

Contour Integration and Segmentation with Self-Organized Lateral Connections

Yoonsuck Choe¹ and Risto Miikkulainen²

1. Department of Computer Science, Texas A&M University, College Station, TX 77843

2. Department of Computer Sciences, The University of Texas at Austin, Austin, TX 78712

choe@tamu.edu, risto@cs.utexas.edu

Abstract

Contour integration in low-level vision is believed to occur based on lateral interaction between neurons with similar orientation tuning. How such interactions could arise in the brain has been an open question. Our model suggests that the interactions can be *learned* through input-driven self-organization, i.e. through the same mechanism that underlies many other developmental and functional phenomena in the visual cortex. The model also shows how synchronized firing mediated by these lateral connections can represent the percept of a contour, resulting in performance similar to that of human contour integration. The model further demonstrates that contour integration performance can differ in different parts of the visual field, depending on what kinds of input distributions they receive during development. The model thus grounds an important perceptual phenomenon onto detailed neural mechanisms, so that various structural and functional properties can be measured, and predictions can be made to guide future experiments.

1 Introduction

Contour integration in low-level vision means forming a coherent percept out of a discontinuous sequence of line segments (figure 1). Contour integration is a special case of perceptual grouping; it takes place early on in the visual processing system and lends itself to precise psychophysical measurements. Thus, understanding the neural mechanisms underlying contour integration can give us insights into how perceptual grouping in general can be implemented.

Psychophysical experiments (Field et al. 1993; Geisler et al. 1999, 2001; Pettet et al. 1998), neurophysiological observations (Gilbert and Wiesel 1990; Bosking et al. 1997; Fitzpatrick et al. 1994), and computational models (Geisler et al. 1999, 2001; Grossberg and Williamson 2001; Li 1998; Ross et al. 2000; VanRullen et al. 2001; Yen and Finkel 1997, 1998) suggest that contour integration in the visual cortex may be due to lateral interaction of neurons with similar orientation tuning. In the preceding models, such interactions are hard-coded based on specific association constraints (Li 1998; Ross et al. 2000; VanRullen et al. 2001), a predetermined set of rules such as relative orientation difference, distance, curvature, and change in curvature (Yen and Finkel 1997, 1998). The models match experimental data quite well. Neurophysiological observations have shown that

specific patterned lateral connections that could implement such rules exist in the visual cortex (Bosking et al. 1997; Fitzpatrick et al. 1994). Grossberg and Williamson (2001) showed computationally how such lateral connections could adapt to achieve stable contour integration with fixed receptive fields.

However, it is currently unknown how such detailed afferent and lateral connection patterns could emerge during development. Furthermore, since the models are based on rules uniformly applied over the whole model cortex, they do not explain differences in contour integration performance across different areas of the visual field. For example, contour integration has been found to be absent in human peripheral vision (Hess and Dakin 1997), and convexity of illusory contours are harder to discriminate in the upper hemifield compared to the lower hemifield (Rubin et al. 1996). The goal of this paper is to show that all these phenomena can be due to input-driven self-organization of the visual cortex. This way, contour integration can be seen as a necessary effect of the same developmental process that is responsible for the organization, plasticity, and several functional phenomena of the visual cortex.

Several models of self-organization have been proposed to explain how the orientation maps in the cortex could form (Bartsch and van Hemmen 2001; Burger and Lang 1999; Goodhill and Cimponeriu 2000; Kohonen 1981, 1982, 1995; Miller 1994; Obermayer et al. 1990; von der Malsburg 1973; see Swindale 1996 for a review). In most of these models, only the afferent connections self-organize while the lateral interactions are represented as a fixed, uniform interaction kernel. In those models where the lateral connections adapt as well, the final connectivity pattern is elongated, but not patchy like the patterns in the visual cortex. Therefore, such models cannot account for functional phenomena that depend on the specific patterns of lateral connections. With this goal in mind, we recently developed a model with *explicit self-organizing lateral connections*, showing that patches of strong lateral connections develop between neurons with similar orientation preference, and that these connections can serve as a foundation for segmentation and binding (RF-SLISSOM, or Receptive Field Spiking Laterally Interconnected Synergetically Self-Organizing




Figure 1: **Contour Integration Task.** This figure shows a typical input image used in psychophysical experiments on contour integration. Human subjects are instructed to find the longest continuous contour consisting of separate line segments embedded in a background of randomly oriented distractors. In this example, the contour consists of six segments, running diagonally from middle-left to bottom-right.

Map; Bednar and Miikkulainen 2000b; Choe and Miikkulainen 1998; Miikkulainen et al. 1997; Sirosh 1995; Sirosh et al. 1996; Sirosh and Miikkulainen 1997). Self-organization of laterally connected maps is the first main principle of the contour integration model presented in this paper.

Patterned lateral interactions are strongly believed to contribute to contour integration, but how does the visual system represent a contour as a coherent object with its neural activity? A separate line of research has produced a possible answer to this question. Experiments have shown that feature binding and segmentation in the visual system may be based on temporal coding produced by synchronous and desynchronous population activity (Eckhorn et al. 1988; Engel et al. 1991; Gray and Singer 1987; Gray et al. 1989; see Singer and Gray 1995, Gray 1999, and Singer 1999 for a review). Locally synchronous firing has been observed for example in the visual cortex of cats and monkeys. Recordings of single-unit activities, multi-unit activities (MUA), and local field potentials (LFP) in different areas of the visual cortex were taken, and neurons with non-overlapping receptive fields were found most likely to be synchronized when the receptive field properties were similar, or when the firing represented global stimulus properties. Computational models also demonstrated that such a behavior can be obtained in a network of neurons with temporal dynamics (Eckhorn et al. 1990; Gerstner 1998; Grossberg and Grunewald 1997; Horn and Opher 1998; Reitboeck et al. 1993; von der Malsburg 1986, 1987; von der Malsburg and Buhmann 1992; Wang 1995, 1996, 2000). Therefore, segmentation and binding by synchronized firing is the second main principle of our contour integration model.

In this paper, the above two principles are unified into a single model. We demonstrate that the orientation map and the lateral connections self-organize so that the functional statistics of lateral connections become similar to edge co-occurrence statistics in natural images. These connections mediate synchronized firing of neurons, so that the contour

integration performance of the model closely matches psychological data. This way, the model (1) shows how the circuitry for contour integration could arise from general self-organization mechanisms in the brain, (2) provides further computational evidence for synchronization as the substrate for segmentation and binding, (3) grounds an important perceptual grouping phenomenon on a detailed neural architecture, where various functional properties can be measured, and predictions made to guide future experiments.

2 Model Description

2.1 Motivation and Overview

The contour integration model is based on the RF-SLISSOM model of self-organization and segmentation in the primary visual cortex (Choe and Miikkulainen 1998). In this model, each cortical neuron receives afferent connections from the input layer and lateral excitatory and inhibitory connections from neighboring neurons in the cortex. The connection strengths self-organize based on correlations in the activity. In the final ordered map, the lateral excitation has a short range, and causes neurons responding to the same connected input object to fire synchronously, effectively binding the spikes into a single coherent representation. The lateral inhibitory connections have a long range, and establish competition between representations of different objects.¹ Neurons representing different objects fire at different times, and the input is thereby segmented into different objects.

This previous model showed how self-organization and segmentation can be achieved in a single unified framework. The lateral interactions play a crucial role in both behaviors: they establish competition that drives self-organization, and they establish desynchronization that drives segmentation. The model did not include any long-range excitatory connections because they were not found necessary to model the above behaviors. However, it turns out that such a parsimonious model cannot account for filling-in phenomena such as contour integration. The network has to be able to bind together representations that are separated by gaps: that is, it has to have long-range excitatory connections that link together the representations of the different segments of a fragmented contour.

The model is extended in this paper with such long-range excitatory connections (figure 2). The extended model is called PGLISSOM (or Perceptual Grouping LISSOM; Choe 2001). The cortical network is divided into two separate components: MAP1 and MAP2. MAP1 is similar to the RF-SLISSOM cortex with short-range excitatory and long-range inhibitory connections. This map has the task of driving the

¹ Although long-range connections in the cortex appear to be mostly excitatory, their effect can be inhibitory through inhibitory interneurons (Grinvald et al. 1994; Hata et al. 1988; Hirsch and Gilbert 1991; Weliky et al. 1995). RF-SLISSOM abstracts such interneurons and models the overall inhibitory effect as one connection.

self-organization of the network into an ordered map. MAP2 performs the task of long-range segmentation and binding, with long-range excitatory connections that perform contour integration, and long-range inhibitory connections that implement segmentation of separate objects. The two maps are assumed to be overlaid in one cortical network: such a functional specialization across laminar layers of the visual cortex is consistent with known neuroanatomical data from layers IV, VI, and II/III of the visual cortex (Grossberg and Williamson 2001). In other words, the model is based on the hypothesis that some of the neurons in each cortical column are involved in establishing and maintaining organization, whereas others perform visual segmentation and binding.

2.2 Neuron Model

The details of the neuron model are illustrated in figure 2a. Each connection is a leaky integrator that performs exponentially decayed summation of incoming spikes (i.e. convolution with an exponential kernel; Eckhorn et al. 1990):

$$s(t) = \sum_{n=0}^t x(t-n)e^{-\lambda n}, \quad (1)$$

where $s(t)$ is the current decayed sum at time step t , $x(t-n)$ is the input spike (either 0 or 1) n time steps in the past, and λ is the decay rate. Different types of connections have distinct decay rates: λ_e for excitatory and λ_i for inhibitory lateral connections, and λ_c for intra-columnar connections. The sum can be defined in a computationally more practical form as a recurrence equation, which is used in the current implementation:

$$s(t) = x(t) + s(t-1)e^{-\lambda}, \quad (2)$$

where $s(t)$ and $s(t-1)$ are the current and previous decayed sums, $x(t)$ is the current input spike and λ is the decay rate (Eckhorn et al. 1990). The leaky integrator models the Post-Synaptic Potential (PSP) that decays exponentially over time in biological neurons. By adjusting the decay rate λ , the synapse can function as either a coincidence detector or as a temporal integrator. When the synaptic decay rate is high, the neuron can only fire when there is a sufficient number of inputs coming in from many synapses simultaneously. On the other hand, when the decay rate is low, the neuron accumulates the input. Thus pre-synaptic neurons can have a lingering influence on the post-synaptic neuron. By varying the decay rates for different types of connections, the relative time scales of the different connection types can be controlled to obtain desirable synchronization behavior.

The spike generator compares the input to a threshold and decides whether to fire a spike. The threshold is dynamic, depending on the previous firing activity at the neuron, in order to model the refractory period and to improve synchronization. It consists of three terms:

$$\theta(t) = \theta_{\text{base}} + \theta_{\text{abs}}(t) + \tau\theta_{\text{rel}}(t), \quad (3)$$

where θ_{base} is the base threshold, $\theta_{\text{abs}}(t)$ implements the absolute refractory period during which the neuron cannot fire, $\theta_{\text{rel}}(t)$ implements the relative refractory period during which firing is possible but requires extensive input, and τ is a scaling constant. The absolute refractory component $\theta_{\text{abs}}(t)$ is defined as:

$$\theta_{\text{abs}}(t) = \begin{cases} \infty & \text{if } y(t-i) = 1 \text{ for any } i \leq \kappa_{\text{abs}} \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

where κ_{abs} determines the length of the absolute refractory period, and $y(t)$ represents whether a spike occurred at time t :

$$y(t) = \mathcal{H}(\sigma(t) - \theta(t-1)), \quad (5)$$

where $\mathcal{H}(\cdot)$ is the Heaviside step function, $\sigma(t)$ is the weighted input sum (equation 7), and $\theta(t-1)$ is the dynamic threshold. The relative refractory component $\theta_{\text{rel}}(t)$ is implemented as an exponentially decayed sum of the output spikes (figure 2a), i.e. a leaky integrator similar to the leaky synapses:

$$\theta_{\text{rel}}(t) = y(t) + \theta_{\text{rel}}(t-1)e^{-\lambda_{\text{rel}}}, \quad (6)$$

where λ_{rel} is the decay rate.

Eckhorn et al. (1990) and Reitboeck et al. (1993) described a similar dynamic threshold mechanism, consisting of θ_{base} and θ_{rel} only. The absolute refractory period was included to our model to ensure that the neurons do not fire too rapidly. An added benefit is that synchronization is more robust against noise (Choe 2001).

2.3 Network Activation and Learning

The organization of the network is shown in figure 2b. The input $\sigma_{i,j}(t)$ to the spike generator of the cortical neuron (in each map) at location (i, j) at time t consists of (1) the input from a fixed-size receptive field in the retina², centered at the location corresponding to the neuron's location in the cortical network, (2) from neurons around the same relative location in the opposite map, and (3) excitation and (4) inhibition from neighboring neurons in the same map:

$$\begin{aligned} \sigma_{i,j}(t) = & g(\gamma_a \sum_{r_1, r_2} \xi_{r_1, r_2} \mu_{ij, r_1 r_2} \\ & + \gamma_c \sum_{p_1, p_2} \zeta(t-1)_{p_1, p_2} \nu_{ij, p_1 p_2} \\ & + \gamma_e \sum_{k, l} \eta_{kl}(t-1) E_{ij, kl} \\ & - \gamma_i \sum_{k, l} \eta_{kl}(t-1) I_{ij, kl}), \end{aligned} \quad (7)$$

where $\gamma_a, \gamma_c, \gamma_e$, and γ_i are the relative strengths of the afferent, intra-columnar, and excitatory and inhibitory lateral contributions, ξ_{r_1, r_2} is the input level of retinal neuron (r_1, r_2) , $\mu_{ij, r_1 r_2}$ is the corresponding afferent connection

² The preprocessing in the retinal ganglion cells and lateral geniculate nucleus (LGN) was bypassed for simplicity.

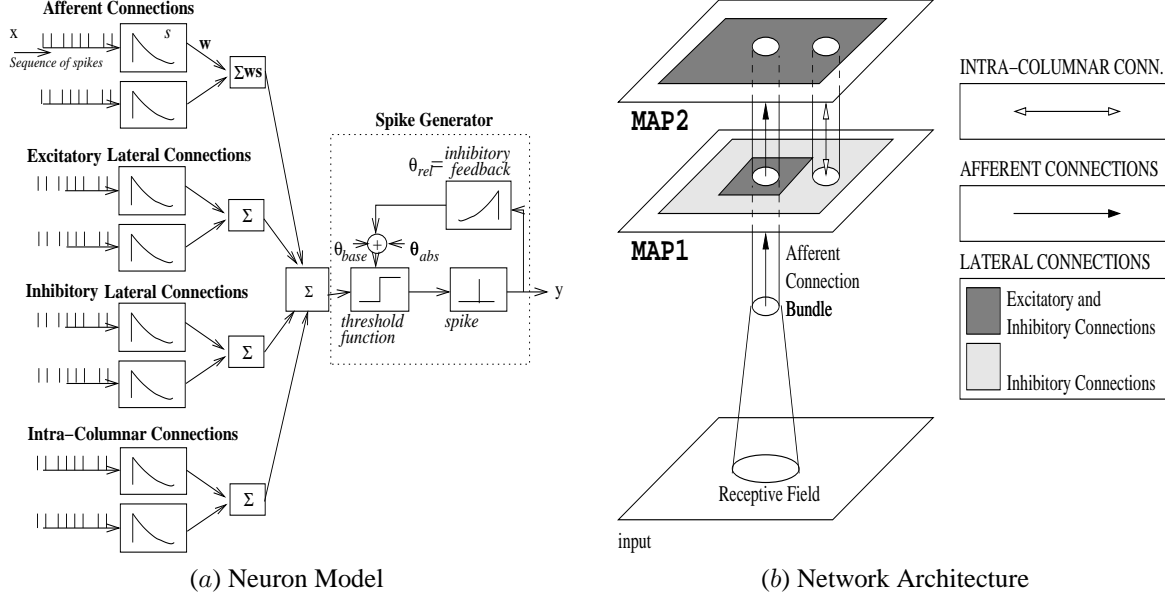


Figure 2: Overview of the PGLISSOM Model. (a) The neuron model. Leaky integrators at each synapse perform decayed summation of incoming spikes, and the outgoing spikes are generated by comparing the sum of weighted sums to the dynamic spiking threshold. Four types of inputs contribute to the activity: afferent, excitatory lateral, inhibitory lateral, and intra-columnar connections. The dynamic threshold consists of the base threshold θ_{base} , the absolute refractory component θ_{abs} , and the relative refractory component θ_{rel} . (b) The overall organization of the network. The cortical network consists of two layers (or maps): MAP1 has short-range excitation and long-range lateral inhibition, and drives the self-organization of the model. In MAP2, both excitation and inhibition have a long range, establishing binding and segmentation. Both maps receive input from a model retina, and neurons in the vertically corresponding locations on the two maps are connected via intra-columnar connections representing a cortical column.

weight, ζ_{p_1, p_2} is the decayed sum of spikes of the cortical neuron (p_1, p_2) of the other cortical map, $\nu_{ij, p_1 p_2}$ is the corresponding intra-columnar connection weight, $\eta_{kl}(t-1)$ is the decayed sum of spikes from the map neuron (k, l) at time $t-1$, $E_{ij, kl}$ is the corresponding excitatory and $I_{ij, kl}$ the inhibitory lateral connection weight, and $g(\cdot)$ is a piece-wise linear approximation of the sigmoid function that squashes the net input sum between 0.0 and 1.0:

$$g(x) = \begin{cases} 0 & \text{if } x < \delta \\ 1 & \text{if } x > \beta \\ \frac{x-\delta}{\beta-\delta} & \text{otherwise,} \end{cases} \quad (8)$$

where δ is the threshold and β is the ceiling.

The inputs to the model consist of activation patterns with activation values ranging between 0 and 1. A fixed such input is presented on the retina at each iteration and the cortical neurons are allowed to generate and exchange spikes. After several iterations, the short-term spiking rate of the neurons in a small window is calculated:

$$V(t) = \tau_{avg} V(t-1) + (1 - \tau_{avg}) y(t), \quad (9)$$

where τ_{avg} is the window size, $V(t-1)$ is the previous average firing rate, and $y(t)$ is the output spike at time t . The afferent, lateral and intra-columnar weights are then modified

according to the normalized Hebbian learning rule:

$$w_{ij, mn}(t) = \frac{w_{ij, mn}(t-1) + \alpha V_{ij}(t) V_{mn}(t)}{\sum_{ij} [w_{ij, mn}(t-1) + \alpha V_{ij}(t) V_{mn}(t)]}, \quad (10)$$

where $w_{ij, mn}(t)$ is the connection weight from neuron (m, n) to (i, j) , $w_{ij, mn}(t-1)$ is the previous weight, α is the learning rate (α_a for afferent, α_c for intra-columnar, α_e for excitatory, and α_i for inhibitory connections), and $V_{ij}(t)$ and $V_{mn}(t)$ are the average spiking rates of the neurons.

This process of weight adaptation is repeated with inputs at random locations and orientations, and the neurons gradually become sensitive to particular orientations at particular locations, resulting in a global retinotopic orientation map similar to that found in the visual cortex. The self-organized map will then synchronize and desynchronize the firing of neurons to indicate binding and segmentation of visual features to different objects present in the scene. The lateral connections that survive connection death play an important role in this process, by mediating synchronization and desynchronization among populations of neurons.

3 Experiments

A Stacked RF-SLISSOM network with a 46×46 retina and a 136×136 cortex was trained for 40,000 iterations with straight elongated Gaussian bars at random locations in the retina. Although natural images could in principle be used




Figure 3: **Orientation Preferences in MAP1 and MAP2** . The orientation preference at each location on the cortex is coded in color, according to the color key on the right. The orientation preference of each neuron was calculated by taking a dot product of its afferent weight matrix and six different elongated Gaussians: the preference was the vector sum of six polar vectors each consisting of the angle of one Gaussian and its dot product (Bednar 1997; Blasdel 1992). The same organization of orientation preferences developed in both maps. The global and local features such as pinwheel centers and fractures in each map closely match those found in the visual cortex.

as well (Bednar 2002), such abstract input is computationally more efficient while still representing the essential local features of natural stimuli after the edge detection and enhancement mechanisms in the retina and LGN. During each training presentation, the network was allowed to settle for 15 time steps (through equation 7) and all connections except the inhibitory lateral connections in MAP2 were updated according to equation 10. The fixed inhibition in MAP2 provides a baseline similar to global inhibition in other cortical models (Eckhorn et al. 1988; Kammen et al. 1989; Terman and Wang 1995; von der Malsburg and Buhmann 1992; Wang 1995, 1996, 2000): it allows input elements to be segmented by default, unless lateral excitation binds them together. The details of the model and the simulation details are given in the Appendix.

3.1 Orientation Map and Functional Connection Statistics

A well-formed orientation map emerged in the training process (figure 3). The map is qualitatively similar to the orientation map in the primary visual cortex with features such as linear zones where orientation preference changes continuously along one direction, pinwheel centers around which a full 360 degrees of orientation preferences can be observed, and fractures where orientation preference changes abruptly (Blasdel 1992; Blasdel and Salama 1986). Because of the intra-columnar connections, similar activity patterns formed on both maps during self-organization, and they developed almost identical global organizations (figure 3). After training, lateral connections with weights less than 0.001 were killed, leaving a patchy connection profile (figure 5a-c).

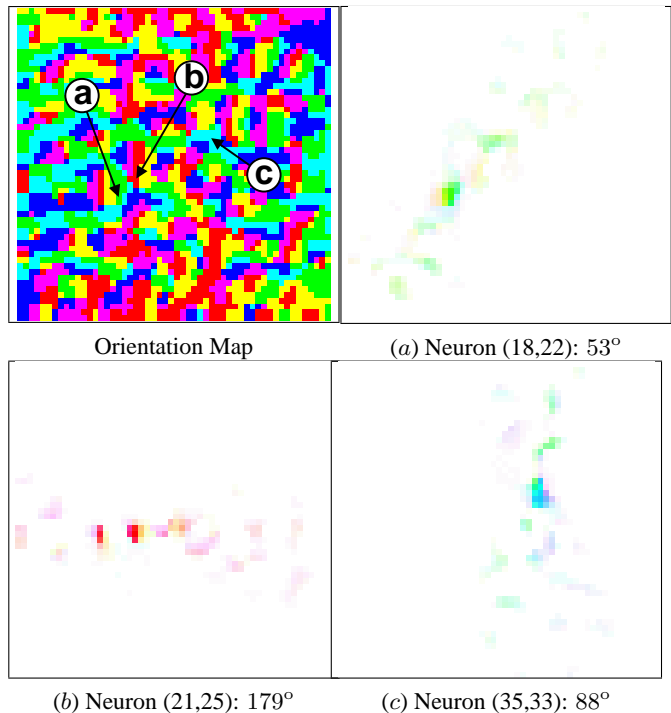


Figure 4: **Excitatory Lateral Connections** . The excitatory lateral connections from three source neurons in MAP2 (marked by arrows in the orientation map) are shown. The hue represents the orientation preference of the target neuron, and the intensity represents the strength of the connection. The neurons are numbered in Cartesian coordinates, where the lower left corner is neuron (1,1) and the upper right corner is neuron (54,54). (a) The excitatory lateral connections of neuron (18,22), with an orientation preference of 53 degrees. (b) The excitatory lateral connections of neuron (21,25), with an orientation preference of 179 degrees. (c) The excitatory lateral connections of neuron (35,33), with an orientation preference of 88 degrees. The lateral connections link neurons with similar orientation tuning (similar hue), and the target zones are aligned along the orientation preference of the source neuron, as is the case in experimental observations (Bosking et al. 1997). Specific connections like these are crucial for implementing perceptual grouping tasks such as contour integration.

Like connectivity patterns found in the visual cortex (Bosking et al. 1997; Fitzpatrick et al. 1994), the remaining lateral connections target those neurons that have a similar orientation preference as their source neuron, and they are distributed mainly along the direction of the source neuron's preferred orientation. In other words, connections link areas with highly-correlated activity, such as those along a continuous contour.

To quantify the grouping rules implemented by the lateral excitatory connections, their distributions in final MAP2 were measured in detail (figure 5). Since these distributions are obtained from the receptive fields of the neurons, they describe the *functional connectivity* of the neurons in the *retinal* (i.e. *visual*) space rather than simple cortical wiring statistics. The results confirm that (1) the lateral connections more

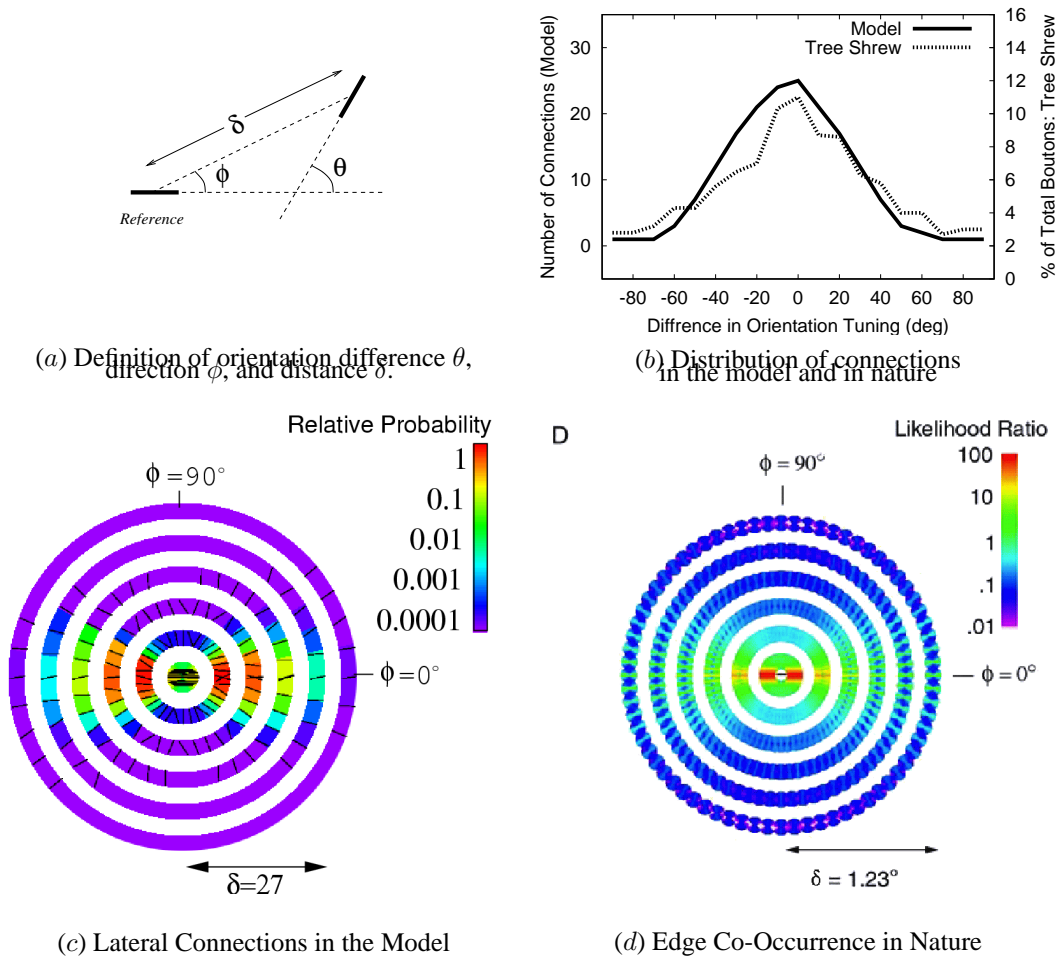


Figure 5: Lateral Connection and Edge Co-Occurrence Statistics. The distributions of lateral connections in the model are compared to the edge co-occurrence statistics in nature to see how well they match perceptual requirements. (a) The plot summarizes the notation used in figure b–d. For each pair of neurons connected with lateral excitatory connections, the afferent weight matrix was used to calculate (1) the orientation preference of the neuron (shown as oriented bars) and (2) the location of the receptive field in retinal space (as the center of gravity of the afferent weight matrix). From these values, the direction ϕ , radial distance δ , and orientation preference difference θ between all pairs of neurons shown as the color of the bar were calculated. Notice that these values measure the spatial relationship between the two neurons in the retinal (or visual) space, not in the cortical space, and therefore allow comparing connectivity with human performance data (Geisler et al. 2001). (b) The number of connections as a function of orientation difference in the model (solid line) and experimental data (dotted line) are shown. The number of excitatory lateral connections in MAP2 that exceeded a threshold value of 0.001 were counted, and normalized by the number of neurons. The corresponding measurements in tree shrew visual cortex were obtained by staining methods (Bosking et al. 1997). This plot shows that strong excitatory lateral connections mostly link neurons with similar orientation tuning. (c) The distributions of θ , ϕ , and δ in MAP2. Each location in polar coordinates (ϕ, δ) displays two values: (1) the black oriented bars represent the most probable orientation θ of the target receptive field at that direction and distance (ϕ, δ) . These orientations are aligned along co-circular paths emanating from the center. (2) The color scale in the background shows the relative log-probability of finding a target receptive field at direction ϕ and distance δ . The locations with high relative probability (the red and yellow areas) form a bow-tie shaped flank along the horizontal axis. Such an arrangement is very similar to the association field (or local grouping functions) suggested by psychophysical research (Field et al. 1993; Geisler et al. 1999, 2001). (d) The Bayesian edge co-occurrence statistic, (Geisler et al. 2001, with permission from the author). Each location (ϕ, δ) contains a small round disk, representing the likelihood ratios of all possible orientations θ at that location by color coding. The θ with the highest ratio is shown in the foreground. Each likelihood ratio represents the conditional probability of a pair of edge elements in configuration (θ, ϕ, δ) belonging to a *same physical contour* vs. *different physical contours* in natural images. A high likelihood ratio indicates that a pair of edges in that configuration is more likely to occur on a *common natural contour* than on *separate contours*. The conditional probabilities were determined through manual labeling of contours in real world images. There is a strong correspondence between this data and the connection statistics in the model, suggesting that the model is well-suited for encoding grouping relations in natural images.

often connect neurons with similar orientation tuning (figure 5b), and (2) connections go to target neurons with receptive fields aligned along the preferred orientation of the source neuron, with a small flank (figure 5c). In other words, neurons whose receptive fields fall on a smooth (co-circular) contour are most likely to be connected with strong lateral excitatory connections in MAP2.

Interestingly, these connection statistics are very similar to the edge co-occurrence statistics in natural images (Geisler et al. 2001; figure 5d). Combined with transitive grouping rules, such edge co-occurrence statistics can accurately predict human contour integration performance (Geisler et al. 1999, 2001). Therefore, we expect the model to perform like humans as well. If this prediction is confirmed, it lends computational support to the idea that self-organized lateral connectivity in V1 underlies contour integration performance in humans.

3.2 Contour Integration

Psychological experiments by Field et al. (1993) and Geisler et al. (1999, 2001) have shown that contour integration accuracy is maximal when orientation jitter in the physical contour is 0° , and the accuracy decreases as a function of increasing orientation jitter. The lateral connection statistics in the previous section are consistent with such behavior, but does the model actually perform that way? To answer this question, we ran several contour integration experiments with varying degrees of orientation jitter (figure 6).

To measure the performance of the model, for each input bar, the number of spikes generated by the area of the cortex that responded to the bar was counted at each time step. This quantity is called the Multi-Unit Activity of the response, or MUA, and it can be used to identify which area of the cortex is active at each time step. In order to determine the degree of synchronization between two areas, the linear correlation coefficient r between their MUA sequences was calculated as follows:

$$r = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2} \sqrt{\sum_i (y_i - \bar{y})^2}}, \quad (11)$$

where x_i and y_i , $i = 1, \dots, N$ are the MUA values at time i for the two areas representing the two different objects in the scene, and \bar{x} and \bar{y} are the mean of each sequence.

Using r as the measure, the contour integration performance of the network in the four different input configurations (degrees of orientation jitter; figure 6) was calculated. The network was presented with each input for 600 iterations and the MUAs from the activity areas on MAP2 corresponding to the nine input segments were obtained. The MUA sequences are shown in figure 6. The correlation coefficients for MUA pairs belonging to the same contour measure the degree of contour integration. The higher these values, the more synchronized are the areas, thus representing a strong percept

of a salient contour. The average of the within-contour correlations is used as a measure of overall performance of the model. The results are summarized in figure 7, plotted against the human performance data from Geisler et al. (1999, 2001). The plot clearly shows that at low orientation jitter, the model and human performance are both high, but as the jitter increases, they both deteriorate in a similar manner. Correlation coefficients between MUA pairs corresponding to two background segments, or pairs between a background and a segment in the contour remained low, usually near 0 (not shown), thus they were not perceptually salient. Such a performance profile is closely predicted by the lateral connection statistics, as described in the previous section. This way, the perceptual phenomenon of contour integration can be grounded on the circuit-level description given by the model.

3.3 Contour Segmentation

Importantly, the synchronization process that establishes the contour percept can also separate different contours to different percepts. In this experiment, two collinear contours were presented as input and the correlations between and across the MUAs representing each input segment were calculated (figure 8). All simulation parameters were the same as in the previous experiment. By comparing the rows in the plot, we can see that in the beginning (at stimulus onset) all areas are synchronized, but as the lateral interactions start to take effect, the MUAs start to form two major groups firing in two alternating phases. The correlation coefficients of areas in the same contour are consistently high (0.86) while those in different contours are very low (-0.11), signifying integration within each contour and segmentation across the two contours. This result suggests that the same circuitry responsible for contour integration can also be responsible for contour segmentation when there are multiple salient contours.

3.4 Influence of Input Distribution on Structure and Performance

The results in previous sections suggest that lateral connections play a significant role in the contour integration and segmentation process. Because these connections are formed in an input-driven self-organizing process, different input distributions result in different patterns of lateral connectivity, which in turn result in different performance in behavioral tasks. This explanation can potentially account for the observations by Hess and Dakin (1997) and Rubin et al. (1996) showing that contour integration performance differs between fovea and periphery, and between upper and lower visual hemifield.

To date, the distributions of input features across the visual field have not been fully characterized, and it is not possible to verify this hypothesis directly with experimental data. However, with the current model we can test the basic principle that different input distributions result in different performance. In this section, we will systematically vary the input

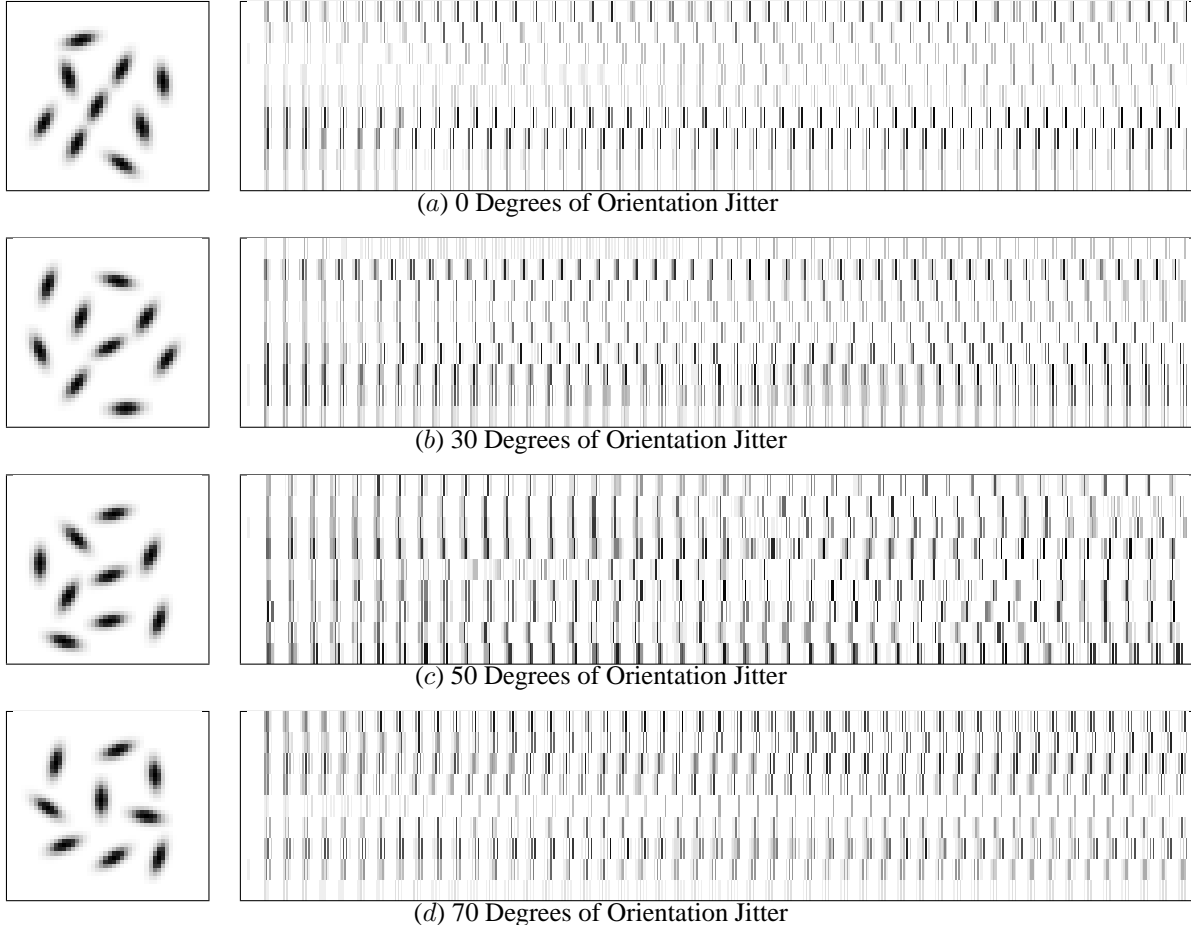


Figure 6: Contour Integration with Varying Degrees of Orientation Jitter. (a–d) The input presented to the network (left) and the resulting multi-unit activity plot (right) are shown. Each contour element is an oriented Gaussian of length $a^2 = 3.5$ and width $b^2 = 1.5$ (appendix A.2). The activity levels of the retinal units are plotted in gray scale from white to black. Each contour was composed of three contour elements, and embedded in a background of six other randomly-oriented elements. Each contour runs diagonally from lower left to top right with varying degrees of orientation jitter. At right, the Multi-Unit Activities (MUAs) of the active areas are shown in gray scale. Time (i.e. simulation iteration) is on the x -axis and the y -axis consists of nine rows, each plotting the MUAs corresponding to one input. The bottom three rows represent the MUAs of the salient contour, and the top six rows represent MUAs of the random background contour elements. Synchronization is very strong for 0 and 30 degrees but relatively weak for 50 and 70 degrees, that is, the contours get harder to detect as the orientation jitter increases. In all cases (a to d), the background MUAs are unsynchronized. A quantitative summary of these results are shown in figure 7.

distribution and show how it affects the lateral connectivity and contour integration in the PGLISSOM model.

There are several ways in which the input distribution could vary between two areas: one area could receive more training with oriented inputs than the other, or it could receive longer or more sharply defined edges, or edges with higher curvature or preferred orientation, or edges organized into a texture. We chose to test PGLISSOM in two representative ways, by (1) changing the input presentation frequency, and by (2) changing the curvature of the input. In the first experiment, inputs were presented in one area twice as often as in the other. In the second, the angles between the line segments constituting an input was varied uniformly randomly

within the interval $[0^\circ, 10^\circ]$ in one area vs. $[0^\circ, 25^\circ]$ in the other. Under each of these conditions, a PGLISSOM network was trained in the same way as in section 3.1, and tested as in section 3.2 using the same set of parameters. These conditions represent the general idea that the inputs e.g. in the fovea and in the lower hemifield are likely to be more numerous and more complex than in the upper hemifield and in the periphery.

After training each network for 40,000 iterations, orientation maps comparable to those in the previous self-organization experiments emerged in each case. However, their lateral connection patterns were quite different, as shown by the (θ, ϕ, δ) statistics similar to those in section 3.1

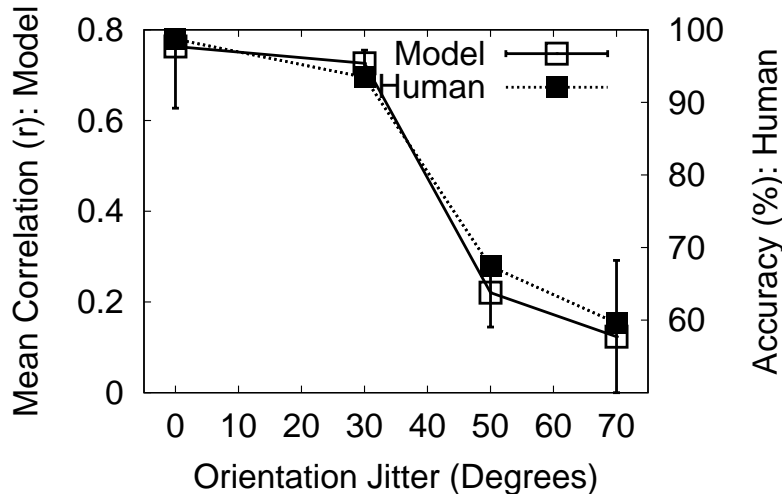


Figure 7: **Contour Integration in Humans vs. the Model.** The model’s performance was measured as the average correlation coefficient between the MUA sequences in the salient contour, and that of humans as the percentage of correctly identified contours (Geisler et al. 1999, 2001; RMS amplitude 12.5, fractal exponent 1.5; the error bars indicate standard deviation in the model). Error bars are \pm standard deviation. The x -axis is the orientation jitter in degrees (θ). In both cases, performance is robust up to 30 degrees, but then quickly breaks down as the orientation jitter increases.




Figure 8: **Contour Segmentation.** The input for the contour segmentation experiment consisted of two contours, vertical and diagonal (shown at left). The gray scales and the input sizes are identical to those in figure 6. The MUA sequences from the six areas of MAP2 responding to each input are shown at right. The bottom three rows correspond to the vertical contour and the top three rows to the diagonal contour. The average correlation coefficients between pairs of MUA sequences is high (0.86) within the same contour and low (-0.11) across different contours. Neurons within a contour are synchronized, while neurons belonging to different contours are desynchronized, segmenting two contours.

(figures 9 and 10).

In the frequency experiment, two major differences emerged (figure 9): (1) The high probability areas extend out longer in the high-frequency map (*a*) than in the low-frequency map (*b*), i.e. the map with more exposure to oriented edges can group together more distant inputs. (2) The most probable θ for a given (ϕ, δ) location tends to be co-circular in the high-frequency map (*a*), while in the low-frequency map (*b*) it is more collinear (i.e. the black edges in the high probability areas are more parallel). Collinearity is the most prominent feature in the input, and is therefore learned first. With enough input presentations, it is extended to large distances. Co-circularity develops slower than collinearity because the response levels are lower in the co-circular arrangement. The more frequently stimulated map had enough input presentations and was able to learn the sec-

ondary (co-circularity) property as well.

In the curvature experiment (figure 10), high probability areas are broader in the map trained with a broader range of curvatures (*a*) compared to the one trained with a narrower range (*b*). As expected, the input-driven self-organizing process has encoded the input distribution differences into the lateral connections. As a result, the map with exposure to higher curvature should be better at integrating co-circular contours.

The difference in structure predicts that contour integration performance between the networks should also differ. To test this prediction, contour integration experiments like those in section 3.2 were performed on each network: each network was activated for 600 iterations and the MUA sequences corresponding to the three contour elements were measured. The results are summarized in figure 11.




Figure 9: **Distributions of Excitatory Lateral Connections with High and Low Frequency of Input Presentation.** During training, PGLISSOM network (a) received oriented training inputs twice as often as network (b). As a result, the lateral connection profiles in MAP2 differ in two significant ways: (1) the high probability areas (red and yellow) extend longer in (a) than in (b), and (2) the most probable θ (black oriented bars) are co-circular in (a), but mostly collinear in (b). These results predict that contours should be easier to detect in visual areas that see oriented inputs more often.




Figure 10: **Distributions of Excitatory Lateral Connections Trained with Different Curvature Ranges.** Two PGLISSOM networks were trained with different input curvature distributions: (a) with $[0^\circ, 25^\circ]$ and (b) with $[0^\circ, 10^\circ]$. The MAP2s in these networks developed different lateral connection distributions as a result. (The number of tracks in the probability plots were increased, by reducing the histogram bin size, to highlight this difference.) The high probability areas (red and yellow) are broader in (a) than in (b), suggesting that contours with more curvature and higher orientation jitter should be easier to detect in network (a).

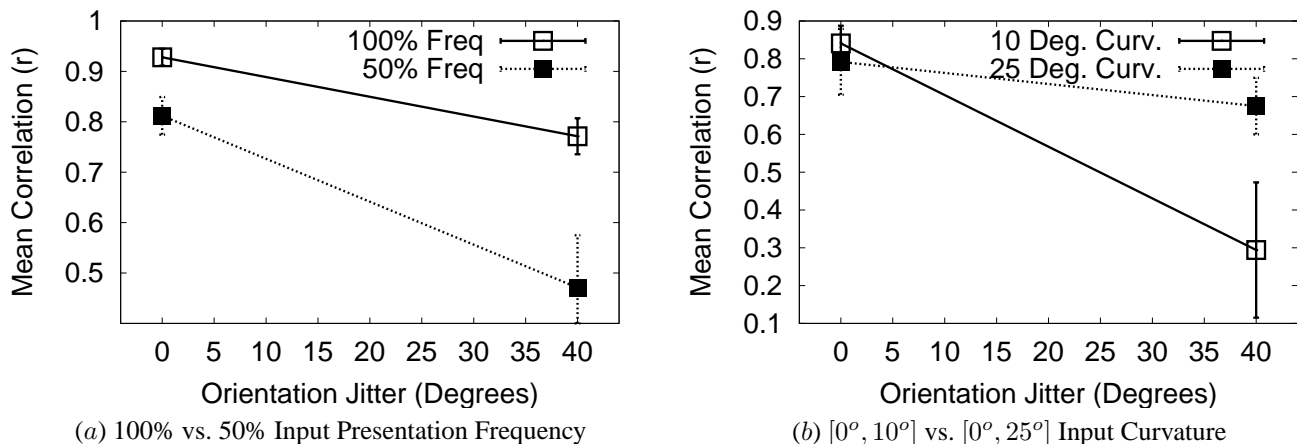


Figure 11: **Contour Integration Performance for Different Input Distributions.** For each input consisting of 3 contour elements, the correlation coefficients were calculated between each pair of MUA sequences, and the average was used as the measure of performance, as in figure 7. Error bars indicate \pm standard deviation. Plot (a) shows the contour integration performance for the frequency experiments. For both 0-degree and 40-degree orientation jitter, the high-frequency map had higher correlation than the low-frequency map. The difference is more pronounced in the 40-degree case, as predicted by the lateral connection distributions in figure 9. Plot (b) shows the contour integration performance for the $[0^\circ, 10^\circ]$ vs. $[0^\circ, 25^\circ]$ curvature experiments. At 0-degree orientation jitter, the performance is comparable, but with 40 degrees of jitter, the map trained with higher curvature input performs significantly better, as predicted by the connection distributions in figure 10. (The performance is generally lower in (b) than in (a) because to cover the whole range of curvatures, each particular curvature can only be shown with low frequency.)

The MUAs were more synchronized in the high-frequency network than in the low-frequency one for both 0 and 40 degrees of orientation jitter (figure 11a). Moreover, this performance gap was wider in the 40-degree case, indicating that the more frequently stimulated map did not just learn to same task more accurately—it actually learned to detect co-circular contours in addition to the collinear ones.

In the curvature comparison experiment, the MUAs of both conditions were equally synchronized in the 0-degree orientation jitter case. However, with 40 degrees of jitter, the map trained with higher curvature ($[0^\circ, 25^\circ]$) synchronized the components of the contour significantly better (figure 11b). The more co-circular lateral connections allowed this map to synchronize line segments that were less perfectly aligned.

These results show that if the input distribution varies across different areas of the visual field, the input-driven self-organization process will shape the connections accordingly, and such structural differences will lead to different performance in contour integration. This is an important prediction of the model that in the future can be tested with input variation in natural visual input. Such studies can eventually lead to a computational explanation of why visual performance differs across the visual field, and perhaps to some extent even in different species.

4 Discussion

Our results show that the specific lateral connectivity necessary for contour integration can be due to input-driven self-

organization. The same self-organization mechanism has previously been shown potentially responsible for orientation, ocular dominance, and frequency columns and patchy connections between them, for repair after cortical and retinal damage, and for tilt aftereffects (Miikkulainen et al. 1997), providing a unified explanation of several different phenomena in the visual cortex. The main new idea advanced in this paper is that long-range *excitatory* lateral connections can also self-organize into highly specific patterns that serve a perceptual grouping function.

The connection patterns that emerge in the model closely approximate those found in neurophysiological experiments (Bosking et al. 1997; Fitzpatrick et al. 1994), and are very similar to the local contour grouping statistics found in natural images (Geisler et al. 1999, 2001). They also generally agree with connection patterns hypothesized in hand-coded computational models (Li 1998; Ross et al. 2000; Yen and Finkel 1997, 1998). We also demonstrated that synchronized firing of neuronal populations can represent the percept of contour very well, by comparing correlations to human contour integration accuracy with varying degrees of orientation jitter (Field et al. 1993; Geisler et al. 1999, 2001).

The input patterns studied in this paper are decidedly simple for two reasons: (1) this way it is possible to characterize and measure model behavior clearly, without confounding factors, and (2) more complex patterns would require larger networks, which are computationally too expensive to simulate at the moment. For example, the current self-organization

simulations required about 200MB of memory and took about 20 hours on a 1.7 GHz Pentium PC. To represent more complex inputs, the number of rows and columns would have to be scaled up by a factor of four, resulting in a simulation with over 40GB of memory and a training time of over 4800 hours. However, there is a good reason to believe that the model will scale up well: it is based on regular patterns of connectivity that can be duplicated horizontally, resulting in a larger-scale model with similar behavior. In a parallel line of research, we have developed methods for such incremental scaling of self-organizing firing-rate models (Bednar et al. 2002); applying these methods to the contour integration task is a most interesting direction of future work. The temporal behavior of the model should also scale up well. Campbell et al. (1999) recently showed that time to synchronization in locally connected integrate-and-fire neurons is logarithmically proportional to the network size. Since the dynamic threshold neuron used in the current model is equivalent to integrate-and-fire neurons, we expect our model to show similar, manageable temporal scaling behavior as the network size is increased. In the near future, sufficient computational power might exist to train the model with natural images. Based on analogous results with firing-rate models (Bednar et al. 2002), we expect the results with more complex images to be similar to those of the current model.

Whether contour integration in the model occurs or not depends on whether the appropriate lateral connections exist or not. Integration is possible only if focused (i.e. patchy) lateral connections link neurons with similar orientation preferences. Even though the integration and adaptation mechanisms might be the same throughout the cortex, if the input to the different areas differs during development, different contour integration performance results. The model therefore suggests why the performance e.g. in the upper vs. lower hemifield (or in fovea vs. periphery) might differ: if the upper visual field does not receive sufficiently dense visual input during development, its lateral connections remain diffuse, resulting in weaker integration. We plan to test this hypothesis in the future with a model that also takes into account the structural differences in these areas, such as different receptor densities. In this way, the observed differences in contour integration performance can possibly be explained as an effect of input-driven self-organization.

Statistics of images projected on the retina indeed support the idea that input distributions may differ among different visual areas. Reinagel and Zador (1999) showed that human gaze most often falls upon areas with high contrast and low pixel correlation. As a result, sharper images may project more often on the fovea than the periphery, allowing more specific connections to form. A similar method can be used to find out if there is a difference in statistical distribution of image features in the lower vs. upper hemisphere. It seems likely, based on the observation that primates mostly manipulate objects in their lower hemifield (Previc 1990). Such




Figure 12: **Simultaneous Activation of Neurons.** The plot shows two representative cases of coactivation (i.e. when two neurons are activated simultaneously), when a long (dashed line) input is presented across the two receptive fields. (a) Collinear arrangement: the two receptive fields (thick bars) are precisely aligned. If a long input is presented along the same direction, the two neurons will respond maximally, and the connection between them becomes stronger. (b) Co-circular arrangement: even though the two receptive fields are slightly misaligned, they are still weakly activated and their connection is strengthened, although less so than in (a).

statistical differences together with Hebbian self-organization would then result in different contour integration capability in different visual areas, as was demonstrated in section 3.4.

A competing hypothesis would be that the differences between hemifields (as well as those between fovea and periphery) are genetically determined. One way of distinguishing between these hypotheses would be to rear an animal with eye glasses that flip the input to the upper and lower hemifield. After the critical period, the animal's performance on contour detection task could be measured, and the connectivity patterns formed in the upper and lower hemifield compared to normally reared control animals. With genetic determination there should be no noticeable difference, whereas PGLISSOM predicts that high connectivity and good integration would occur in the upper hemifield, instead of the lower hemifield as in control animals.

The fact that even simple patterns such as straight Gaussian bars shape the circuitry for contour integration is an interesting result. It supports a previous proposal by Bednar and Miikkulainen (1998, 2000a) that simple internally generated patterns in the developing nervous system may pretrain the cortex before birth, explaining why a certain degree of organization and functionality already exists in a newborn cortex. However, since the PGLISSOM model was trained with straight Gaussian bars, one would expect only collinear properties to emerge in the connection profile, instead of the co-circular patterns actually observed (figure 5). Such an unexpected result follows from Hebbian learning on graded responses (figure 12). Neurons co-activate even if their receptive fields are not perfectly aligned, allowing co-circular connections to develop along with the collinear ones. Such graded training in general matches the regularities in the visual environment, forming a robust starting point for learning more refined regularities in the visual input.

As we have seen in this paper, connection statistics, feature co-occurrence statistics, and performance are very closely related. It may be possible to measure co-occurrence statistics of visual features other than orientation as well, and such statistics can be used to derive hypotheses about the *func-*

tional connectivity of visual cortical areas. Thus, perceptual grouping rules employed by the brain can be systematically investigated by examining the statistical structure in natural scenes.

5 Conclusion

This paper shows how the specific connection patterns that may facilitate contour integration and segmentation in the visual cortex can be due to the same general process of input-driven self-organization as many other cortical structures. The contour integration performance measured by the degree of synchronization in the model matches human performance data very well, lending further support for the idea that segmentation and binding could be due to synchronized firing of neuronal groups. The model also suggests that differently distributed input presentations and the resulting lateral connections may be the cause for the different degrees of contour integration observed in the different visual areas. It should be possible to account for other low-level Gestalt phenomena with similar computational principles.

Acknowledgments.

Thanks to Wilson Geisler for the human data and to James Bednar for insightful comments on an earlier version of the manuscript. We would also like to thank the anonymous reviewer for helpful comments on section 3.4. This research was supported in part by National Science Foundation grant IRI-9811478, National Institutes of Mental Health (Human Brain Project) grant 1R01-MH66991, and Texas Higher Education Coordinating Board grant 000512-0217-2001.

A Appendix: Simulation Setup

This section describes the simulation setup in detail for accurate replication of the results presented in this paper. The code and simulation configuration files can be found on the world-wide-web at <http://www.cs.tamu.edu/faculty/choe>.

A.1 Network

While MAP1 consisted of 136×136 neurons, MAP2 was reduced to 54×54 to save simulation time and memory. The intra-columnar connections between MAP1 and MAP2 were proportional to scale, so that the relative locations of corresponding neurons in the two maps were the same. However, different parameter values were required for the two maps, corresponding to their different sizes. Excitatory lateral connections in MAP1 had an initial radius of 7 and gradually reduced to 3, and inhibitory lateral connections had a fixed radius of 10. Initially, large areas have correlated activity so that global order can be formed, and later on, the reduced lateral excitatory connections help fine-tune the local order in the map (Kohonen 1982, 1989, 1993; Sirosh and Miikkulainen 1997). In MAP2, excitatory lateral connections had a radius of 40 and inhibitory connections 54. Afferent connections to the retina had a radius of 6 in both maps, and intra-columnar connections a radius of 2 in both maps. The retina

consisted of 46×46 receptors, except for section 3.4 where it was 72×72 to have sufficiently large lower and upper hemifields for the experiments. As long as the relative sizes of the map, the retina, and the lateral connection radii are similar to these values, the maps self-organize well (see Bednar et al. (2002) for precise equations that allow scaling maps to different sizes).

A.2 Self-Organization

This section describes the simulation setup used in sections 3.1 and 3.4. The input in the training experiment consisted of straight oriented Gaussians:

$$\xi_{r_1, r_2} = \exp\left(-\frac{((r_1 - x)\cos(\phi) - (r_2 - y)\sin(\phi))^2}{a^2} - \frac{((r_1 - x)\sin(\phi) + (r_2 - y)\cos(\phi))^2}{b^2}\right), \quad (12)$$

where ξ_{r_1, r_2} is the desired activity of the retinal neuron at location (r_1, r_2) , a^2 and b^2 specify the length along the major and minor axes of the Gaussian, and ϕ specifies its orientation. These axis lengths were $a^2 = 15.0$ and $b^2 = 0.6$ for the first 1,000 iterations, and they were increased to 45.0 and 0.45 thereafter, except for section 3.4. In that case, the axis lengths were $a^2 = 15.0$ and $b^2 = 1.3$ at the beginning and increased to 50.0 and 0.8 by iteration 10,000, to compensate for the enlarged retina size. All weights were initialized with uniform random numbers within $[0..1]$. The relative contributions of afferent, lateral excitatory, lateral inhibitory, and intra-columnar connections were $\gamma_a = 1.1$, $\gamma_e = 0.8$, $\gamma_i = 0.9$, and $\gamma_c = 0.5$ for MAP1 and $\gamma_a = 1.1$, $\gamma_e = 0.2$, $\gamma_i = 2.5$, and $\gamma_c = 0.9$ for MAP2. The learning rates of afferent, lateral excitatory, lateral inhibitory, and intra-columnar connections were $\alpha_a = 0.012$, $\alpha_e = 0.008$, $\alpha_i = 0.008$, and $\alpha_c = 0.012$ for MAP1 and $\alpha_a = 0.012$, $\alpha_e = 0.008$, $\alpha_i = 0.0$, and $\alpha_c = 0.012$, for MAP2. At 5,000 iterations, α_a and α_c in both maps were decreased to 0.008 so that the global order in the map could start stabilizing. Initial base threshold θ_{base} for both maps was 0.05. At the beginning of each settling iteration, the θ_{base} was adjusted to 50% of $\max_{i,j}(\sigma_{i,j}(t))$ so that the network would not become too active or totally silent. Later, the percentile was increased to 57.5% at 15,000 iterations for MAP1, and 65% at 5,000 for MAP2. While organized maps can be obtained without such parameter adaptation, it generally leads to better results. Interestingly, biological evidence also supports such adaptation processes during learning, including both threshold adaptation (Azouz and Gray 2000; Prince and Huguenard 1988) and synaptic plasticity (Caleo and Maffei 2002).

The synaptic decay rates were different for different types of connections. Previous sum was decayed by $e^{-\lambda}$, where $\lambda = 3.0, 0.5, \text{ and } 1.0$ for lateral excitatory, inhibitory, and intra-columnar connections for both maps. The decay rate in the spike generator's inhibitory feedback $\lambda_{\text{rel}} = 0.5$ in both maps. The relative contribution of the inhibitory feedback in

dynamic threshold calculation $\tau = 0.4$ in both maps. The threshold and ceiling of the linear approximation of the sigmoid function $g(\cdot)$ were $\delta = 0.01$ and $\beta = 1.3$ in both maps. For the average spiking rate of neurons, a running average with the rate $\tau_{\text{avg}} = 0.92$ was calculated.

A.3 Contour Integration and Segmentation

This section describes the setup for contour integration experiments in sections 3.2, 3.3, and 3.4. The input to the network consisted of oriented Gaussians of length $a^2 = 3.5$ and width $b^2 = 1.5$ (equation 12). Examples are shown in figure 6. For the training, a long Gaussian was necessary, but for the contour integration experiments, they were short enough to fit into a single receptive field (afferent connection radius = 6).

The network configuration and parameters were the same as in appendix A.2 except for the following changes: The lateral excitatory connections in MAP2 with weights less than 0.001 were deleted, modeling death of unused connections (Katz and Callaway 1992), and γ_e in MAP2 was increased to 0.8. In addition, the excitatory learning rate α_e in MAP2 was set to 0.1. Although not strictly necessary for grouping, this fast learning makes the patchy connections more uniform, and helps promote synchrony among the connected regions; it does not affect the patchy structure of the connections nor the organization of the map. Fast adaptation has been proposed to be useful in several forms (von der Malsburg 1981; Crick 1984; Wang 1996), but it remains to be verified in biological systems, and its role in smoothing the response constitutes a further prediction of the model. To help desynchronization (segmentation) and model synaptic noise, MAP2 γ_i was increased to 5.0 and 4% noise was added. Previously, for self-organization, the absolute refractory period (κ_{abs}) was set to 0. The firing rates of the neurons were high as a result and the simulation proceeded in a fast time-scale. For the contour integration experiments, a finer degree of temporal resolution was necessary, so κ_{abs} was increased to 4.

References

- Azouz, R., and Gray, C. M. (2000). Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proceedings of the National Academy of Sciences, USA*, 97:8110–8115.
- Bartsch, A. P., and van Hemmen, J. L. (2001). Combined Hebbian development of Genuiculocortical and lateral connectivity in a model of primary visual cortex. *Biological Cybernetics*, 84:41–55.
- Bednar, J. A. (1997). *Tilt Aftereffects in a Self-Organizing Model of the Primary Visual Cortex*. Master’s thesis, Department of Computer Sciences, The University of Texas at Austin. Technical Report AI-TR-97-259.
- Bednar, J. A. (2002). *Learning to See: Genetic and Environmental Influences on Visual Development*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin. Technical Report AI-TR-02-294.
- Bednar, J. A., Kelkar, A., and Miikkulainen, R. (2002). Modeling large cortical networks with growing self-organizing maps. In Bower, J. M., editor, *Computational Neuroscience: Trends in Research, 2002*, 315–321. New York: Elsevier.
- Bednar, J. A., and Miikkulainen, R. (1998). Pattern-generator-driven development in self-organizing models. In Bower, J. M., editor, *Computational Neuroscience: Trends in Research, 1998*, 317–323. New York: Plenum.
- Bednar, J. A., and Miikkulainen, R. (2000a). Self-organization of innate face preferences: Could genetics be expressed through learning?. In *Proceedings of the 17th National Conference on Artificial Intelligence*, 117–122. Cambridge, MA: MIT Press.
- Bednar, J. A., and Miikkulainen, R. (2000b). Tilt aftereffects in a self-organizing model of the primary visual cortex. *Neural Computation*, 12(7):1721–1740.
- Blasdel, G. G. (1992). Orientation selectivity, preference, and continuity in monkey striate cortex. *Journal of Neuroscience*, 12:3139–3161.
- Blasdel, G. G., and Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321:579–585.
- Bosking, W. H., Zhang, Y., Schofield, B., and Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *Journal of Neuroscience*, 17(6):2112–2127.
- Burger, T., and Lang, E. W. (1999). An incremental Hebbian learning model of the primary visual cortex with lateral plasticity and real input patterns. *Zeitschrift für Naturforschung C—A Journal of Biosciences*, 54:128–140.
- Caleo, M., and Maffei, L. (2002). Neurotrophins and plasticity in the visual cortex. *Neuroscientist*, 8:52–61.
- Campbell, S. R., Wang, D. L., and Jayaprakash, C. (1999). Synchrony and desynchrony in integrate-and-fire oscillators. *Neural Computation*, 11:1595–1619.
- Choe, Y. (2001). *Perceptual Grouping in a Self-Organizing Map of Spiking Neurons*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX. Technical Report AI01-292.
- Choe, Y., and Miikkulainen, R. (1998). Self-organization and segmentation in a laterally connected orientation map of spiking neurons. *Neurocomputing*, 21:139–157.
- Crick, F. (1984). Function of the thalamic reticular complex: The searchlight hypothesis. *Proceedings of the National Academy of Sciences, USA*, 81:4586–4950.
- Eckhorn, R., Bauer, R., Jordan, W., Kruse, M., Munk, W., and Reitboeck, H. J. (1988). Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics*, 60:121–130.

- Eckhorn, R., Reitboeck, H. J., Arndt, M., and Dicke, P. (1990). Feature linking via synchronization among distributed assemblies: Simulations of results from cat visual cortex. *Neural Computation*, 2:293–307.
- Engel, A. K., Kreiter, A. K., König, P., and Singer, W. (1991). Synchronization of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. *Proceedings of the National Academy of Sciences, USA*, 88:6048–6052.
- Field, D. J., Hayes, A., and Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local association field. *Vision Research*, 33:173–193.
- Fitzpatrick, D., Schofield, B. R., and Strote, J. (1994). Spatial organization and connections of iso-orientation domains in the tree shrew striate cortex. In *Society for Neuroscience Abstracts*, vol. 20, 837.
- Geisler, W. S., Perry, J. S., Super, B. J., and Gallogly, D. P. (2001). Edge Co-occurrence in natural images predicts contour grouping performance. *Vision Research*. 711–724.
- Geisler, W. S., Thornton, T., Gallogly, D. P., and Perry, J. S. (1999). Image structure models of texture and contour visibility. In *NATO Workshop on Search and Target Acquisition*.
- Gerstner, W. (1998). Spiking neurons. In Maass, W., and Bishop, C. M., editors, *Pulsed Neural Networks*, chapter 1, 3–54. MIT Press.
- Gilbert, C. D., and Wiesel, T. N. (1990). The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Research*, 30(11):1689–1701.
- Goodhill, G. J., and Cimoneriu, A. (2000). Analysis of the elastic net model applied to the formation of ocular dominance and orientation columns. *Network*, 11:153–168.
- Gray, C. M. (1999). The temporal correlation hypothesis of visual feature integration: Still alive and well. *Neuron*, 24:31–47.
- Gray, C. M., König, P., Engel, A., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit intercolumnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337.
- Gray, C. M., and Singer, W. (1987). Stimulus specific neuronal oscillations in the cat visual cortex: A cortical functional unit. In *Society for Neuroscience Abstracts*, vol. 13, 404.3.
- Grinvald, A., Lieke, E. E., Frostig, R. D., and Hildesheim, R. (1994). Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. *Journal of Neuroscience*, 14:2545–2568.
- Grossberg, S., and Grunewald, A. (1997). Cortical synchronization and perceptual framing. *Journal of Cognitive Neuroscience*, 9:106–111.
- Grossberg, S., and Williamson, J. R. (2001). A neural model of how horizontal and interlaminar connections of visual cortex develop into adult circuits that carry out perceptual grouping and learning. *Cerebral Cortex*, 9:878–895.
- Hata, Y., Tsumoto, T., Sato, H., Hagihara, K., and Tamura, H. (1988). Inhibition contributes to orientation selectivity in the visual cortex of the cat. *Nature*, 335:815–817.
- Hess, R. F., and Dakin, S. C. (1997). Absence of contour linking in peripheral vision. *Nature*, 390:602–604.
- Hirsch, J. A., and Gilbert, C. D. (1991). Synaptic physiology of horizontal connections in the cat’s visual cortex. *Journal of Neuroscience*, 11:1800–1809.
- Horn, D., and Opher, I. (1998). Collective excitation phenomenon and their applications. In Maass, W., and Bishop, C. M., editors, *Pulsed Neural Networks*, chapter 11, 297–320. MIT Press.
- Kammen, D. M., Holmes, P. J., and Koch, C. (1989). Origin of oscillations in visual cortex: Feedback versus local coupling. In Cotterill, R. M. J., editor, *Models of Brain Functions*, 273–284. Cambridge, England: Cambridge University Press.
- Katz, L. C., and Callaway, E. M. (1992). Development of local circuits in mammalian visual cortex. *Annual Review of Neuroscience*, 15:31–56.
- Kohonen, T. (1981). Automatic formation of topological maps of patterns in a self-organizing system. In *Proceedings of the 2nd Scandinavian Conference on Image Analysis*, 214–220. Espoo, Finland: Pattern Recognition Society of Finland.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69.
- Kohonen, T. (1989). *Self-Organization and Associative Memory*. Berlin; New York: Springer. Third edition.
- Kohonen, T. (1993). Physiological interpretation of the self-organizing map algorithm. *Neural Networks*, 6:895–905.
- Kohonen, T. (1995). *Self-Organizing Maps*. Berlin; New York: Springer.
- Li, Z. (1998). A neural model of contour integration in the primary visual cortex. *Neural Computation*, 10:903–940.
- Miikkulainen, R., Bednar, J. A., Choe, Y., and Sirosh, J. (1997). Self-organization, plasticity, and low-level visual phenomena in a laterally connected map model of the primary visual cortex. In Goldstone, R. L., Schyns, P. G., and Medin, D. L., editors, *Perceptual Learning*, vol. 36 of *Psychology of Learning and Motivation*, 257–308. San Diego, CA: Academic Press.
- Miller, K. D. (1994). 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. *Journal of Neuroscience*, 14:409–441.
- Obermayer, K., Ritter, H. J., and Schulten, K. J. (1990). A principle for the formation of the spatial structure of cortical feature maps. *Proceedings of the National Academy of Sciences, USA*, 87:8345–8349.
- Pettet, M. W., McKee, S. P., and Grzywacz, N. M. (1998). Constraints on long range interactions mediating contour detection. *Vision Research*, 38:865–879.
- Previc, F. H. (1990). Functional specialization in the lower and upper visual fields in humans: Its ecological origins and neurophysiological implications. *Behavioral and Brain Sciences*, 13:519–575.
- Prince, D. A., and Huguenard, J. R. (1988). Functional properties of neocortical neurons. In Rakic, P., and Singer, W., editors, *Neurobiology of Neocortex*, 153–176. New York: Wiley.
- Reinagel, P., and Zador, A. M. (1999). Natural scene statistics at the center of gaze. *Network: Computation in Neural Systems*, 10:1–10.
- Reitboeck, H., Stoecker, M., and Hahn, C. (1993). Object separation in dynamic neural networks. In *Proceedings of the IEEE International Conference on Neural Networks* (San Francisco, CA), vol. 2, 638–641. Piscataway, NJ: IEEE.
- Ross, W. D., Grossberg, S., and Mingolla, E. (2000). Visual cortical mechanisms of perceptual grouping: Interacting layers, networks, columns, and maps. *Neural Networks*, 13:571–588.
- Rubin, N., Nakayama, K., and Shapley, R. (1996). Enhanced perception of illusory contours in the lower versus upper visual hemifields. *Science*, 271:651–653.
- Singer, W. (1999). Neuronal synchrony: A versatile code for the definition of relations? *Neuron*, 24:49–65.
- Singer, W., and Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, 18:555–586.
- Sirosh, J. (1995). *A Self-Organizing Neural Network Model of the Primary Visual Cortex*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX. Technical Report AI95-237.
- Sirosh, J., and Miikkulainen, R. (1997). Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. *Neural Computation*, 9:577–594.
- Sirosh, J., Miikkulainen, R., and Choe, Y., editors (1996). *Lateral Interactions in the Cortex: Structure and Function*. Austin, TX: The UTCS Neural Networks Research Group. Electronic book, ISBN 0-9647060-0-8, <http://www.cs.utexas.edu/users/nn/webpubs/htmlbook96>.
- Swindale, N. V. (1996). The development of topography in the visual cortex: A review of models. *Network – Computation in Neural Systems*, 7:161–247.
- Terman, D., and Wang, D. (1995). Global competition and local cooperation in a network of neural oscillators. *Physica D*, 81:148–176.
- VanRullen, R., Delorme, A., and Thorpe, S. J. (2001). Feed-forward contour integration in primary visual cortex based on asynchronous spike propagation. *Neurocomputing*, 38–40:1003–1009.
- von der Malsburg, C. (1973). Self-organization of orientation-sensitive cells in the striate cortex. *Kybernetik*, 15:85–100.
- von der Malsburg, C. (1981). The correlation theory of brain function. Internal Report 81-2, Department of Neurobiology, Max-Planck-Institute for Biophysical Chemistry, Göttingen, Germany.
- von der Malsburg, C. (1986). A neural cocktail-party processor. *Biological Cybernetics*, 54:29–40.
- von der Malsburg, C. (1987). Synaptic plasticity as basis of brain organization. In Changeux, J.-P., and Konishi, M., editors, *The Neural and Molecular Bases of Learning*, 411–432. New York: Wiley.
- von der Malsburg, C., and Buhmann, J. (1992). Sensory segmentation with coupled neural oscillators. *Biological Cybernetics*, 67:233–242.
- Wang, D. (1996). Synchronous oscillations based on lateral connections. In Sirosh, J., Miikkulainen, R., and Choe, Y., editors, *Lateral Interactions in the Cortex: Structure and Function*. Austin, TX: The UTCS Neural Networks Research Group. Electronic book, ISBN 0-9647060-0-8, <http://www.cs.utexas.edu/users/nn/webpubs/htmlbook96>.
- Wang, D. (2000). On connectedness: A solution based on oscillatory correlation. *Neural Computation*, 12:131–139.
- Wang, D. L. (1995). Emergent synchrony in locally coupled neural oscillators. *IEEE Transactions on Neural Networks*, 6:941–948.
- Weliky, M., Kandler, K., Fitzpatrick, D., and Katz, L. C. (1995). Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. *Neuron*, 15:541–552.
- Yen, S.-C., and Finkel, L. H. (1997). Identification of salient contours in cluttered images. In *Computer Vision and Pattern Recognition*, 273–279.
- Yen, S.-C., and Finkel, L. H. (1998). Extraction of perceptually salient contours by striate cortical networks. *Vision Research*, 38:719–741.