Self-Organization and Segmentation with Laterally Connected Spiking Neurons

Yoonsuck Choe
Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712 USA

Risto Miikkulainen
Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712 USA

Abstract

A self-organizing model of spiking neurons with dynamic thresholds and lateral excitatory and inhibitory connections is presented and tested in the image segmentation task. The model integrates two previously separate lines of research in modeling the visual cortex. Laterally connected self-organizing maps have been used to model how afferent structures and lateral connections could 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 activity. Although these approaches differ in how they model the neuron, they have 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 a 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. [1996b] for an overview). 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. This process can be modeled computationally, showing how structures such as ocular dominance and orientation columns and patterned lateral connections between them form based on input-driven Hebbian learning process (the Laterally Interconnected Synergetically Self-Organizing Map, or LISSOM [Miikkulainen *et al.*, 1997; Sirosh, 1995; Sirosh and Miikkulainen, 1994; 1996; 1997; Sirosh *et al.*, 1996a]).

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 are presented [Gray and Singer, 1987; Eckhorn *et al.*, 1988; Gray *et al.*, 1989]. Several models have been proposed to explain this phenomenon [von der Malsburg, 1987; von der Malsburg and Buhmann, 1992; Eckhorn *et al.*, 1990; Reitboeck *et al.*, 1993; Wang, 1996]. The model of Reitbock et al. [1993] is particularly interesting because of its sophisticated model of the neuron: the synapses are leaky integrators that sum incoming signals over time with exponential decay. A network of such neurons can segment multiple objects in a scene by synchronizing neuronal activity. Spikes of neurons representing the same object are synchronized, and those of neurons representing different objects are desynchronized.

This paper shows how the leaky integrator model of the spiking neuron can be integrated with the LISSOM model of self-organization. The architecture is named Spiking Laterally Interconnected Synergetically Self-Organizing Map, or SLISSOM. SLISSOM (1) forms a topological map from an initially random network through synergetic self-organization and (2) generates synchronized and desynchronized neuronal activity that can be used for segmenting multiple objects in the scene. The results suggest that lateral connections play a central role in both the development and function of the visual cortex.

2 The SLISSOM Architecture

SLISSOM consists of two layers of interconnected neurons: the "retina" and the "cortex" (figure 1*a*). The overall organization of SLISSOM is based on the LISSOM architecture [Miikkulainen *et al.*, 1997; Sirosh, 1995; Sirosh and Miikkulainen, 1994; 1996; 1997; Sirosh *et al.*, 1996a], and the neuron model on the leaky integrator neurons of Eckhorn el al. [1990] and Reitbock et al. [1993]. LISSOM provides a self-organizing principle and the leaky integrator neuron introduces temporal dynamics to the SLISSOM model.

Each cortical neuron receives afferent connections from the input layer and lateral (excitatory and in-

To appear in *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence* **(IJCAI-97, Nagoya, Japan), 1997**

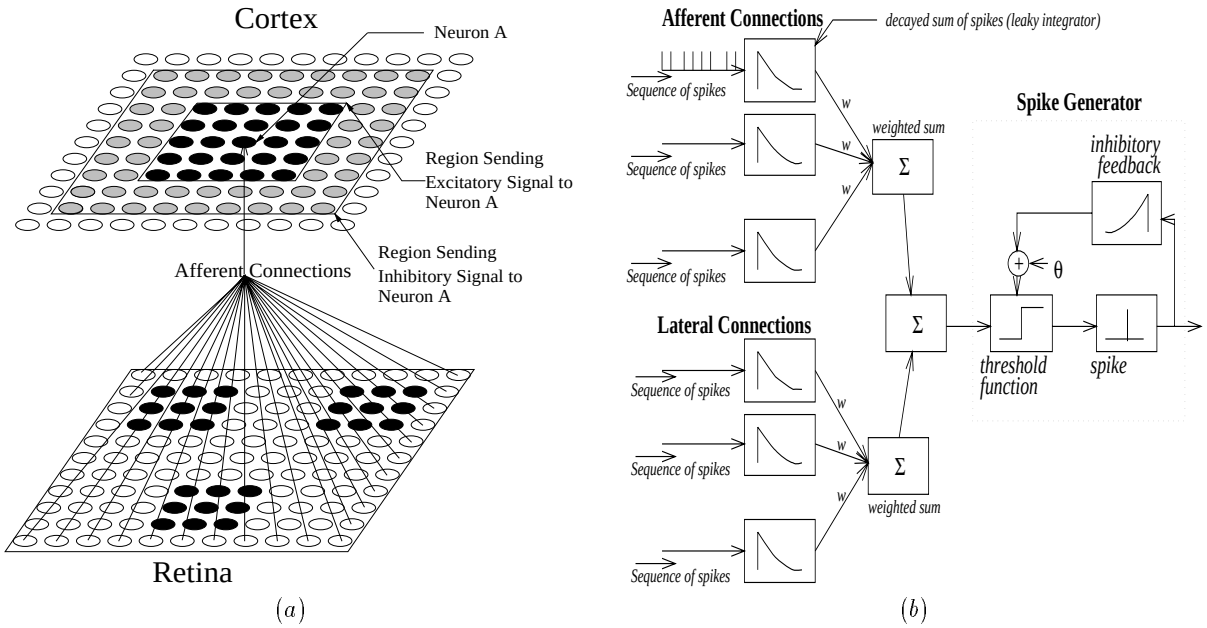

Figure 1: **The SLISSOM Architecture.** (a) The organization of the SLISSOM network. The bottom layer is the retina, and the top layer models the cortical neurons. There are short-range lateral excitatory connections and long-range lateral inhibitory connections between cortical neurons. Each of these neurons receives input from all neurons in the retina. A sample input (consisting of three 3×3 input spots) is shown on the retina. (b) The structure of a single neuron in SLISSOM. Leaky integrators at each synapse perform decayed summation of incoming spikes. The spike generator 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.

hibitory) 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 1b). 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. Active spiking therefore increases the effective threshold, making further spiking less likely and keeping the activation of the system within a reasonable range [Eckhorn *et al.*, 1988; 1990].

The overall organization of the SLISSOM model is shown in figure 1a. The cortical neurons receive input from all retinal neurons. The excitatory lateral connections are short-range (neurons marked black in the cortex), and inhibitory connections are long-range (neurons marked gray). 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 $\frac{1}{e^{\lambda_q}}$, where λ_q is the decay parameter, as they are shifted through the queue. The inhibitory feedback loop in the spike generator (figure 1b) is a similar queue that receives spikes from the spike generator itself, with decay $\frac{1}{e^{\lambda_s}}$.

The input to the network consists of squares of fixed size (figure 1a). Spikes are generated at the active retinal neurons and sent through the afferent connections to the cortical neurons. The net input σ_{ij} 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 contributions:

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

where γ_{a}, γ_{e}, and γ_{i} are the scaling factors for the afferent, excitatory, and inhibitory contributions, ξ_{r_1,r_2} is

[1] This differs from Eckhorn et al. [1990] and Reitbock et al. [1993] 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.

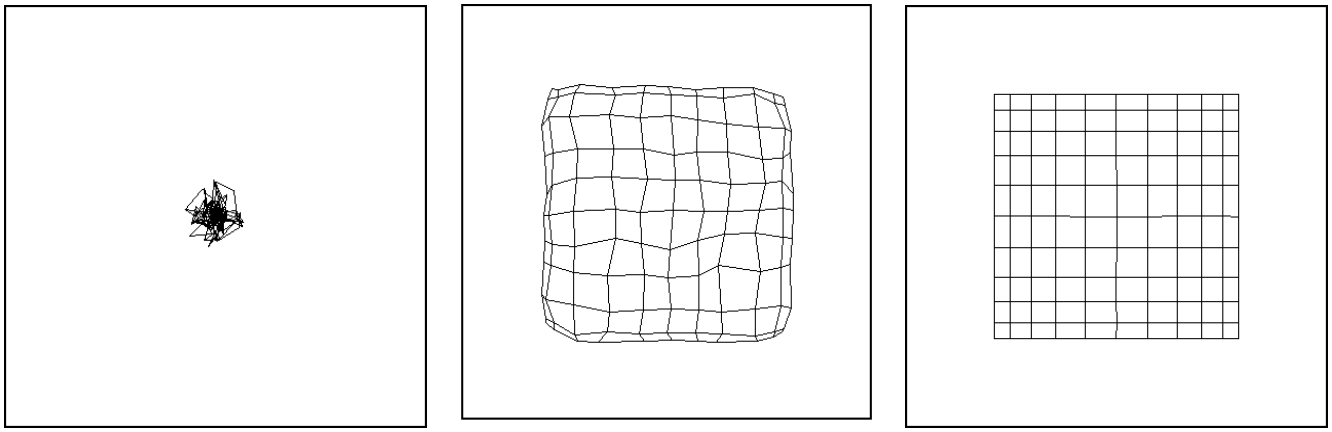

Figure 2: **Self-Organization of the SLISSOM Map.** The nodes in the grid show the centers of gravity of the receptive fields of the cortical neurons. Nodes representing immediate neighbors in the network are connected with a line. (a) The afferent weights are initially randomized and their centers of gravity are about the same. (b) After 5500 iterations, the network forms a well-formed mapping of the input space, comparable to (c), the ideal grid where each node represents a gaussian receptive field located directly below the map unit.

the decayed sum of the incoming queue from the retinal neuron at (r_1, r_2), μ_{ij,r_1r_2} is the corresponding afferent connection weight, $\eta_{kl}(t-1)$ is the decayed sum of the incoming queue from the map neuron at (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 $\sigma_{ij} > \theta + \vartheta_{ij}$, where θ is the base threshold and ϑ_{ij} the output of the spike generator's leaky integrator.

In the standard LISSOM model, the input is kept constant while the cortical response settles to a focused, redundancy-removed activation pattern through the lateral connections. SLISSOM goes through a similar settling process. The input is kept constant 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. Both the afferent and the lateral weights are modified according to the Hebbian principle:

$$w_{ij,mn}(t) = \frac{w_{ij,mn}(t-1) + \alpha V_{ij} X_{mn}}{\mathcal{N}}, \qquad (2)$$

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 (cf. Sirosh and Miikkulainen [1994]). Each neuron receives input from all receptors in the retina, and has excitatory connections with neighboring neurons and inhibitory connections with a larger area of the map.[2] The radius of the lateral excitation is gradually reduced, resulting in fine tuning of the map [Miikkulainen *et al.*, 1997; Sirosh and Miikkulainen, 1997]. The weights are adapted both during self-organization and segmentation.

3 Experiments

The SLISSOM experiment consists of two parts: (1) self-organization, and (2) object segmentation. During self-organization, lateral and afferent connection weights are adapted to form a topological map of the input space. After the network has stabilized, multiple objects (3×3 squares) are presented to the retina. The weights adapted to the input and the network segments the objects by temporally alternating the activity on the map.

The retina and the cortex both consisted of 11×11 units. The afferent weights were initialized to have receptive fields of size 3×3 on the retina, centered right below each neuron, and then 65% noise was added to their values (figures 2a and 3a). The lateral connection weights were randomly initialized within $[0,1]$ (figure 3c). Inhibitory connections covered the whole map, and excitatory connections linked to a square area centered at each neuron (figure 1a), with initial radius of 8, gradually decreasing to 1 in 3,500 iterations. At the same time, the lateral inhibitory learning rate α_{i} gradually increased from 0.001 to 0.1. Slow adaptation in the beginning captures long-term correlations within the inputs, which is necessary for self-organization. Fast adaptation towards

[2] Although most long-range synapses in the cortex are excitatory, they can have inhibitory overall effects through interneurons [Grinvald *et al.*, 1994; Hata *et al.*, 1993; Hirsch and Gilbert, 1991]. The LISSOM model predicts that such long-range inhibition is computationally necessary for self-organization to occur [Sirosh and Miikkulainen, 1997].

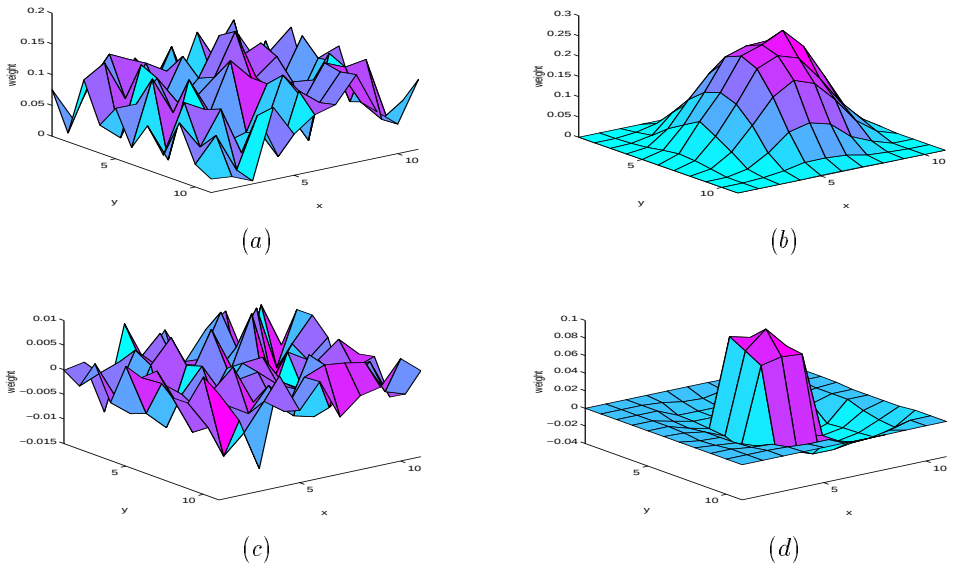

Figure 3: **Connection Weights of the Map Neuron** $(6,6)$**.** (a) Initial afferent weights, (b) final afferent weights, (c) initial combined $(excitatory - inhibitory)$ lateral interaction profile, and (d) final combined lateral interaction profile. The final weights are shown after 5500 self-organization iterations. The x and y axes in (a) and (b) represent the location on the retina, and in (c) and (d), the location on the cortical network.

the end facilitates quick modulation of the activity necessary for segmentation (section 4).

During self-organization, single 3×3 square objects were presented to the network. The retinal neurons representing objects were spiking at each time step, and the settling consisted of 15 cycles of cortical activity update (equation 1). After settling, connection weights were modified according to equation 2, based on the average firing rate over the last 10 cycles. Each such presentation was counted as an iteration. After 5500 iterations, both the afferent and the lateral weights stabilized into smooth profiles. Afferent weights formed smooth gaussian receptive fields most sensitive to input from the retinal neuron right below the map neuron, as shown in figure $3b$. Lateral weights formed smooth Mexican-hat profiles, as shown in figure $3d$. Figure 2 shows the global organization of the map during the process. The final map (figure $2b$) closely resembles the ideal map of the input space (figure $2c$).

Once the SLISSOM network had formed smooth and concentrated receptive fields and lateral interaction profiles, segmentation experiments were conducted on it. Several input spots (again, 3×3 squares) were presented to the retina at the same time. The spots constantly spiked on the retina for 500 time steps. For each spot, a separate 5×5 area on the map responded and the other areas remained silent. The lateral connection weights were adapted at each time step according to equation 2, with $\alpha_i = 0.1$, based on the average firing rate over the last 10 steps.

Segmentation is evident in the total number of spikes generated within each area per time step (i.e. the multi-unit activity, or MUA; figure 4). A high MUA value implies that most neurons in the area are firing together, and a zero value implies that the area is silent. Initially, the three areas corresponding to the three input spots are equally active, but as time goes on, they start to alternate. The spikes within the same area become synchronized (the neurons turn on and off together), and the spikes across the different areas become desynchronized (while one area is active, the other two are silent). Such synchronized and alternating activity indicates that there are three separate objects in the input; in other words, it constitutes a mechanism for binding and segmentation. This result is very robust and works repeatedly for different locations on the retina and for different numbers of objects, as long as the input spots are spatially separate (see section 4).

4 Discussion

Several studies have shown that fast adaptation of synaptic efficacy is necessary for feature binding through temporal coding [von der Malsburg, 1987; Wang, 1996]. Similarly in the experiments with 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

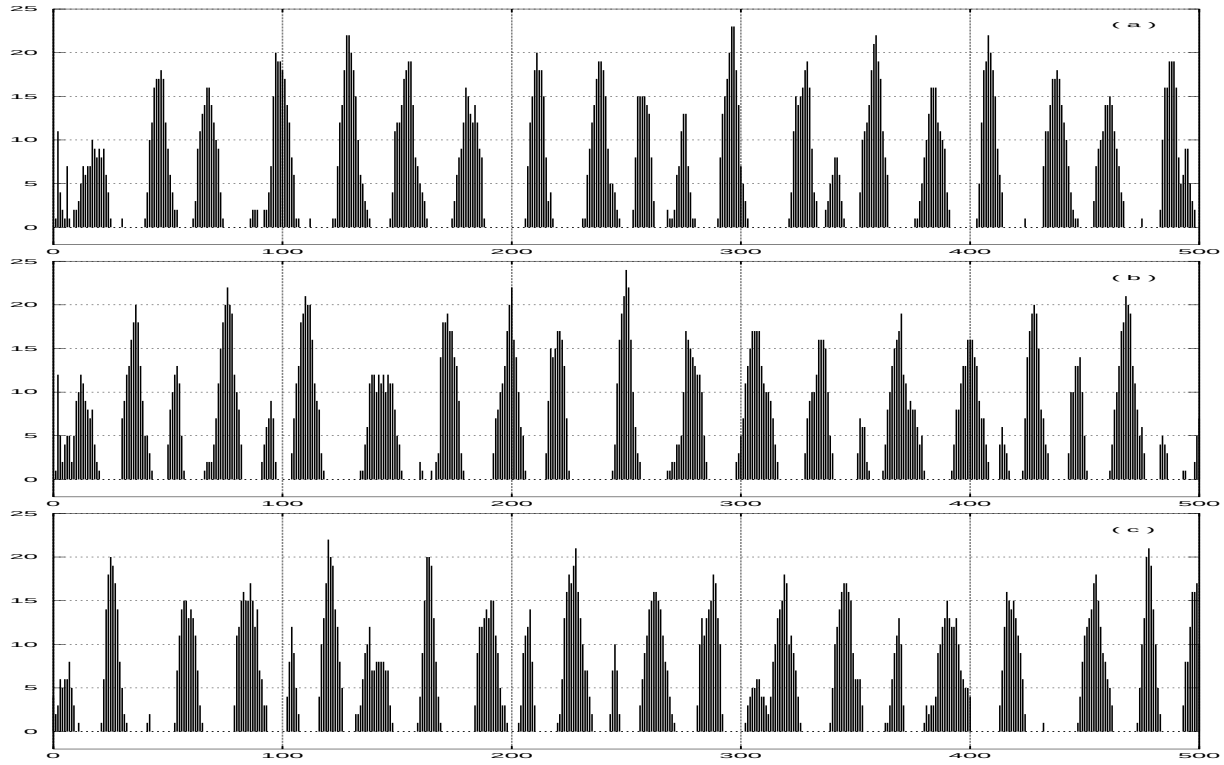

Figure 4: **The Multi-Unit Activities of Areas Responding to Three Different Objects.** The total number of spikes per time step in each of the three 5×5 areas are plotted over 500 time steps. Although initially there is simultaneous activity in all areas, they quickly desynchronize and activation rotates from one area to another.

adaptation [Wang, 1996]. The other is to vary the learning rate of the synapse. It is unknown which approach is more biologically plausible; this question has yet to be settled physiologically. In this work, the learning process starts out with a slow learning rate and gradually the synapses become more plastic. This scheme does not disturb the self-organization since the activity on the map becomes more consistent and predictable as the training goes on, and the need for keeping track of the long-term correlations disappears. 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 MUAs show some overlap even when the input is successfully segmented (figure 4). This is due to the slightly overlapping receptive fields in the model. Gray et al. [1989] 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 edges. The overlap is unavoidable in the current small SLISSOM network, but could be reduced in larger-scale simulations. Such simulations with a large number and variety of objects constitute the most immediate direction of future research. Segmentation in a more detailed self-organized model of the visual cortex, with orientation columns and patterned lateral connections will also be studied, and it may be possible to account for phenomena such as Gestalt effects based on the patterned lateral connections.

5 Conclusion

In this paper, the SLISSOM model of dynamic spiking in a synergetically self-organizing map was presented. Adapting lateral connections were shown to play an essential role in both self-organization and image segmentation, showing how the development and function of the visual cortex could be accounted for by a single unified architecture.

Acknowledgments

Thanks to Andrea Haessly and Joseph Sirosh for initial simulations that led to the SLISSOM architecture. 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

[Eckhorn *et al.*, 1988] 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.

[Eckhorn *et al.*, 1990] 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.

[Gray and Singer, 1987] 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.

[Gray *et al.*, 1989] 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.

[Grinvald *et al.*, 1994] 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.

[Hata *et al.*, 1993] 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.

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

[Miikkulainen *et al.*, 1997] R. Miikkulainen, J. A. Bednar, Y. Choe, and Joseph. 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, *Psychology of Learning and Motivation*, volume 36. Academic Press, San Diego, CA, 1997. In press.

[Reitboeck *et al.*, 1993] 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.

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

[Sirosh and Miikkulainen, 1996] 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.

[Sirosh and Miikkulainen, 1997] J. Sirosh and R. Miikkulainen. Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. *Neural Computation*, 9(3):577–594, 1997. In press.

[Sirosh *et al.*, 1996a] 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.

[Sirosh *et al.*, 1996b] 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.

[Sirosh, 1995] 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.

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

[von der Malsburg, 1987] 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.

[Wang, 1996] 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.