In press , Neurocomputing (1998) **URL:***http://www.cs.utexas.edu/users/nn/*

Self-Organization and Segmentation in a Laterally Connected Orientation Map of Spiking Neurons

Yoonsuck Choe and Risto Miikkulainen
Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712 USA
yschoe,risto@cs.utexas.edu

Abstract

The RF-SLISSOM model integrates two separate lines of research on computational modeling of the visual cortex. Laterally connected self-organizing maps have been used to model how afferent structures such as orientation columns and patterned lateral connections can simultaneously self-organize through input-driven Hebbian adaptation. Spiking neurons with leaky integrator synapses have been used to model image segmentation and binding by synchronization and desynchronization of neuronal group activity. Although these approaches differ in how they model the neuron and what they explain, they share the same overall layout of a laterally connected two-dimensional network. This paper shows how both self-organization and segmentation can be achieved in such an integrated network, thus presenting a unified model of development and functional dynamics in the primary visual cortex.

1 Introduction

Several models of the visual cortex that take into account lateral interactions between neurons have recently been proposed (see Sirosh et al. [34] for an overview). These models can potentially account for a wider set of developmental and functional phenomena than Self-Organizing Map (SOM) models without explicit lateral connections [23, 24, 25, 26]. In the early stages of the development of the visual cortex, lateral connections are believed to self-organize in synergy with the afferent connections to form a topological map of the input space [8, 40]. This process can be modeled computationally, showing how structures such as oriented receptive fields, orientation columns, and patterned lateral connections form based on input-driven Hebbian learning process (the Receptive Field - Laterally Interconnected Synergetically Self-Organizing

Keywords: *Self-organization, segmentation, binding, synchronization, spiking neurons, lateral connections*

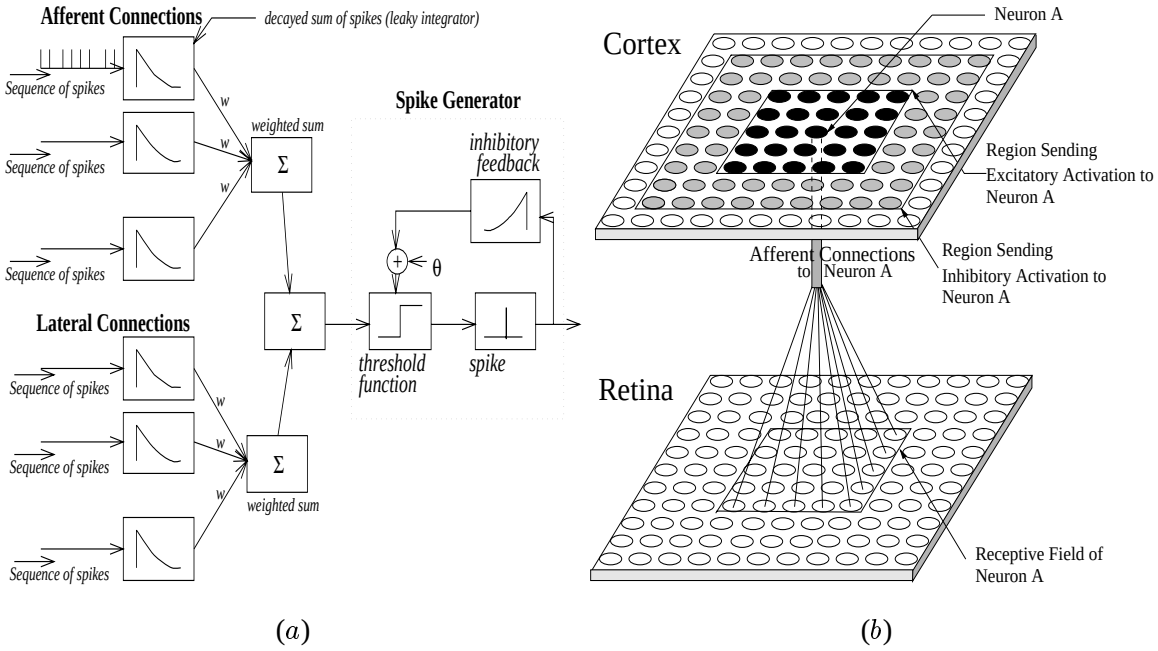

Figure 1: **The RF-SLISSOM Architecture.** (a) The structure of the spiking neuron. Leaky integrators at each synapse perform decayed summation of incoming spikes. The spike generator at each neuron compares the weighted sum of the integrator outputs to a dynamic threshold, firing a spike if the sum exceeds the threshold. Each spike increases the threshold, with exponential decay, making it harder for the neuron to fire again in the near future. (b) The organization of the RF-SLISSOM network. The bottom layer models the retina, and the top layer models the cortical neurons (transformations in the LGN are bypassed for simplicity). The cortical neurons are connected with short-range lateral excitatory connections (from neurons marked black in the cortex) and long-range lateral inhibitory connections (from neurons marked black plus the ones marked gray). Each cortical neuron (i, j) receives input from a receptive field of retinal neurons centered at the corresponding location on the retina.

Map, or RF-LISSOM [22, 29, 31, 32, 33]).

Lateral connections may also play a central role in the function of the visual cortex, by modulating the spiking behavior of neuronal groups. They could cause synchronization and desynchronization of spiking activity, thus mediating feature binding and segmentation. Such synchronization of neuronal activity emerges in the visual cortex of the cat when light bars of various orientation and length are presented [9, 14, 15]. Several models have been proposed to explain this phenomenon [10, 28, 36, 37, 38]. The model of Reitbock et al. [28] is particularly interesting because of its sophisticated model of the neuron with biologically plausible temporal dynamics. The synapses in the model are leaky integrators that sum incoming activation over time with exponential decay. A network of such neurons can segment multiple objects in a scene by synchronizing neuronal group activity. Spikes of neurons representing the same object are syn-

chronized, and those of neurons representing different objects are desynchronized, thus forming a temporal coding of the objects in the scene.

In preliminary work we showed how the leaky integrator model of the spiking neuron can be integrated with a simplified RF-LISSOM model of the cortex. The Spiking Laterally Interconnected Synergetically Self-Organizing Map, or SLISSOM [6, 7, 22] formed a topological map of the retinal locations and generated synchronized and desynchronized neuronal activity that was used for segmenting multiple objects in the scene. The model was simplified in that the cortical neurons were fully connected to the retina and the inhibitory lateral connections covered the whole cortex. In other words, there were no receptive fields, and input features such as orientation could not be taken into account. Instead, the network was trained with small squares at random locations on the retina and the map organized to represent retinotopic location only.

In this paper, the spiking neuron idea is integrated with the full RF-LISSOM model, including local receptive fields and patchy lateral connections in the network, and oriented Gaussian bars as inputs. The resulting Receptive Field SLISSOM (or RF-SLISSOM) model shows how (1) oriented receptive fields, (2) a global map of orientation selective columns, and (3) patterned lateral connections develop from initially random afferent and lateral connections through self-organization. Functionally, (4) the map synchronizes and desynchronizes neuron activity through adapting lateral connections as before, and thereby implements segmentation and binding. However, in RF-SLISSOM also (5) the orientation affects whether two inputs are perceived as different or the same: similarly oriented nearby lines are perceived as parts of the same object, constituting a rudimentary perception of texture. The model therefore suggests specific ways in which the lateral connections can play a central role in both the development and function of the visual cortex.

2 The RF-SLISSOM Architecture

RF-SLISSOM consists of two layers of interconnected neurons: the "retina" and the "cortex". The overall organization of RF-SLISSOM is based on the RF-LISSOM architecture [22, 29, 31, 32, 33], and the neuron model on the leaky integrator neurons of Eckhorn el al. [10] and Reitbock et al. [28]. RF-LISSOM provides the self-organizing structure and the leaky integrator neuron introduces temporal dynamics into the model.

Each cortical neuron receives afferent connections from the input layer and lateral (excitatory and inhibitory) connections from other neurons in the cortex. Each connection is a leaky integrator that performs decayed summation of incoming spikes, thereby establishing not only spatial summation, but also temporal summation of activity (figure 1a). Each new spike is added to the sum of the previous ones, and the sum

is exponentially decayed over time. The current sums are multiplied by the connection weight and added together[1] to form the net input to the neuron. The spike generator compares the net input to a threshold and decides whether to fire a spike. The threshold is a sum of two factors: the base threshold θ and the decayed sum of past spikes, formed by a similar leaky integrator as in the input synapses. The spike generator mechanism models the absolute and relative refractory periods of a neuron. Active spiking increases the effective threshold, making further spiking less likely and keeping the activation of the system within a reasonable range [9, 10].

The overall organization of the RF-SLISSOM model is shown in figure $1b$. Each cortical neuron receives input from a fixed-size receptive field in the retina, centered at the location corresponding to the neuron's location in the cortical network. The neuron also has excitatory connections with neighboring neurons and inhibitory connections with a larger area of the map.[2] Each connection has a queue that stores previous spikes. In calculating the postsynaptic potential, the latest spike has the value of 1.0 and older ones are decayed by $1/e^{\lambda_q}$, where λ_q is the decay parameter, as they are shifted through the queue. The inhibitory feedback loop in the spike generator (figure $1a$) is a similar queue that receives spikes from the spike generator itself, with decay $1/e^{\lambda_s}$.

The input to the network consists of oriented Gaussian bars, defined by

$$\xi_{r_1,r_2} = exp(-\frac{((r_1-x)cos(\phi)-(r_2-y)sin(\phi))^2}{a^2} - \frac{((r_1-x)sin(\phi)+(r_2-y)cos(\phi))^2}{b^2}), \quad (1)$$

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. Each retinal neuron (r_1, r_2) spikes at a constant rate such that the decayed sum of its spikes is equal to ξ_{r_1,r_2}. The oriented Gaussian approximate the spatio-temporal firing patterns in the developing visual pathway [39].

Input is generated at the active retinal neurons and sent through the afferent connections to the cortical neurons. The net input $\sigma_{i,j}$ to the spike generator of the cortical neuron at location (i, j) at time t is calculated by summing the afferent and excitatory lateral contributions and subtracting the inhibitory lateral

[1] This differs from Eckhorn et al. [10] and Reitbock et al. [28] who *multiplied* the weighted sums from afferent connections and those from lateral connections. Multiplying exerts better modulation on the neuronal activity, but disturbs self-organization by rapid fluctuation. In our experiments, modulation turned out to be possible with *additive* neurons as well.

[2] Although most long-range synapses in the cortex are excitatory, they can have inhibitory overall effects through interneurons [16, 17, 18]. The RF-LISSOM model predicts that such long-range inhibition is computationally necessary for self-organization to occur [32].

contributions:

$$\sigma_{i,j}(t) = \gamma_{\mathrm{a}} \sum_{r_1,r_2} \xi_{r_1,r_2} \mu_{ij,r_1r_2} + \gamma_{\mathrm{e}} \sum_{k,l} \eta_{kl}(t-1) E_{ij,kl} - \gamma_{\mathrm{i}} \sum_{k,l} \eta_{kl}(t-1) I_{ij,kl}, \tag{2}$$

where γ_{a}, γ_{e}, and γ_{i} are the relative strengths of the afferent, excitatory, and inhibitory contributions, ξ_{r_1,r_2} is the decayed sum of spikes of the retinal neuron (r_1, r_2), μ_{ij,r_1r_2} is the corresponding afferent connection weight, $\eta_{kl}(t-1)$ is the decayed sum of spikes from the map neuron (k, l) at time $t-1$, and $E_{ij,kl}$ is the corresponding excitatory and $I_{ij,kl}$ the inhibitory lateral connection weight. The spike generator fires a spike if $\sigma_{i,j} > \theta + \vartheta_{i,j}$, where θ is the base threshold and $\vartheta_{i,j}$ the output of the spike generator's leaky integrator.

In the standard RF-LISSOM model, the input is kept constant while the cortical response settles through the lateral connections, forming a concentrated, redundancy-reduced activation pattern (as measured by the kurtosis of the map activity [11, 29]). The RF-SLISSOM goes through a similar settling process. The retinal neurons are spiking constantly at each iteration and the cortical neurons are allowed to exchange spikes. After a while, the neurons reach a stable rate of firing, and this rate is used to modify the weights.[3] Both the afferent and the lateral weights are modified according to the Hebbian principle:

$$w_{ij,mn}(t) = \frac{w_{ij,mn}(t-1) + \alpha V_{ij} X_{mn}}{\mathcal{N}}, \tag{3}$$

where $w_{ij,mn}(t)$ is the connection weight between neurons (i, j) and (m, n), $w_{ij,mn}(t-1)$ is the previous weight, α is the learning rate (α_{a} for afferent, α_{e} for excitatory, and α_{i} for inhibitory connections), V_{ij} and X_{mn} are the average spiking rates of the neurons, and $\mathcal{N}$ is the normalization factor, $\sum_{mn} [w_{ij,mn}(t-1) + \alpha V_{ij} X_{mn}]^2$ for afferent connections and $\sum_{mn} [w_{ij,mn}(t-1) + \alpha V_{ij} X_{mn}]$ for lateral connections. Weight normalization is necessary in Hebbian learning to keep the weights from growing without bounds, but it can also be viewed as a process of redistributing the synaptic resources of each neuron. The exact form of this process, and whether it maintains a fixed balance or regulates it to maintain stable firing patterns, is controversial [35]. For simplicity and uniformity, the sum normalization method given in equation 3 is used in RF-LISSOM [30].

The radius of the lateral excitation is gradually reduced, following a preset schedule, resultin in fine tuning of the map (for a theoretical motivation for this process, see [19, 20, 21, 23, 32], for neurobiological evidence, see [8, 17]). The weight adaptation processes are active both during self-organization and segmentation. The RF-SLISSOM model therefore integrates the spiking leaky integrator model with the

[3] It would also be possible to adapt the weights continuously throughout the settling. However, the settled activation is more focused, and using it speeds up self-organization and helps form a better mapping of the input space.

RF-LISSOM structure, modeling self-organization and functional dynamics of the visual cortex at a more accurate level than earlier models.

3 Experiments

The RF-SLISSOM experiment consists of two parts: (1) self-organization, and (2) segmentation and binding. During self-organization, single Gaussian bars (equation 2) are presented to the network as input. The afferent and lateral connection weights adapt to form a structured orientation map. After the organization has stabilized, multiple objects consisting of Gaussian bars of various orientations are presented on the retina. The connection weights continue adapting, and the network segments the objects by temporally alternating the activity on the map from one object to another.

3.1 Self-Organization of the Orientation Map

The retina consisted of 12×12 units and the cortex had 36×36 neurons. The magnification factor, or the ratio of cortical neurons per retinal neurons, was thus 9, allowing several differently-oriented receptive fields to develop for each location in the retina. The afferent receptive fields were 7×7 squares, comparable in size to the major axis of the Gaussian. The receptive field of neuron (i, j) was centered at $(i/3, j/3)$, and the afferent weights were randomly initialized within $[0.25, 1]$ in the 3×3 center and $[0, 0.75]$ in the perimeter of the receptive field and then normalized as in equation 3 (figure $2a$).

The lateral connection weights were randomly initialized within $[0, 1]$ and normalized (equation 3; figure $2c$). Each neuron received inhibition from a square 31×31 neighborhood, and excitation from a neighborhood initially of size 17×17, gradually decreasing to 7×7 in 2,500 input presentations. At the same time, the lateral inhibitory learning rate α_{i} gradually increased from 0.001 to 0.1, and the afferent learning rate α_{a} from 0.02 to 0.5, while α_{e} was 0.001. Slow adaptation in the beginning captures long-term correlations within the inputs, which is necessary for self-organization. Fast adaptation towards the end facilitates quick modulation of the activity necessary for segmentation (section 4). The scaling factors were $\gamma_{\mathrm{a}} = 0.9$, $\gamma_{\mathrm{e}} = 1.1$, and $\gamma_{\mathrm{i}} = 1.2$. The simulations were not very sensitive to these values. These values were found to give good results experimentally, but the simulations are not very sensitive to the parameters, as long as their relationships remain roughly the same.

During self-organization, single oriented Gaussian spots with $a = 4, b = 1$ were presented at random locations on the retina. The orientation ϕ was randomly chosen from 8 alternatives: vertical, horizontal, the

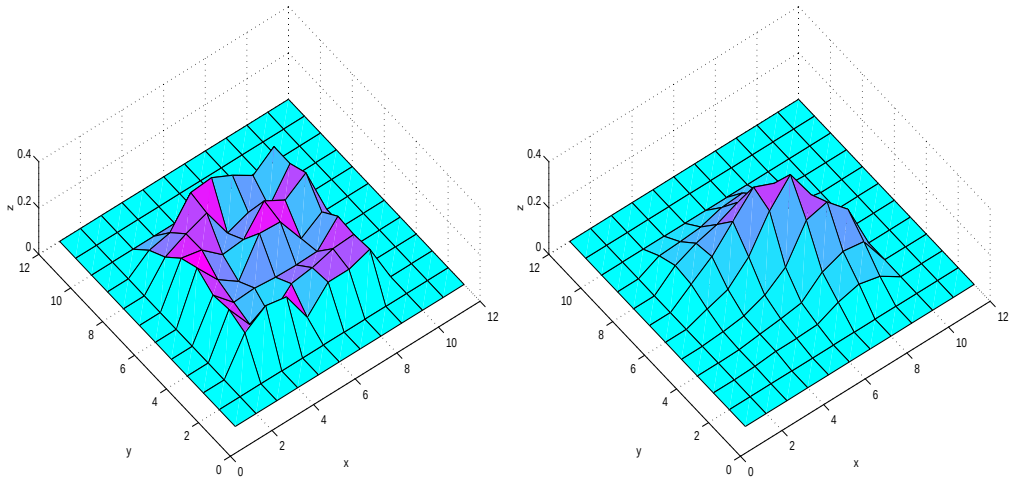

(a) Initial Afferent Weights (b) Final Afferent Weights

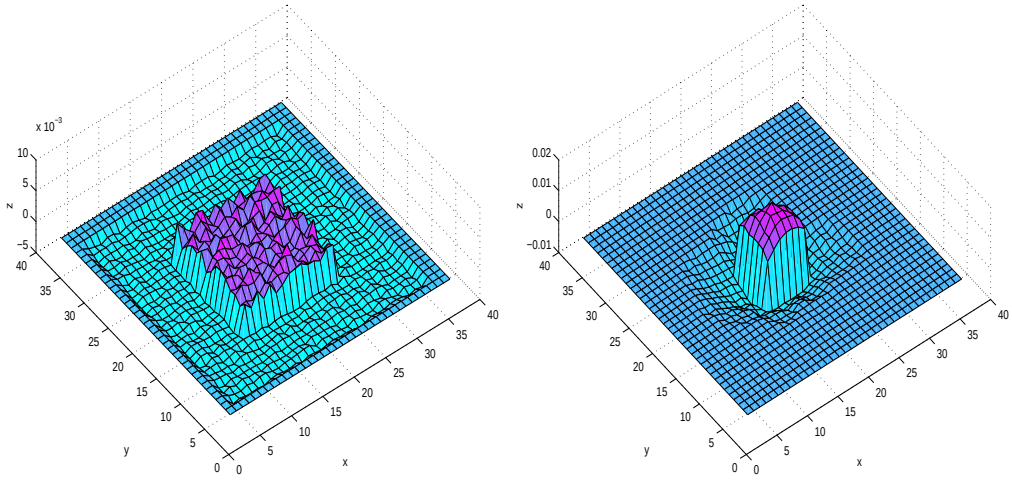

(c) Initial Lateral Interaction (d) Final Lateral Interaction

Figure 2: **Self-Organization of Afferent and Lateral Connections.** Connection weights of cortical neuron $(17, 17)$ are shown here before and after self-organization. The x and y axes in (a) and (b) represent the location of the retinal neuron, and in (c) and (d), the location of the cortical neuron from which the connections come to neuron $(17, 17)$. The connection weights are plotted in the z dimension. (a) The receptive field of this neuron is a 7×7 area on the retina, centered at the retinal neuron $(5, 5)$. The weights are initially random (with a slight emphasis at the center 3×3 area) and normalized, and thus the weight plot shows a rough landscape. (b) The final afferent weights form a well-shaped elongated Gaussian, and are most responsive to Gaussian input with an orientation of 157.5^o. (c) In this plot, the lateral inhibitory weights are subtracted from the excitatory weights, resulting in a combined lateral interaction profile. Since the excitatory connections are short-range (15×15) and the inhibitory connections long-range (31×31), and both weights are initially random and normalized, the profile shows a rugged Difference of Boxes landscape. (d) During self-organization, the excitatory radius was reduced to 3, resulting in a smaller and higher (due to normalization) excitatory center. The long-range inhibitory weights have become patchy, as displayed more clearly in figure 4.

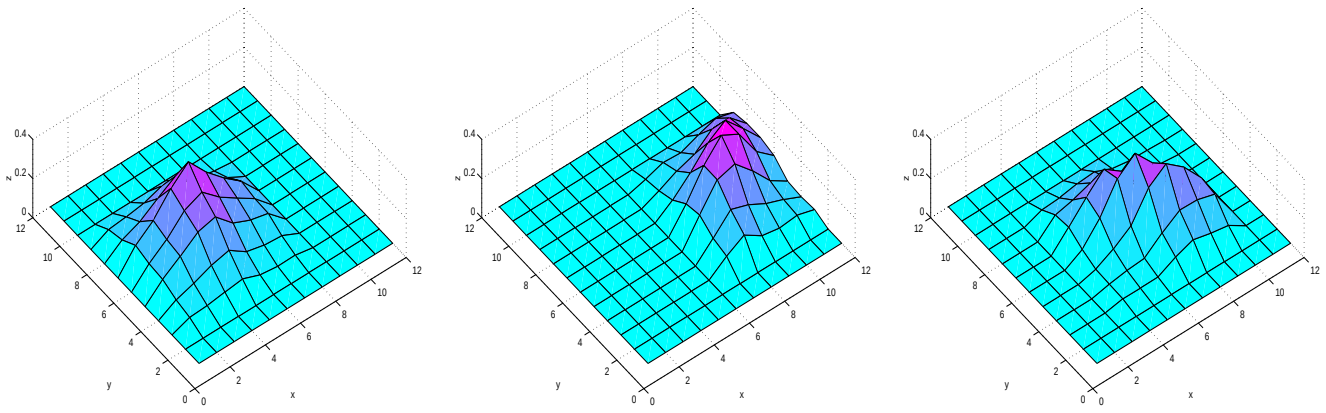

(a) Unoriented RF (b) RF tuned to 45^o (c) RF tuned to 135^o

Figure 3: **Self-Organized Oriented Receptive Fields.** A range of orientation preferences and selectivities develop in the self-organizing process. Three examples are shown in this figure: (a) Neuron $(12, 16)$ has an unoriented, symmetric receptive field. (b) Neuron $(27, 17)$ is sharply tuned to 45^o. (c) Neuron $(19, 16)$ is sharply tuned to 135^o. As in figures $2a$ and b, the x and y axes indicate the location of the retinal neuron from which the afferent weight comes from, and the z axis indicates the connection weight.

two diagonals, and the 4 angles between these orientations. The available orientations were limited so that a small magnification factor could be used, making the network small enough to be simulated computationally.

The retinal neurons were spiking at a continuous rate, and the settling consisted of 13 cycles of cortical activity update (equation 2). After settling, the connection weights were modified according to equation 3, based on the average firing rate over the last 10 cycles. Each such presentation was counted as an iteration. After 5,500 iterations, both the afferent and the lateral weights stabilized into smooth profiles (figure $2b$ and d).

The self-organization of afferents resulted in a variety of oriented receptive fields similar to those found in the visual cortex (figure 3). Some are highly selective to inputs of a particular orientation, others unselective. The global organization of such receptive fields can be visualized by labeling each neuron by the preferred angle and the degree of selectivity to inputs at that angle (figure 4). The resulting orientation map is remarkably similar in structure to those observed in the primary visual cortex by recent imaging techniques [3, 4] and contains structures such as pinwheels (where orientation preferences gradually change 180^o around a single point; figure 4*b*, Box 1), linear zones (where preferences change gradually; Box 2), and fractures (where orientation preference changes abruptly from one orientation to a very different one; Box 3).

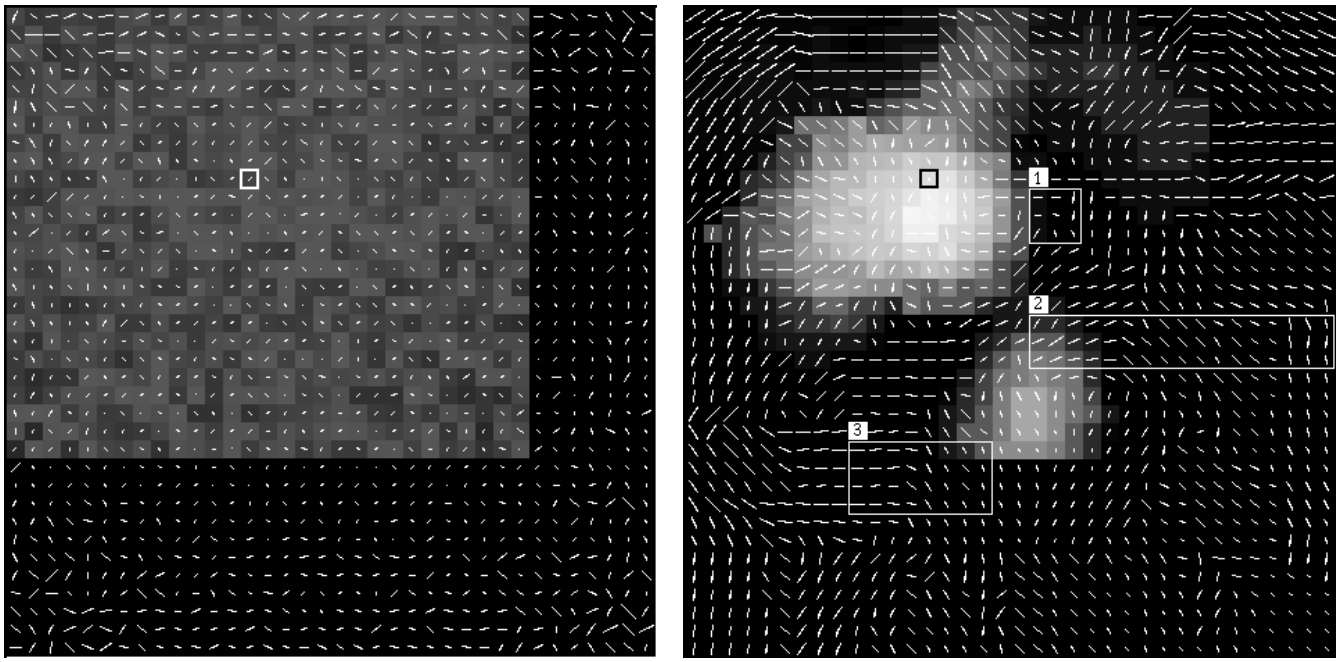

(a) (b)

Figure 4: **Self-Organization of the Orientation Map and Lateral Connections.** Each neuron in the cortical map is represented by a line segment. The orientation of the line indicates the orientation of the Gaussian input that results in the largest response of the neuron, and the length of the line represents how selective the neuron is to the orientation. The length is proportional to O^3/C, where O is the dot product of the afferent weights with the oriented Gaussian bar, and C the dot product of the afferent weights with an unoriented Gaussian pattern. The gray scale background shows the inhibitory lateral connection weights of the neuron bounded by a white box in (a) and a black box in (b). (a) The afferent weights are initially random, and the neurons tend to be unselective. The lateral inhibition to neuron (13,9) initially comes from a wide area with random weights. (b) After 5,500 iterations, the network has formed a structured orientation map. Neurons with similar orientation preferences are organized into intertwined columnar areas, as observed experimentally in the biological visual cortex. Features such as (1) pinwheel centers, (2) linear zones, and (3) fractures are found in this map as well as in biological cortex. The lateral connections of neuron (13,9), which prefers input of 117^o, have become patchy, connecting primarily to neurons with similar orientation preference, and primarily along the 117^o direction.

The lateral connection weights were initially spread over long distances and covered a substantial part of the network (figure 4*a*). During self-organization, the connections between uncorrelated regions became weaker, and in the end, the strong connections link areas of similar orientation preferences (figure 4*b*). Furthermore, the connection patterns are elongated along the direction that corresponds to the neuron's preferred stimulus orientation, as they do in the cortex [13]. This organization reflects the activity correlations caused by the elongated Gaussian input pattern: such a stimulus activates primarily those neurons that are

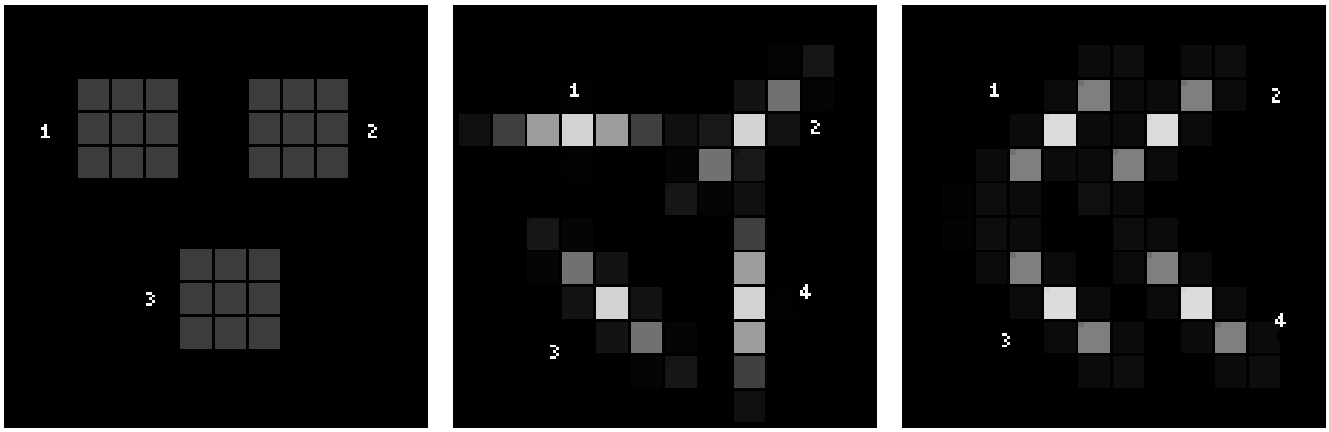

(a) (b) (c)

Figure 5: **Input Used in Segmentation Experiments.** (a) Three separate 3×3 box objects. (b) Four separate Gaussian bars with different orientations. (c) Two textures, each consisting of a pair of equally oriented bars (1 & 2 of 45^o , and 3 & 4 of 135^o; the textures were limited to a pair of bars because that was all that could be fit in the small retina). The brightness of each location shows the leaky integrator activation ξ_{r_1,r_2} resulting from the spiking of each retinal neuron, from black (0.0) to white (1.0). Each input component has the same total activation.

tuned to the same orientation as the stimulus, and located along its length. Similar results have been obtained before with the RF-LISSOM model. The current simulations show that the same results hold for networks with spiking neurons, and therefore RF-SLISSOM constitutes a significant step towards computational understanding of the development in the visual cortex.

3.2 Segmentation and Binding of Objects and Textures

Once the RF-SLISSOM network formed a map of oriented receptive fields and patterned lateral interaction profiles, three different segmentation experiments were conducted on it:

1. Segmentation of three separate unoriented 3×3 boxes (figure $5a$), to see whether the oriented model could still segment unoriented objects,

2. Segmentation of four separate Gaussian bars of different orientations (figure $5b$), to determine whether the model could segment oriented objects, and

3. Binding of two adjacent Gaussian bars into a texture, and segmentation of two adjacent differently-oriented textures into different objects, as a first experiment in texture segmentation (figure $5c$).

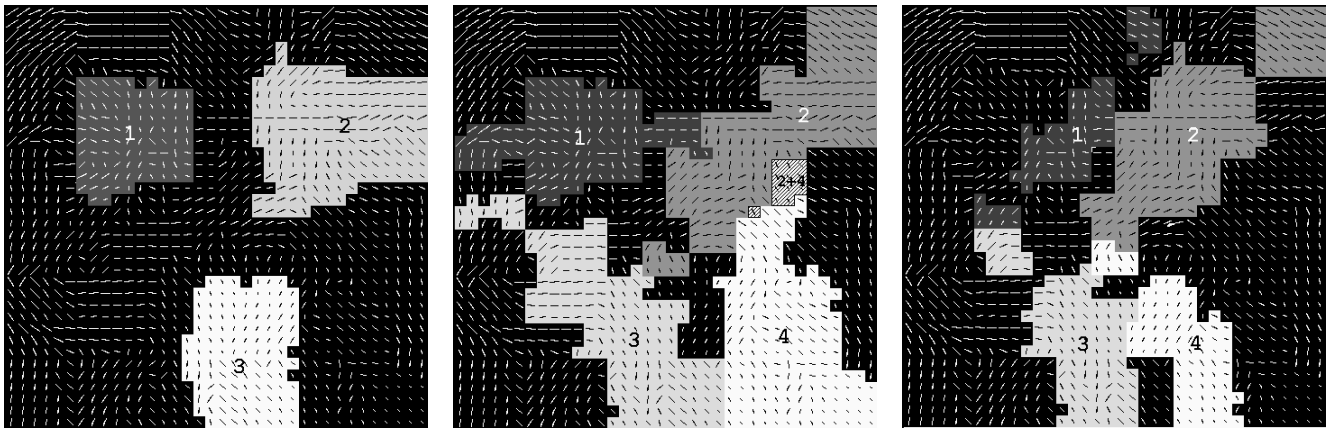

(a) (b) (c)

Figure 6: **Cortical Areas Responding to the Different Input Components.** Figures (a), (b), and (c) correspond to the three different experiments shown in figure 5. In each figure, the units where the multi-unit activity (MUA) was recorded for each different input component (identified by the number) are shown with a gray-scale shading. In forming these areas, all units whose receptive field centered on a unit with an activity of 0.1 or higher were included, regardless of their orientation selectivity.

In each experiment, the input was presented on the retina for 500 time steps, and the lateral connection weights were adapted at each time step according to equation 3. The afferent connection weights were not modified.

In each experiment, each component of the input pattern activated a different area in the cortex, with a slight overlap (figure 6). Activity in any such area indicates that the cortex is currently representing the corresponding input component. The total number of spikes generated within each area per time step (i.e. the multi-unit activity, or MUA) can be used as a measure of how strong the representation is. A high MUA value indicates a strong representation (most neurons in the area are firing together), and a zero value implies that the representation is currently not active.

By plotting the MUAs of the different areas over time, it is possible to see how the cortex performs binding and segmentation. The neurons that represent the same object tend to be active at the same time and silent at the same time, that is, their activity is *synchronized*. Neurons representing different objects are not active at the same time, *desynchronizing* their activations. This way only one area has a high MUA value at any one time, and the activation rotates from one area to another.

For example in the first experiment (figure 7), the three areas do not initially show much synchrony or desynchrony: a small percentage of neurons spike per time step in all areas. After about 25 steps however, the MUAs of different areas start to peak at different times. The neurons within each area exchange primarily

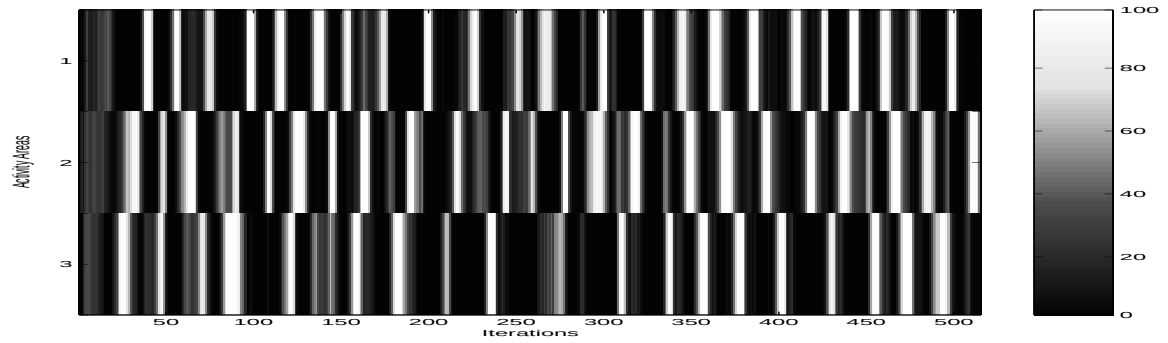

Figure 7: **Segmentation of Unoriented Objects.** The Multi-Unit Activities (MUAs, or the percentage of units spiking per time step) of the three areas of figure 6*a*, responding to the three different objects of figure 5*a*, are shown in gray-scale coding, according to the scale shown at right. Each row represents one area, and the MUA values over time are shown as a linear sequence from left to right. Initially the activity is weak and there is no temporal correlation among the MUAs. The lateral weights adapt, and after about 25 time steps the MUAs of the different areas start to peak at different time steps, indicating segmentation of the input into the three different objects.

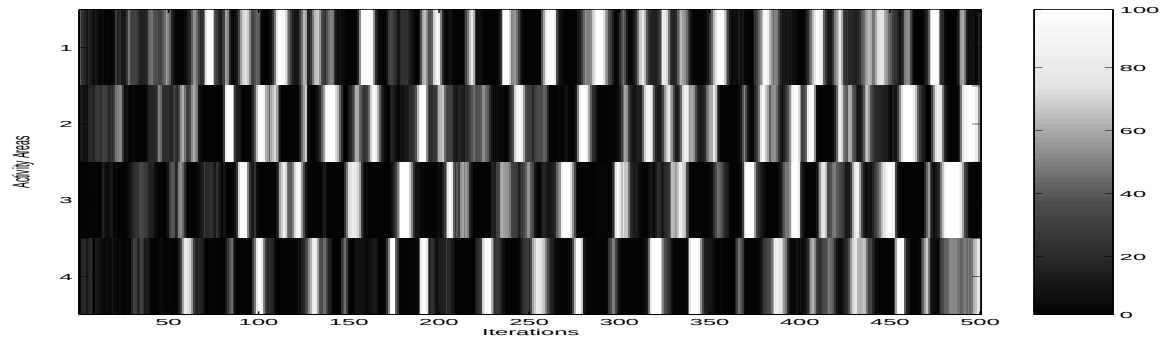

Figure 8: **Segmentation of Four Oriented Objects.** The MUAs of the four areas of figure 6*b*, responding to the four Gaussian bars of different orientation in figure 5*b*, are shown. As in figure 7, the initial activity is equally weak in all areas, but quickly synchronizes with each area and desynchronizes between them, indicating segmentation to four objects.

excitatory activation, whereas the different areas are connected with inhibitory connections. At this time, the connection weights have increased enough to cause activation within each area to synchronize, and between different areas to desynchronize.

Similarly in the second experiment (figure 8), the four areas responding to the four different oriented Gaussian bars are each not very active initially, but as time goes on, they start to peak at different times. Such synchronized and alternating activity indicates that there are four separate objects in the input; in other words, it constitutes a mechanism for binding and segmentation.

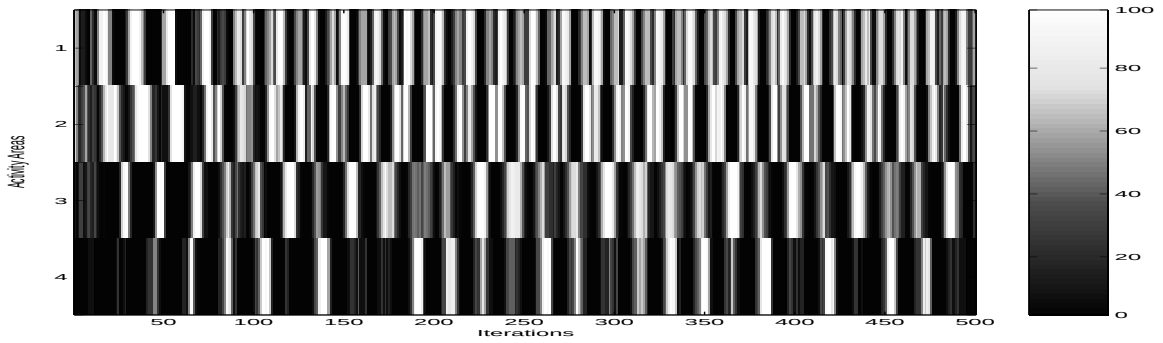

Figure 9: **Binding and Segmentation of Two Texture Areas.** The MUAs of the four areas of figure 6*c*, responding to the four Gaussian bars of figure 5*c*, are shown. All areas are initially desynchronized, but around iteration 100, as excitation increases between bars 1 and 2 on one hand and 3 and 4 on the other, and inhibition increases between the top and the bottom areas, the bars belonging to the same texture are bound together and segmented from the other texture.

In the third experiment (figures 9 and 5*c*), the top two Gaussian bars (1 and 2) on one hand and the bottom two (3 and 4) on the other have the same orientation and a small separation, and can be perceived to belong to the same textured object. Initially, the responses to the different bars are desynchronized, but around step 200, as the lateral connections adapt, short-range excitation between the similarly oriented bars increases, as does the long-range inhibition between the top and bottom responses. This way the bars belonging to the same texture are bound together, and the two textures are segmented. Although the textures are very primitive in this experiment (due to computational limitations on the network size), the mechanisms are general and should scale up well to larger textures.

The quickly adapting lateral connections are crucial in obtaining the segmentation and binding results outlined above. The lateral connections quickly adapt to activity correlations in the map. The short-range excitation establishes binding between the neurons in the same area, causing them to fire in synchrony, and the long-range inhibition segments the areas by making it harder for them to be active at the same time. The leaky integrator threshold is also very important for establishing desynchronization. A neuron that has fired frequently in the past builds up its threshold and eventually becomes inactive. Thus an actively spiking area becomes silent after a while, and neurons that have been silent due to high inhibition from the spiking area cross their thresholds and start spiking. Thus, the adapting lateral connections together with leaky integrator thresholds produce temporally alternating synchronization that amounts to low-level segmentation and binding.

The MUAs in figures 7, 8, and 9 show some overlap even when the input is successfully segmented. This is due to the slightly overlapping receptive fields in the model. Gray et al. [14] observed that in the cat visual cortex, strong phase-locking occurred when the receptive fields were clearly separate. Apparently when they overlap slightly, phase locking becomes less well defined at the border of the areas. The overlap is unavoidable in the current small RF-SLISSOM network, but could be reduced in larger-scale simulations.

4 Discussion

Several studies have shown that fast adaptation of synaptic efficacy is necessary for feature binding through temporal coding [36, 38]. Similarly in the experiments with RF-SLISSOM, rapid adaptation of lateral weights was found necessary for oscillatory behavior. On the other hand, self-organization requires slow adaptation so that long-term correlations can be learned. If the weights are initially random and change rapidly, they will fluctuate a lot and an ill-formed map will result. There are two possible solutions to this problem. One way is to have two sets of lateral connections, one for fast adaptation and the other for slow adaptation [38]. The fast connections become active after the slow connections have converged. The other is to vary the learning rate of the synapse. The learning process starts out with a slow learning rate and gradually the synapses become more plastic. The two solutions are mathematically equivalent and there is no sufficient neurobiological evidence to distinguish between them at this point. The second one is simpler and was therefore chosen for this paper.

The gradually increasing plasticity does not hamper self-organization since the activity on the map becomes more consistent and predictable as the training goes on. After self-organization, only the lateral weights adapt, and the overall organization of the map is therefore preserved. The pattern of patchy lateral connections is disturbed to some degree, but this could be avoided by pruning the weak connections towards the end of the self-organization, as is done in the RF-LISSOM model [5, 27, 30].

One of the fundamental properties of the RF-LISSOM model and also inherited by RF-SLISSOM, is that the map generates a sparse coding of visual input [22, 29, 31, 32, 33]. The self-organized lateral connections represent correlations in the input. Because they are mostly inhibitory, redundant activity is reduced, and the number of active neurons for each input minimized. Such coding is thought to be well suited for the detection of suspicious coincidences, associative memory, and feature grouping [1, 2, 11, 12]. The RF-SLISSOM model introduces temporal dynamics to sparse coding. Leaky integrator neurons with dynamic thresholds desynchronize representations of different objects, forming a temporal sparse coding as

well. The model therefore predicts that both spatial and temporal sparse coding is used by the visual cortex to efficiently represent and process the large amounts of visual information presented by the environment.

The current RF-SLISSOM simulations are limited by the size of the network. Due to the large number of lateral connections in the model, the memory requirements and execution time grows very quickly as the network size is increased, roughly in the order of $O(n^4)$. Although the retina-to-cortex magnification factor of 9 was sufficient for the network to form an organized orientation map with the appropriate features, the map was not quite as smooth and regular as those obtained in RF-LISSOM simulations on massively parallel supercomputers with a magnification factor of 64 [22, 29, 31, 32, 33]. Larger magnification factor allows a larger number of cortical neurons to represent the features of the input at each location, resulting in a better mapping of the input space. Thus, a more accurate representation of gradually changing orientation preference requires a simulation with a larger magnification factor. Similarly, representing other features such as ocular dominance and spatial frequency [22] is possible but increases the computational requirements even further. Such simulations should be possible with the newest generation of parallel supercomputers, and constitute the most immediate direction of future work.

The extent of lateral connections is a crucial factor in binding and segmentation. The excitatory lateral connections determine how large an object can be identified as a single entity through binding. If the object generates an activity area wider than the lateral excitation radius, the activity tends to desynchronize and settle into multiple areas. This phenomenon has been observed in simulations where the excitatory connection radius was reduced to 1. In such cases, the inhibitory input exceeds the excitation, and the network is unable to hold the activity areas together. On the other hand, the extent of the inhibitory lateral connections determines the distance of objects that can be segmented. If the objects are further apart, they do not inhibit each other, and desynchronization may not occur. Thus, segmentation and binding depends on actual lateral connections in the model, predicting that these processes occur only locally at the early stages of visual processing. Binding and segmentation in the global context would have to be done at a higher level. Future work therefore includes experiments with hierarchically arranged RF-SLISSOM maps, where temporally and spatially sparse features detected by lower level RF-SLISSOM maps are combined into global representations of the visual environment at higher levels.

5 Conclusion

The RF-SLISSOM model demonstrates how a network of dynamic spiking neurons in a laterally connected synergetically self-organizing map can be used to explain a number of developmental and functional phenomena in the visual cortex. Based on rudimentary visual input, the model self-organizes and develops orientation-selective receptive fields and patterned lateral connections, forming a global orientation map with similar features as observed in neurobiological experiments. The temporal dynamics of the leaky integrator neurons combined with fast adaptation of the lateral weights provides a mechanism for low-level segmentation and binding. Thus, self-organization and segmentation are integrated into a single model, using the same architectural elements and the same Hebbian learning rule. We believe the model incorporates crucial principles of the biological visual cortex, and if simulated in a larger scale and tested systematically, can be used to understand the structure and function of the visual cortex at the computational level.

Acknowledgments

Thanks to Andrea Haessly and Joseph Sirosh for initial simulations that led to the RF-SLISSOM architecture, and to two anonymous reviewers for help in clarifying the paper. This research was supported in part by the National Science Foundation under grant #IRI-9309273 and by the Texas Higher Education Coordinating Board under grant #ARP-444.

References

[1] H. B. Barlow. Single units and sensation: A neuron doctrine for perceptual psychology? *Perception*, 1:371–394, 1972.

[2] H. B. Barlow. The twelfth Bartlett memorial lecture: The role of single neurons in the psychology of perception. *Quarterly Journal of Experimental Psychology*, 37A:121–145, 1985.

[3] G. G. Blasdel. Orientation selectivity, preference, and continuity in monkey striate cortex. *Journal of Neuroscience*, 12:3139–3161, August 1992.

[4] G. G. Blasdel and G. Salama. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321:579–585, 1986.

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

[6] Y. Choe and R. Miikkulainen. Self-organization and segmentation with laterally connected maps of spiking neurons. In *Workshop on Self-Organizing Maps*, pages 20–31, Espoo, Finland, 1997. Helsinki University of Technology.

[7] Y. Choe and R. Miikkulainen. Self-organization and segmentation with laterally connected spiking neurons. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, pages 1120–1125, San Mateo, CA, 1997. Morgan Kaufmann. In press.

[8] M. B. Dalva and L. C. Katz. Rearrangements of synaptic connections in visual cortex revealed by laser photostimulation. *Science*, 265:255–258, July 1994.

[9] R. Eckhorn, R. Bauer, W. Jordan, M. Kruse, W. Munk, and H. J. Reitboeck. Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics*, 60:121–130, 1988.

[10] R. Eckhorn, H. J. Reitboeck, M. Arndt, and P. Dicke. Feature linking via synchronization among distributed assemblies: Simulations of results from cat visual cortex. *Neural Computation*, 2:293–307, 1990.

[11] D. J. Field. Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America*, 4:2379–2394, 1987.

[12] D. J. Field. What is the goal of sensory coding? *Neural Computation*, 6:559–601, 1994.

[13] D. Fitzpatrick, B. R. Schofield, and J. Strote. Spatial organization and connections of iso-orientation domains in the tree shrew striate cortex. In *Society for Neuroscience Abstracts*, volume 20, page 837, 1994.

[14] C. M. Gray, P. Konig, A. Engel, and W. Singer. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337, 1989.

[15] C. M. Gray and W. Singer. Stimulus specific neuronal oscillations in the cat visual cortex: A cortical functional unit. In *Society of Neuroscience Abstracts*, volume 13, page 404.3, 1987.

[16] A. Grinvald, E. E. Lieke, R. D. Frostig, and R. Hildesheim. 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, 1994.

[17] Y. Hata, T. Tsumoto, H. Sato, K. Hagihara, and H. Tamura. Development of local horizontal interactions in cat visual cortex studied by cross-correlation analysis. *Journal of Neurophysiology*, 69:40–56, January 1993.

[18] J. A. Hirsch and C. D. Gilbert. Synaptic physiology of horizontal connections in the cat's visual cortex. *Journal of Neuroscience*, 11:1800–1809, June 1991.

[19] T. Kohonen. Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69, 1982.

[20] T. Kohonen. *Self-Organization and Associative Memory*. Springer, Berlin; New York, third edition, 1989.

[21] T. Kohonen. Physiological interpretation of the self-organizing map algorithm. *Neural Networks*, 6:895–905, 1993.

[22] R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh. Self-organization, plasticity, and low-level visual phenomena in a laterally connected map model of the primary visual cortex. In R. L. Goldstone, P. G. Schyns, and D. L. Medin, editors, *Perceptual Learning*, volume 36 of *Psychology of Learning and Motivation*, pages 257–308. Academic Press, San Diego, CA, 1997.

[23] K. Obermayer, G. G. Blasdel, and K. J. Schulten. Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Physical Review A*, 45:7568–7589, 1992.

[24] K. Obermayer, H. J. Ritter, and K. J. Schulten. Large-scale simulations of self-organizing neural networks on parallel computers: Application to biological modelling. *Parallel Computing*, 14:381–404, 1990.

[25] K. Obermayer, H. J. Ritter, and K. J. Schulten. A neural network model for the formation of topographic maps in the CNS: Development of receptive fields. In *Proceedings of the International Joint Conference on Neural Networks* (San Diego, CA), Piscataway, NJ, 1990. IEEE.

[26] K. Obermayer, H. J. Ritter, and K. J. Schulten. A principle for the formation of the spatial structure of cortical feature maps. *Proceedings of the National Academy of Sciences, USA*, 87:8345–8349, 1990.

[27] D. Purves and J. W. Lichtman. *Principles of Neural Development*. Sinauer, Sunderland, MA, 1985.

[28] H. Reitboeck, M. Stoecker, and C. Hahn. Object separation in dynamic neural networks. In *Proceedings of the IEEE International Conference on Neural Networks* (San Francisco, CA), volume 2, pages 638–641, 1993.

[29] J. Sirosh. *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, 1995. Technical Report AI95-237.

[30] J. Sirosh and R. Miikkulainen. Cooperative self-organization of afferent and lateral connections in cortical maps. *Biological Cybernetics*, 71:66–78, 1994.

[31] J. Sirosh and R. Miikkulainen. Self-organization and functional role of lateral connections and multisize receptive fields in the primary visual cortex. *Neural Processing Letters*, 3:39–48, 1996.

[32] J. Sirosh and R. Miikkulainen. Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. *Neural Computation*, 9:577–594, 1997.

[33] J. Sirosh, R. Miikkulainen, and J. A. Bednar. Self-organization of orientation maps, lateral connections, and dynamic receptive fields in the primary visual cortex. In J. Sirosh, R. Miikkulainen, and Y. Choe, editors, *Lateral Interactions in the Cortex: Structure and Function*. The UTCS Neural Networks Research Group, Austin, TX, 1996. Electronic book, ISBN 0-9647060-0-8, http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96.

[34] J. Sirosh, R. Miikkulainen, and Y. Choe, editors. *Lateral Interactions in the Cortex: Structure and Function*. The UTCS Neural Networks Research Group, Austin, TX, 1996. Electronic book, ISBN 0-9647060-0-8, http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96.

[35] G. Turrigiano, L. F. Abbott, and E. Marder. Activity-dependent changes in the intrinsic properties of cultured neurons. *Science*, 264:974–977, 1994.

[36] C. von der Malsburg. Synaptic plasticity as basis of brain organization. In J.-P. Changeux and M. Konishi, editors, *The Neural and Molecular Bases of Learning*, pages 411–432. Wiley, New York, 1987.

[37] C. von der Malsburg and J. Buhmann. Sensory segmentation with coupled neural oscillators. *Biological Cybernetics*, 67:233–242, 1992.

[38] D. Wang. Synchronous oscillations based on lateral connections. In J. Sirosh, R. Miikkulainen, and Y. Choe, editors, *Lateral Interactions in the Cortex: Structure and Function*. The UTCS

Neural Networks Research Group, Austin, TX, 1996. Electronic book, ISBN 0-9647060-0-8, http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96.

[39] R. O. Wong. The role of spatio-temporal firing patterns in neuronal development of sensory systems. *Current Opinion in Neurobiology*, 3:595–601, 1993.

[40] R. O. Wong. Synaptic activity and the construction of cortical circuits. *Science*, 274:1133–1138, 1996.