart, editors, *Neuroscience and Connectionist Theory*, pages 267–353. Lawrence Erlbaum Ass., Hillsdale, NJ, 1990.

[40] K.D. Miller. Derivation of linear Hebbian equations from a nonlinear Hebbian model of synaptic plasticity. *Neural Comput.*, 2:321–333, 1990.

[41] K.D. Miller. A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs. *J. Neurosci.*, 14:409–441, 1994.

[42] K.D. Miller, J.B. Keller, and M.P. Stryker. Ocular dominance column development: Analysis and simulation. *Science*, 245:605–615, 1989.

[43] K.D. Miller and D.J.C. MacKay. The role of constraints in Hebbian learning. *Neural Comput.*, 6:100–126, 1994.

[44] 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*, pages 255–350. MIT Press/Bradford, Cambridge, MA, 1990.

[45] M. Miyashita and S. Tanaka. A mathematical model for the self-organization of orientation columns in visual cortex. *NeuroReport*, 3:69–72, 1992.

[46] K. Obermayer, G.G. Blasdel, and K. Schulten. A statistical mechanical analysis of self-organization and pattern formation during the development of visual maps. *Phys. Rev. A*, 45:7568–7589, 1992.

[47] E. Oja. A simplified neuron model as a principal component analyzer. *J. Math. Biol.*, 15:267–273, 1982.

[48] H. Ritter, T. Martinetz, and K. Schulten. *Neural Computation and Self-Organizing Maps: An Introduction.* Addison-Wesley, Reading, MA, 1992.

[49] J. Rubner and K. Schulten. Development of feature detectors by self-organization. *Biol. Cybern.*, 62:193–199, 1990.

[50] T.D. Sanger. An optimality principle for unsupervised learning. In D. Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 1, pages 11–19. Morgan Kaufmann, San Mateo, CA, 1989.

[51] C.J. Shatz. The developing brain. *Scientific Am.*, 267:60–67, 1992.

[52] C.J. Shatz and M.P. Stryker. Ocular dominance in layer IV of the cat's visual cortex and the effects of monocular deprivation. *J. Physiol.*, 281:267–283, 1978.

[53] J. Sirosh and R. Mikkulainen. A unified neural network model for the self-organization of topographic receptive fields and lateral interactions. *Neural Comput.*, (to appear), 1995.

[54] M.P. Stryker and S.L. Strickland. Physiological segregation of ocular dominance columns depends on the pattern of afferent electrical activity. *Inv. Opthal. Supp.*, 25:278, 1984.

[55] N.V. Swindale. A model for the coordinated development of columnar systems in primate striate cortex. *Biol. Cyb.*, 66:217–230, 1992.

[56] S. Tanaka. Theory of ocular dominance column formation: Mathematical basis and computer simulation. *Biol. Cybern.*, 64:263–272, 1991.

[57] C. von der Malsburg. Self-organization of orientation selective cells in the striate cortex. *Kybernetik*, 14:85–100, 1973.

[58] C. von der Malsburg. Network self-organization in the ontogenesis of the mammalian visual system. Internal Report 93-06, Ruhr-Universität Bochum, Institut für Neuroinformatik, 44780 Bochum, Germany, 1993.

[59] C. von der Malsburg and D.J. Willshaw. A mechanism for producing continuous neural mappings: ocularity dominance stripes and ordered retino-tectal projections. *Exp. Brain Res.*, Supp. 1:463–469, 1976.

[60] T.N. Wiesel and D.H. Hubel. Comparison of the effects of unilateral and bilateral eye closure on cortical unit responses in kittens. *J. Neurophysiol.*, 28:1029–1040, 1965.

[61] T.N. Wiesel and D.H. Hubel. Ordered arrangement of orientation columns in monkeys lacking visual experience. *J. Comp. Neurol.*, 158:307–318, 1974.

[62] F. Wolf, H-U. Bauer, and T. Geisel. Formation of field discontinuities and islands in visual cortical maps. *Biol. Cyb.*, 70:525–531, 1994.

[63] F. Wolf, K. Pawelzik, T. Geisel, D.S. Kim, and T. Bonhoeffer. Optimal smoothness of orientation preference maps. In *Computation in Neurons and Neural Systems*, pages 97–102, Boston, 1994. Kluwer.

[64] R.O. Wong, M. Meister, and C.J. Shatz. Transient period of correlated bursting activity during development of the mammalian retina. *Neuron*, 11:923–938, 1993.