

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Lateral Connections In The Visual Cortex Can Self-Organize Cooperatively With Multisize RFs Just As With Ocular Dominance and Orientation Columns

Permalink

<https://escholarship.org/uc/item/6rd723jp>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 18(0)

Authors

Sirosh, Joseph

Miikkulainen, Risto

Publication Date

1996

Peer reviewed

Lateral Connections In The Visual Cortex Can Self-Organize Cooperatively With Multisize RFs Just As With Ocular Dominance and Orientation Columns

Joseph Sirosh¹ and Risto Miikkulainen

Department of Computer Sciences
The University of Texas at Austin, Austin, TX-78712
sirosh,risto@cs.utexas.edu

Abstract

Cells in the visual cortex are selective not only to ocular dominance and orientation of the input, but also to its size and spatial frequency. The simulations reported in this paper show how size selectivity could develop through Hebbian self-organization, and how receptive fields of different sizes could organize into columns like those for orientation and ocular dominance. The lateral connections in the network self-organize cooperatively and simultaneously with the receptive field sizes, and produce patterns of lateral connectivity that closely follow the receptive field organization. Together with our previous work on ocular dominance and orientation selectivity, these results suggest that a single Hebbian self-organizing process can give rise to all the major receptive field properties in the visual cortex, and also to structured patterns of lateral interactions, some of which have been verified experimentally and others predicted by the model.

Introduction

In their first recordings from the primary visual cortex of the cat, Hubel and Wiesel (1959, 1962) reported that cortical cells were more selective to the width of patterns than were retinal cells. They noted that cortical cells would give no response to a bar covering the whole receptive field (RF), whereas in the retina and the LGN, cells would typically respond to such patterns. Subsequently, detailed studies by Campbell et al. (1969), De Valois et al. (1982) and others showed that cortical cells are narrowly tuned to the spatial frequency of inputs, and had typical bandpass responses, responding only to inputs in a specific frequency range. A continuum of spatial frequencies from low to high were represented in the cortex (Silverman et al. 1989), and cells in each range of spatial frequency were organized into distinct spatial frequency columns (Tootell et al. 1981; Tootell et al. 1988). In essence, cortical cells exhibited an organization of spatial frequency selectivity similar to ocular dominance (OD) and orientation (OR) columns.

Several computational models have been built to demonstrate how other RF properties such as OR preference, OD, and retinotopy can emerge from simple self-organizing processes (e.g. Goodhill 1993; Miller et al. 1989; Obermayer et al.

1992; von der Malsburg 1973). However, to date, only one computational model has included the development of spatial frequency selectivity (Miller 1994). In Miller's model, OR preference and spatial frequency selectivity develop together, and perhaps because of the interactions between these two domains, does not produce a clear columnar organization of spatial frequency selectivity. Although the above models replicate the self-organization of afferent structures quite well, they are based on the simplification that the neuronal response properties are primarily determined by the organization of afferent synapses. Lateral interactions between neurons are approximated by simple mathematical functions (e.g. Gaussians) and assumed to be uniform throughout the network; the structured lateral connectivity of the cortex is not explicitly taken into account. Such models do not explicitly replicate the activity dynamics of the visual cortex, and therefore can make only limited predictions about interactions between receptive fields and cortical function.

Recent experiments have shown that lateral connection patterns closely follow the neuronal response properties (Gilbert and Wiesel 1989; Malach et al. 1993). For example, in the normal visual cortex, long-range lateral connections link areas with similar OR preference (Gilbert and Wiesel 1989). Like neuronal response properties, the connectivity pattern is highly plastic in early development and can be altered by experience (Katz and Callaway 1992). Such patterned lateral connections develop at approximately the same time as the cortical columns (Burkhalter et al. 1993; Katz and Callaway 1992). Together, these observations suggest that the same experience-dependent process drives the development of both neuronal response properties and lateral connectivity.

Previously, we have shown that a single Hebbian self-organizing process can account for the development of patterned lateral connections, afferent receptive fields, topographic maps and OD and OR columns in the cortex (LISSOM, the Laterally Interconnected Synergetically Self-Organizing Map; Sirosh 1995; Sirosh and Miikkulainen 1995a, 1995b, 1996b, 1996a). However, we have not studied the selectivity to different-sized stimuli with LISSOM before, although it is a major component of cortical organization. This article investigates whether the same self-organizing process can give rise to RFs selective to different stimulus sizes. Because size selectivity is closely related to spatial frequency selectivity,

¹Current address: HNC Software Inc., San Diego, CA 92121-3728, and Sloan Center For Theoretical Neurobiology The Salk Institute, La Jolla, CA 92037; sirosh@hnc.com.




Figure 1: The Receptive-Field LISSOM architecture. The afferent and lateral connections of a single neuron in the LISSOM network are shown. The afferents form a local anatomical receptive field on the retina.

such self-organization should account for spatial frequency columns as well.

Several new results are reported in this article. It is shown how afferent RFs of different sizes develop from simple retinal images and organize across the network in a systematic fashion. In addition, lateral connections self-organize cooperatively and simultaneously with the size selectivity properties, producing patterns that follow the receptive field organization. In combination with our previous work, these results suggest that a single unified self-organizing process can give rise to not only all the major receptive field properties in the visual cortex, but also the patterns of lateral interactions.

The Receptive Field LISSOM (RF-LISSOM) model

The LISSOM network is a sheet of interconnected neurons (figure 1). Through afferent connections, each neuron receives input from a “retina”. In addition, each neuron has reciprocal excitatory and inhibitory lateral connections with other neurons. Lateral excitatory connections are short-range, connecting only close neighbors. Lateral inhibitory connections link neurons over both short and long distances, and may even implement full connectivity between neurons in the network.

Neurons receive afferent connections from broad overlapping patches on the retina called anatomical RFs. The $N \times N$ network is projected on to the retina of $R \times R$ receptors, and each neuron is connected to receptors in a square area of side s around the projections. Thus, neurons receive afferents from corresponding regions of the retina. Depending on the location of the projection, the number of afferents to a neuron from the retina could vary from $\frac{1}{2}s \times \frac{1}{2}s$ (at the corners) to $s \times s$ (at the center). Typically, R is much less than N and s is large enough to cover many receptors, resulting in large overlap between receptive fields of nearby neurons.

The input to the model consists of gaussian spots of “light” on the retina:

$$\xi_{a,b} = \exp\left(-\frac{(a-x_i)^2 + (b-y_i)^2}{u^2}\right) \quad (1)$$

where $\xi_{a,b}$ is the activation of receptor (a,b) , u^2 determines the width of the spot, and (x_i, y_i) : $0 \leq x_i, y_i < R$ its center. Without normalization, larger-sized spots would produce stronger activation. Therefore, the retinal activity vector is normalized to constant length. The width u is chosen uniformly randomly in a given range, so that inputs of a variety of sizes are presented to the network.

The external and lateral weights are organized through an unsupervised learning process. At each training step, neurons start out with zero activity. The initial response η_{ij} of neuron (i,j) is based on the scalar product

$$\eta_{ij} = \sigma \left(\sum_{a,b} \xi_{ab} \mu_{ij,ab} \right), \quad (2)$$

where ξ_{ab} is the activation of retinal receptor (a,b) within the anatomical RF of the neuron, $\mu_{ij,ab}$ is the corresponding afferent weight, and σ is a piecewise linear approximation of the familiar sigmoid activation function. The response evolves over time through lateral interaction. At each time step, the neuron combines the above afferent activation $\sum \xi \mu$ with lateral excitation and inhibition:


$$\eta_{ij}(t) = \sigma \left(\sum \xi \mu + \gamma_e \sum_{k,l} E_{ij,kl} \eta_{kl}(t-1) - \gamma_i \sum_{k,l} I_{ij,kl} \eta_{kl}(t-1) \right), \quad (3)$$

where $E_{ij,kl}$ is the excitatory lateral connection weight on the connection from neuron (k,l) to neuron (i,j) , $I_{ij,kl}$ is the inhibitory connection weight, and $\eta_{kl}(t-1)$ is the activity of neuron (k,l) during the previous time step. The constants γ_e and γ_i determine the relative strengths of excitatory and inhibitory lateral interactions. The activity pattern starts out diffuse and spread over a substantial part of the map, and converges iteratively into stable focused patches of activity, or activity bubbles. After the activity has settled, typically in a few iterations of equation 3, the connection weights of each neuron are modified. Both afferent and lateral weights adapt according to the same mechanism: the Hebb rule, normalized so that the sum of the weights is constant:


$$w_{ij,mn}(t + \delta t) = \frac{w_{ij,mn}(t) + \alpha \eta_{ij} X_{mn}}{\sum_{mn} [w_{ij,mn}(t) + \alpha \eta_{ij} X_{mn}]}, \quad (4)$$

where η_{ij} stands for the activity of neuron (i,j) in the final activity bubble, $w_{ij,mn}$ is the afferent or lateral connection weight (μ , E or I), α is the learning rate for each type of connection (α_a for afferent weights, α_E for excitatory, and α_I for inhibitory) and X_{mn} is the presynaptic activity (ξ for afferent, η for lateral).

Both inhibitory and excitatory lateral connections follow the same Hebbian learning process and strengthen by correlated activity. At long-distances, very few neurons have correlated activity and therefore most long-range connections eventually become weak. Such weak connections are eliminated periodically, and through weight normalization, inhibition concentrates in a closer neighborhood of each neuron.



(a) Ocular dominance and lateral connections



(b) Orientation columns and lateral connections

Figure 2: Ocular dominance and orientation columns, and lateral connection patterns. In figure (a), each neuron in a self-organized LISSOM network is labeled with a grey-scale value (*black* \rightarrow *white*) that represents continuously-changing ocular dominance from exclusive left to exclusive right. Small white dots indicate the strongest lateral input connections to the neuron marked with the big white dot, after self-organization. Only the long-range inhibitory connections are shown. The connections of a monocular neuron predominantly link areas of the same ocular dominance. In figure (b), the orientation preference and selectivity of each neuron is represented in grey scale, and the lateral connections of a typical neuron are plotted as in (a). The connections preferentially link similar orientation columns.

The radius of the lateral excitatory interactions starts out large, but as self-organization progresses, it is decreased until it covers only the nearest neighbors (c.f. Self-Organizing Map; Kohonen 1982, 1989). Such pruning of lateral connections produces activity bubbles that are focused and local. As a result, weights change in smaller neighborhoods, and receptive fields become better tuned to local areas of the retina.

Self-Organization of Ocular Dominance and Orientation Columns and Lateral Connections

We have previously used variations of the RF-LISSOM network to model the development of OD and OR columns in the primary visual cortex. Although these phenomena can be modeled with a variety of techniques (Goodhill 1993; Miller et al. 1989; Obermayer et al. 1992; von der Malsburg 1973), RF-LISSOM model is unique in that it also shows how the lateral connections self-organize in the process and what role they play in self-organization and information processing.

In the ocular dominance study (Sirosh and Miikkulainen 1995b; Sirosh and Miikkulainen 1996b), two retinas were connected to the cortical network. Uncorrelated gaussian light spots were used as input, simulating strabismic vision, which is known to result in very pronounced ocular dominance columns (Löwel and Singer 1992). The simulation results are in very good agreement with the biological observations (figure 2a). Sharp and intertwined patterns of ocular dominance form, and neurons are predominantly connected to other neurons with similar selectivity.

In the orientation column experiment (Sirosh and Miik-

kulainen 1995a; Sirosh and Miikkulainen 1996a), elongated gaussians with various orientations and degrees of eccentricity were used as input instead of symmetric gaussians. The cortical network developed a representation of orientation preference and selectivity very similar to those observed in the visual cortex (Blasdel and Salama 1986; Blasdel 1992; figure 2b). Receptive fields of various eccentricity and orientation develop, and neurons are ordered into an orientation map that exhibits linear zones (where orientation preference changes smoothly), pinwheels (singular points around which all orientation preferences appear) and fractures (where preference changes abruptly). The lateral connections again connect neurons that respond to similar inputs. Some of the lateral connection patterns have just recently been discovered, others are predictions of the model.

Self-Organization of Multisize Receptive Fields and Lateral Connections

The hypothesis tested in the present study was whether similar columnar organization and lateral connection patterns would form also when the size of the gaussian light spot was the main dimension of variation in the input. Simulations were carried out on a network of 192×192 neurons, with inputs coming from a 24×24 retina. The anatomical RF size was chosen to be 11×11 , so that there is substantial overlap between the RFs. All the connections were initialized to random weights. A total of 25,000 training steps were used. At each step, a random-size Gaussian spot was presented on the retina as input. The lateral excitatory radius of each neuron started out as

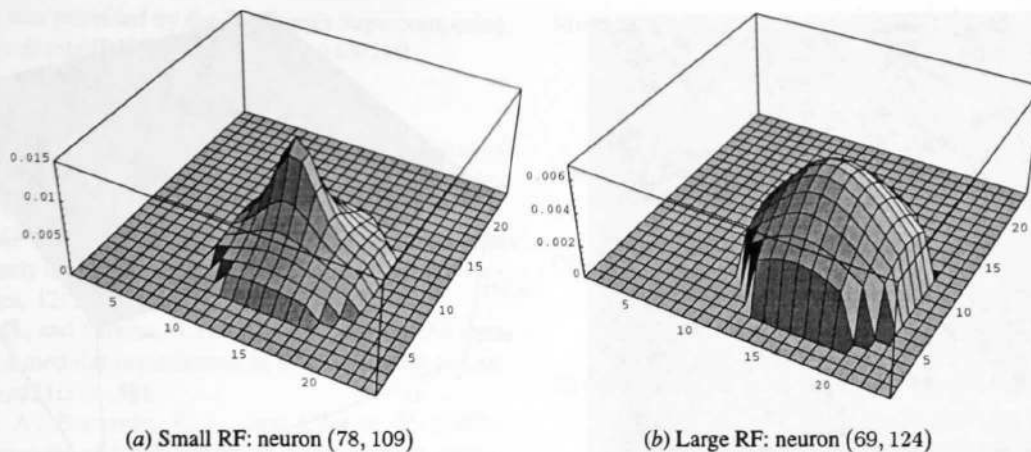


Figure 3: Self-organized receptive fields. The afferent weights of neurons at two different locations in a 192x192 network are shown after self-organization. Initially the weights are completely random, but after self-organization, a smooth hill-shaped weight profile develops. Though the anatomical RFs are the same, the afferent weights are organized into a variety of sizes from narrow, highly peaked receptive fields to large and broad ones.

19, but as training progressed, it was gradually decreased to 1 (as in the Self-Organizing Map algorithm (Kohonen 1989)). The lateral inhibitory connections had a radius of 47, and weak connections were pruned at intervals of 10,000 iterations.

The self-organization of afferents results in smooth, hill-shaped RFs. A variety of RFs of different sizes are produced, some narrow and tuned to small stimuli, others large and most responsive to large stimuli (figure 3). Simultaneously with the RFs, each neuron's lateral connections evolve, and by the Hebbian mechanism, are distributed according to how well the neuron's activity correlates with the activities of the other neurons. Let us examine the nature of such activity correlations. The inputs vary in size from $u = 0.75$ to $u = 8.0$, and are normalized. Therefore, the smallest inputs produce very bright activity in a few receptors. They are also smaller than the size of each anatomical receptive field. Therefore, these inputs predominantly stimulate neurons with small receptive fields and having anatomical RFs in the same position as the spot. Such neurons will have strong activity correlations with other small receptive field neurons, but little correlation with neurons having broader receptive fields². The global organization of size preferences and lateral connections can be visualized by labeling each neuron with a color that indicates the width of its RF, and plotting the patterns of lateral connections on top. As figure 4a shows, the RF organization has the form of connected, intertwined patches, similar to OD columns (e.g. see Sirosh and Miikkulainen 1995b), and the lateral connections of neurons connect to regions of the same size preference. The actual strengths of the connections are shown in figure 4b.


Neurons with larger receptive fields have a slightly differ-

²Note that even small spots produce quite widespread activity in the network, because each retinal receptor connects to a large number of cortical neurons


ent pattern of activity correlations. The larger spots are not localized within the anatomical RF as are the smaller inputs, and extend beyond it. They produce activity over a wider area in the network than the smaller, localized spots. As a result, the inputs that best stimulate larger RF neurons also cause activity in large parts of the network. Therefore, the activity correlations of such neurons are not as strongly determined by size as that of small RF neurons. Therefore, the lateral connections of neurons with larger RFs often link to smaller RF neurons also. In the cortex, neurobiologists have not yet studied how the patterns of lateral connections relate to either size or spatial frequency preferences.

The columnar organization does not develop in small networks. Simulations show that, for a given variance of the stimuli size, the ratio of neurons in the network to receptors in the retina (the magnification factor) has to be greater than a threshold value for a stable columnar organization to appear. Below the threshold, smooth RFs and an ordered topographic map develop, but all the RFs tend to have the same size, corresponding to the average width of the input stimulus. Above the threshold, symmetry breaking occurs, producing a variety of RF sizes. Such symmetry breaking is similar to that of the Self-Organizing Map (Kohonen 1982, 1989), where an input feature is represented in the network only if its variance is greater than a threshold proportional to the magnification factor (Obermayer et al. 1992).

It is not known whether the long-range lateral connections in the cortex are organized according to size or spatial frequency selectivity. So far, the lateral connection patterns have only been studied in relation to the organization of OD and OR preference (Malach et al. 1993; Löwel and Singer 1992; Gilbert and Wiesel 1989). However, considerable psychophysical and neurobiological evidence indicates selective lateral interactions between neurons tuned to different spatial frequencies (De Valois and Tootell 1983; Bauman and Bonds



(a) Columns and lateral connections



(b) Lateral inhibitory weights

Figure 4: Size selective columns, and lateral connection patterns. In figure (a), each neuron in the network is labeled with a grey-scale value (*black* → *white*) that represents continuously-changing size preference from small values to large values. Small white dots indicate the lateral input connections to the neuron marked with the big white dot. The size preferences are organized systematically across the network into connected, intertwined patches, and the strongest lateral connections predominantly link areas of the same size selectivity. Figure (b) shows the weights of the lateral connections plotted in (a). The connection strengths represent the activity correlations of the neuron with the other neurons in the network. The columnar organization of the RFs is reflected in the weights. The connections also are strongest in the immediate vicinity of the neuron (at center) and become weaker with distance. The large areas of zero weights stand for the connections that have been pruned away during self-organization.

1991). As in the RF-LISSOM model, these interactions are also known to be largely inhibitory (De Valois and Tootell 1983; Vidyasagar and Mueller 1994). The model suggests that the long-range lateral connections could be the anatomical substrate for inhibition between spatial frequency channels. The model further predicts that the patterns of lateral connections in the cortex would be influenced not only by OD and OR preference, but also by selectivity to spatial frequency.

Discussion

Combined with our previous work on OD and OR maps and lateral connections, the new results suggest that a single Hebbian mechanism produces the receptive fields and lateral interactions in the primary visual cortex. It also makes several predictions on the lateral connectivity patterns in the cortex. However, it is important to note that the units and connections in the RF-LISSOM model do not correspond one-to-one to neurons and synapses. Instead, each unit should be seen as a vertical column of neurons, and the connections stand for the interactions between these columns. One important prediction of the model is that long-range lateral interactions are inhibitory: This is computationally necessary for the self-organization of receptive fields to occur. However, it doesn't mean that the synapses on long-range connections necessarily have to be inhibitory, as long as their overall effect on the column is inhibitory.

Perhaps most significantly, the RF-LISSOM model suggests a computational role for self-organized structures in the

primary visual cortex. According to the model, two different computations are performed during sensory processing: First, the inputs are projected onto the principal feature dimensions represented by the afferent receptive field structure. Then, the redundancies are filtered out by the inhibitory lateral interactions. The result is an efficient, redundancy-reduced sparse coding of the visual input which is then passed on to higher processing levels. This prediction can probably be verified experimentally by using information theory to analyze the optical images of cortical activity patterns produced in response to simple retinal images. If confirmed, it would constitute a major step in understanding the function of the observed primary visual cortex structures.

Conclusion

The RF-LISSOM model shows how a columnar organization of multisized receptive fields can develop and how lateral connection patterns follow this organization. Combined with our previous work, these results show how a single local and unsupervised self-organizing process can be responsible for the development of both the afferent and lateral connection structures in the primary visual cortex. The model suggests that afferent receptive fields develop a sparse coding of the visual input, and that recurrent lateral interactions eliminate redundancies in cortical activity patterns.

Acknowledgments

This research was supported in part by National Science Foundation under grant #IRI-9309273. Computer time for the

simulations was provided by the Pittsburgh Supercomputing Center under grants IRI930005P and TRA940029P.

References

- Bauman, L., and Bonds, A. (1991). Inhibitory refinement of spatial frequency selectivity in single cells of the cat striate cortex. *Vision Research*, 31(6):933–944.
- Blasdel, G. G. (1992). Orientation selectivity, preference, and continuity in monkey striate cortex. *The 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.
- Burkhalter, A., Bernardo, K. L., and Charles, V. (1993). Development of local circuits in human visual cortex. *Journal of Neuroscience*, 13:1916–1931.
- Campbell, F., Cooper, G., and Enroth-Cugell, C. (1969). The spatial selectivity of the visual cells of the cat. *Journal of Physiology (London)*, 203:223–235.
- De Valois, K. K., and Tootell, R. B. H. (1983). Spatial-frequency-specific inhibition in cat striate cortex cells. *Journal of Physiology (London)*, 336:359–376.
- De Valois, R. L., Albrecht, D. G., and Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, 22:545–559.
- Gilbert, C. D., and Wiesel, T. N. (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *Journal of Neuroscience*, 9:2432–2442.
- Goodhill, G. (1993). Topography and ocular dominance: a model exploring positive correlations. *Biological Cybernetics*, 69:109–118.
- Hubel, D. H., and Wiesel, T. N. (1959). Receptive fields of single neurons in the cat's striate cortex. *Journal of Physiology*, 148:574–591.
- Hubel, D. H., and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology (London)*, 160:106–154.
- 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. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69.
- Kohonen, T. (1989). *Self-Organization and Associative Memory*. Berlin; Heidelberg; New York: Springer. Third edition.
- Löwel, S., and Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science*, 255:209–212.
- Malach, R., Amir, Y., Harel, M., and Grinvald, A. (1993). Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in the primate striate cortex. *Proceedings of the National Academy of Sciences, USA*, 90:10469–10473.
- 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.
- Miller, K. D., Keller, J. B., and Stryker, M. P. (1989). Ocular dominance column development: Analysis and simulation. *Science*, 245:605–615.
- Obermayer, K., Blasdel, G. G., and Schulten, K. J. (1992). Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Physical Review A*, 45:7568–7589.
- Silverman, M. S., Grosz, D. H., De Valois, R. L., and Elfar, S. D. (1989). Spatial frequency organization in primate striate cortex. *Proceedings of the National Academy of Sciences, USA*, 86:711–715.
- 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.
- Sirosh, J., and Miikkulainen, R. (1995a). Cooperative self-organization of orientation maps and lateral connections in the visual cortex. In *Society for Neuroscience Abstracts*, vol. 21, part 3, page 1751.
- Sirosh, J., and Miikkulainen, R. (1995b). Ocular dominance and patterned lateral connections in a self-organizing model of the primary visual cortex. In Tesauro, G., Touretzky, D. S., and Leen, T. K., editors, *Advances in Neural Information Processing Systems 7*. Cambridge, MA: MIT Press.
- Sirosh, J., and Miikkulainen, R. (1996a). Self-organization of orientation maps, lateral connections, and dynamic receptive fields in the primary visual cortex. In Sirosh, J., Miikkulainen, R., and Choe, Y., editors, *Lateral Interactions in the Cortex: Structure and Function*. Electronic book, http://www.cs.utexas.edu/users/nn/lateral_interactions_book/cover.html. ISBN 0-9647060-0-8.
- Sirosh, J., and Miikkulainen, R. (1996b). Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. *Neural Computation*. In press.
- Tootell, R. B., Silverman, M. S., and De Valois, R. L. (1981). Spatial frequency columns in primary visual cortex. *Science*, 214:813–815.
- Tootell, R. B., Silverman, M. S., Hamilton, S. L., Switkes, E., and De Valois, R. L. (1988). Functional anatomy of macaque striate cortex. v. spatial frequency. *Journal of Neuroscience*, 8:1610–1624.
- Vidyasagar, T. R., and Mueller, A. (1994). Function of GABA inhibition in specifying spatial frequency and orientation selectivities in cat striate cortex. *Experimental Brain Research*, 98(1):31–38.
- von der Malsburg, C. (1973). Self-organization of orientation-sensitive cells in the striate cortex. *Kybernetik*, 15:85–100.