Self-Organizing Feature Maps with Lateral Connections: Modeling Ocular Dominance

Joseph Sirosh and Risto Miikkulainen
Department of Computer Sciences
University of Texas at Austin, Austin, Texas-78712, USA.
sirosh,risto@cs.utexas.edu

The neocortex is the largest part of the mammalian brain, and appears to be the least genetically determined. Much of its structure and connectivity depends on electrical activity during development. Because various neocortical areas are very similar in anatomical structure and exhibit similar developmental phenomena, it has been suggested that a common organizing mechanism underlies their ontogeny [18]. By modeling the development of a well-understood substructure such as the primary visual cortex, it might be possible to elucidate this common mechanism.

The primary visual cortex, like many other regions of the neocortex, is a topographic map, and is organized such that adjacent neurons respond to adjacent regions of the retina. This retinotopic map, as well as finer structures within it such as ocular dominance columns, forms by the self-organization of the afferent (input) connections to the cortex. The self-organizing process is driven by external input [4, 20, 19], and appears to be based on correlated (i.e. cooccurring) neuronal activity and the resulting cooperation and competition between neurons [19, 18].

In addition to the afferent connections, the neocortex contains a dense network of long-range lateral connections parallel to the cortical surface [3, 15]. These connections are reciprocal and are believed to mediate the cooperation and competition. They grow exuberantly after birth and reach their full extent in a short period. During subsequent development, they get automatically pruned into well-defined clusters [2, 5]. Pruning happens at the same time as the afferent connections organize into topographic maps. The final clustered distribution corresponds closely to the distribution of afferent connections in the map. For example, in the mature visual cortex, lateral connections primarily run between areas with similar afferent connection structure, such as iso-orientation columns [3]. The structural correspondence and their simultaneous development indicate that the ontogeny of lateral and afferent connections are interdependent.

The development of lateral connections, like that of afferent connections, depends on cortical activity caused by external input. Three observations support this notion: (1) When the primary visual cortex (of the cat) is deprived of visual input during early development, the lateral connectivity remains crude and unrefined [1]. (2) The pattern of lateral connection clusters can be altered by changing the input to the developing cortex. The resulting patterns reflect the correlations in the input [8]. (3) In the mouse somatosensory barrel cortex, sensory deprivation (by sectioning the input nerve) causes drastic decreases in the extent and density of lateral connections [9]. These three observations suggest that the lateral connection structure is not defined genetically, but is acquired.

Lateral connections do not just modulate cortical activity – their development is essential for the self-organization of afferent connections as well. Every neuron receives a large number of lateral connections. Although each individual connection is weak, their total effect on neural activity can be substantial [3], and thereby affect the development of afferent connections. Changes in afferent connections then change the activity patterns on the cortex, which in turn influence the development of lateral connections. The development of both sets of connections thus appears to proceed synergetically and simultaneously, eventually evolving to a state of equilibrium in the adult animal.

The lateral connection structure plays a significant role also in cortical function: (1) by integrating information over large parts of the cortex, lateral connections mediate context-dependent processing of input

To appear in *Proceedings of the Connectionist Models Summer School*, (Boulder, Colorado), June 1993

stimuli [3]; (2) lateral connections may also mediate the synchronization of activity over long distances of the cortex, and thereby help form dynamic representations of coherent input areas [16]; (3) by learning correlations in input during development, they can potentially form long-term representations of input regularities such as gestalt rules [19]; and (4) by combining such representations with input activity, they may also perform feature grouping and segmentation during perception [19]. Therefore, to understand cortical function one must discover not only the structure and function of the afferent connections but also that of the lateral connections. By modeling how lateral connections develop along with the other aspects of cortical structure, a more fundamental understanding of neocortical processes can be achieved.

Several models have been proposed to explain various aspects of afferent connection structure, such as retinotopy, ocular dominance and orientation preference in the visual cortex [20, 12, 11]. The simultaneous development of these three properties has recently been modeled by the Self-Organizing Map (SOM, [6]) algorithm [14, 13]. These models usually assume that the strength of lateral interaction falls off with distance as a Gaussian function, and is uniform throughout the network. The models do not address the development of lateral connections and how they affect the self-organization of afferent connections.

In this paper, a computational model for the synergetic development of the afferent and lateral connections in cortical feature maps is presented. The model is called LISSOM, for Laterally Interconnected Synergetically Self-Organizing Map. LISSOM is based on a simulated network of neurons with afferent connections from the external world and reciprocal lateral connections. Connections adapt based on correlated activity between neurons. The result is a self-organized structure with (1) afferent connection weights that form a map of the input space and (2) lateral connections that store long-term correlations in neuronal activity.

Although LISSOM can be utilized as an abstract self-organizing algorithm, its most promising application is in modeling the development of the neocortex. This paper addresses how ocular dominance and lateral connections develop simultaneously in the visual cortex. The results of these experiments hold promise for substantial insights into cortical development.

The LISSOM Model

Each neuron in a LISSOM network (figure 1) has a set of afferent input connections (from the external input to the map) and a set of lateral input connections (from the other neurons in the map). Each neuron develops an initial response as a weighted sum of the activation in its afferent input connections. The lateral interactions between neurons then focus the initial activation pattern into a localized response on the map. After the pattern has stabilized, all connection weights are modified. As the self-organization progresses, the neurons grow more nonlinear and weak connections die off. Below, these general structures and mechanisms of LISSOM are described in detail, and illustrated in forming a map of a uniform distribution on a square (figures 2, 3 and 4).

The LISSOM network is a sheet of interconnected neurons (figure 1). Through the excitatory afferent connections, every neuron receives the same vector of external input values. In addition, each neuron has reciprocal excitatory and inhibitory lateral connections with other neurons. Lateral excitatory connections are short-range, connecting only close neighbors in the map. Lateral inhibitory connections run for long distances, and may even implement full connectivity between neurons in the map. Each connection has a characteristic strength (or weight), which may be any value between zero and a prescribed limit that depends on the synaptic resources of the neuron.

Input vectors to the network are normalized so that vectors with large norms do not dominate the self-organizing process. In the normalization, the original n-dimensional input distribution is mapped on the surface of an $(n+1)$-dimensional unit hypersphere [10]. The $(n+1)$th-dimension becomes the radius of the hypersphere, and the original set of n dimensions become angles specifying the input point on the surface of the hypersphere. For example, in forming a 2D map as in figure 2, the square area was laid on the surface of a 3D sphere of radius 1.0. In effect, inputs from this area are still 2-dimensional in spherical coordinates because the radius is constant. Each spherical input vector $(x_1, x_2, 1)$, was then transformed into a 3-dimensional cartesian vector $x = (\xi_1, \xi_2, \xi_3)$:

$$\begin{cases} \xi_1 = 1 \cdot \cos(x_1)\cos(x_2), \\ \xi_2 = 1 \cdot \sin(x_1)\cos(x_2), \\ \xi_3 = 1 \cdot \sin(x_2). \end{cases} \tag{1}$$

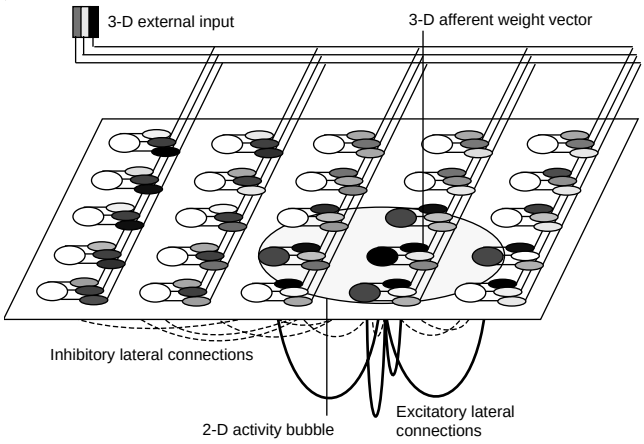

Figure 1: **The LISSOM architecture.** Each neuron receives the same afferent input vector and computes an initial response based on a measure of similarity of the input to the corresponding weight vector. The responses then repeatedly propagate through the lateral connections (only a few connections of the most strongly active unit are shown) and evolve into an activity "bubble". After the activity stabilizes, weights of the active neurons are adapted.

The external and lateral weights are organized through a purely unsupervised learning process. Input items are randomly drawn from the input distribution and presented to the network one at a time. At each training step, the neurons in the network start out with zero activity. The initial response of each neuron η_{ij} in the map is based on the scalar product:

$$\eta_{ij} = \sigma\left(\sum_h \xi_h \mu_{ij,h}\right), \tag{2}$$

where ξ_h are the inputs to the network and $\mu_{ij,h}$ are the corresponding afferent weights. The afferent weights $\mu_{ij,h}$ can be positive or negative. The function σ is a piecewise linear approximation to the sigmoid activation function:

$$\sigma(x) = \begin{cases} 0 & x \leq \delta \\ (x-\delta)/(\beta-\delta) & \delta < x < \beta \\ 1 & x \geq \beta \end{cases}, \tag{3}$$

where δ and β are the lower and upper thresholds. The sigmoid introduces a nonlinearity into the response, and makes the neuron selective to a small range of input vectors that are close to the afferent weight vector.

The response evolves over time through the lateral interaction. At each time step, the neuron combines external activation with lateral excitation and inhibition:

$$\eta_{ij}(t) = \sigma\left(\sum_h \xi_h \mu_{ij,h} + \gamma_e \sum_{k,l} E_{ij,kl}\eta_{kl}(t-1) - \gamma_i \sum_{k,l} I_{ij,kl}\eta_{kl}(t-1)\right), \tag{4}$$

where $E_{ij,kl}$ is the excitatory lateral connection weight on the connection from unit (k,l) to unit (i,j), $I_{ij,kl}$ is the inhibitory connection weight, and $\eta_{kl}(t-1)$ is the activity of unit (k,l) during the previous time step. The lateral connection weights are all positive. The constants γ_e and γ_i are scaling factors on the excitatory and inhibitory weights and determine the strength of lateral interaction. The activity pattern starts out diffuse and spread over a substantial part of the map, and converges iteratively into a stable focused patch of activity, or activity bubble [17, 10]. After the activity has settled, typically in a few iterations, the connection weights of each neuron are modified.

The lateral weights are modified by the Hebb rule, but keeping the sum of the weights constant:

$$\gamma_{ij,kl}(t+\delta t) = \frac{\gamma_{ij,kl}(t) + \alpha_L \eta_{ij}\eta_{kl}}{\sum_{kl}\left[\gamma_{ij,kl}(t) + \alpha_L \eta_{ij}\eta_{kl}\right]}, \tag{5}$$

where η_{ij} stands for the activity of the unit (i,j) in the settled activity bubble, the γs are the lateral interaction weights ($E_{ij,kl}$ or $I_{ij,kl}$) and α_L is the learning rate for the lateral interaction (α_E for excitatory

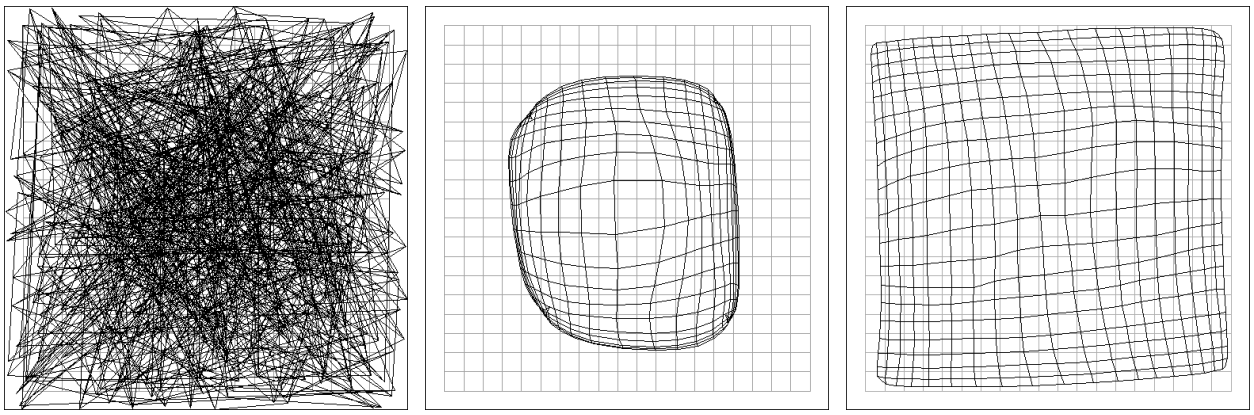

(a) 0 inputs: Random (b) 5,000: Start of pruning (c) 30,000: Final

Figure 2: **Self-organization of the afferent input weights in a 2D map.** The weight vector of each neuron in the 20×20 map is transformed back to the original spherical coordinates and plotted as a point in the input space (represented by the square). Each weight vector is connected to those of the four immediate neighbors by a line. The resulting dark grid depicts the topological organization of the map. The dotted grid shows the best possible approximation of the input space by a 20×20 network. The afferent weight vectors are initially uniformly distributed on the input square (a). Input vectors are randomly drawn from the input area and presented to the network, and the weights are modified according to equations 5 and 6. As the self-organization progresses, the network unfolds and the weight vectors spread out to form a regular topological map of the input space. Before connection pruning, the map is stable, but contracted into the center of the input distribution (b). After connections begin to die off, the map expands and eventually covers the entire input space (c).

weights and α_I for inhibitory). The larger the product of the pre- and post-synaptic activity $\eta_{ij}\eta_{kl}$, the larger the weight change.

The external input weights are modified according to the normalized Hebbian rule:

$$\mu_{ij,h}(t+\delta t) = \frac{\mu_{ij,h}(t) + \alpha\eta_{ij}\xi_h}{\left\{\sum_h \left[\mu_{ij,h}(t) + \alpha\eta_{ij}\xi_h\right]^2\right\}^{1/2}}, \tag{6}$$

which is otherwise similar to (5), but maintains the sum of squares of external weights constant.

The above processes of lateral interaction and weight adaptation are sufficient to form ordered afferent and lateral input connections (figure 2b). However, in order to develop a good coverage of the input space, it is necessary that these processes gradually become more focused and local. As self-organization proceeds in LISSOM, neuronal responses grow more nonlinear and weak lateral connections die off, which results in more focused activity bubbles and weight changes.

The lower threshold of the sigmoidal activation function specifies the minimum input required for a neuron to produce output, and its slope determines how much the neuron amplifies changes in the input. Therefore, neurons can be made more selective and sensitive by increasing the lower threshold, δ, and decreasing the upper threshold, β, resulting in smaller and more refined activity bubbles. In LISSOM, δ is increased and β decreased proportional to the activity of the neuron at each input presentation, up to prescribed maximum and minimum limits:

$$\begin{aligned} \delta_{ij}(t+1) &= \min(\delta_{ij}(t) + \alpha_\delta\eta_{ij}, \delta_{\max}), \\ \beta_{ij}(t+1) &= \max(\beta_{ij}(t) - \alpha_\beta\eta_{ij}, \beta_{\min}). \end{aligned} \tag{7}$$

As a result, the neurons grow nonlinear faster at those parts of the map that see more activity. Such modification automatically takes into account the level of organization around the neuron, and results in regular final maps (figure 2c).

Once the map has organized partially, most of the long-range lateral connections join areas that are almost never active simultaneously. Their weights become small, and they can be pruned without disrupting self-organization. Most long-range inhibitory connections are eliminated this way (figure 3 and 4). Since the total synaptic weight is kept constant, inhibition concentrates in the immediate neighborhood of the

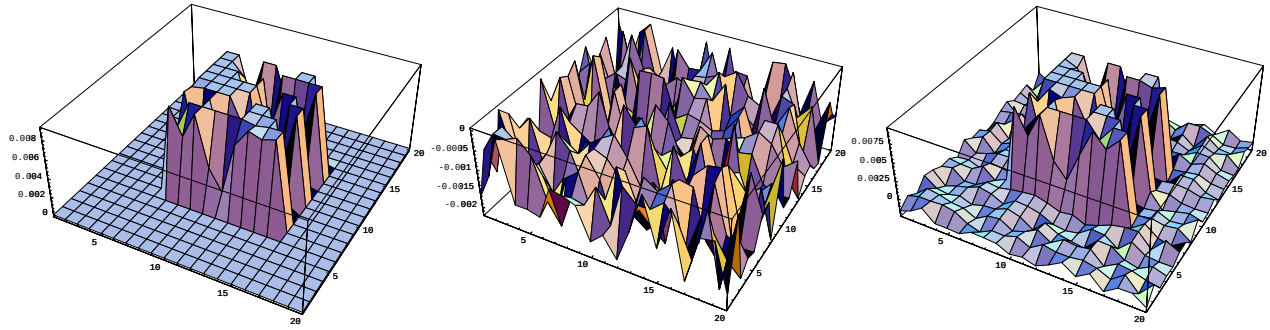

(a) Excitation (b) Inhibition (c) Combined interaction

Figure 3: **Initial lateral interaction in the 2D map.** The lateral excitation and inhibition weights and the combined interaction profile are plotted for the neuron at position $(10, 10)$ in the 20×20 map. The excitation weights (a) are initially randomly distributed within a radius $d = 4$, and zero outside. The inhibition weights (b) are randomly distributed within $d' = 12$, and for neuron $(10, 10)$, cover the entire map. The combined interaction (c) is the sum of the excitatory and inhibitory weights and illustrates the total effect of the lateral connections.

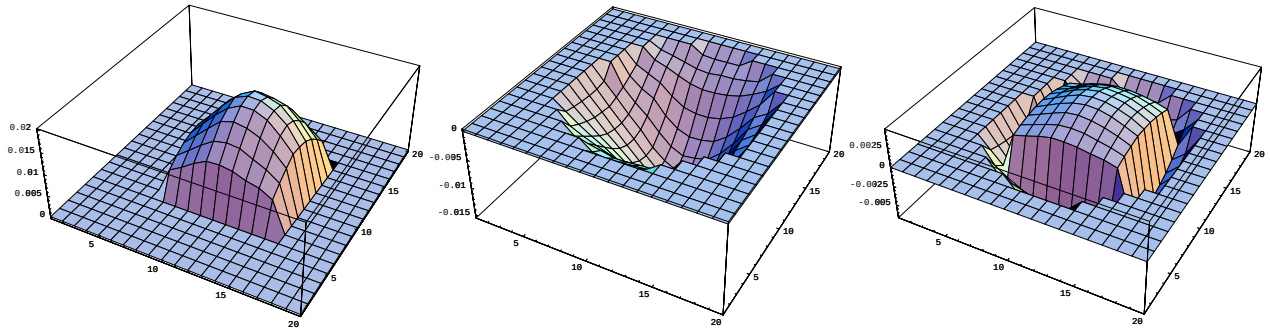

(a) Excitation (b) Inhibition (c) Combined interaction

Figure 4: **Final lateral interaction in the 2D map.** The long-range connections have died off (indicated by zero strength in the figure), resulting in more concentrated and deeper inhibition around the neuron. The combined interaction has a "Mexican hat" shape. Such interaction resulted in smaller activity bubbles, allowing the map to expand as in figure 2c.

neuron. The short-range excitatory connections join neurons that are often part of the same bubble. They have relatively large weights and survive.

Before connection death, activity bubbles are large and weights change in large neighborhoods. Each adaptation makes neighboring weight vectors more parallel, which means that the map contracts to the center of the input space (figure 2b). However, as connections die off, the lateral interaction tends to form smaller, higher-contrast activity bubbles. Weights change in smaller neighborhoods, which allows the map to expand and become a better approximation of the input space (figure 2c).

The afferent connections in LISSOM organize very much like in Kohonen's Self-Organizing Map algorithm. However, in the SOM process, the maximally responding unit is chosen through global supervision, and adaptation neighborhoods are reduced according to a predetermined schedule. In contrast, the LISSOM process is based on purely local rules and the network self-organizes completely without global supervision. Even the shape of the lateral interaction is automatically extracted from the statistical properties of the external input. The self-organization thus "bootstraps" by using external input information to establish the necessary cooperative and competitive interactions.

Modeling Ocular Dominance

The LISSOM process models cortical development at a new level, namely that of explicit activity-based cooperation and competition and developing lateral connections. At this level, it is possible to investigate

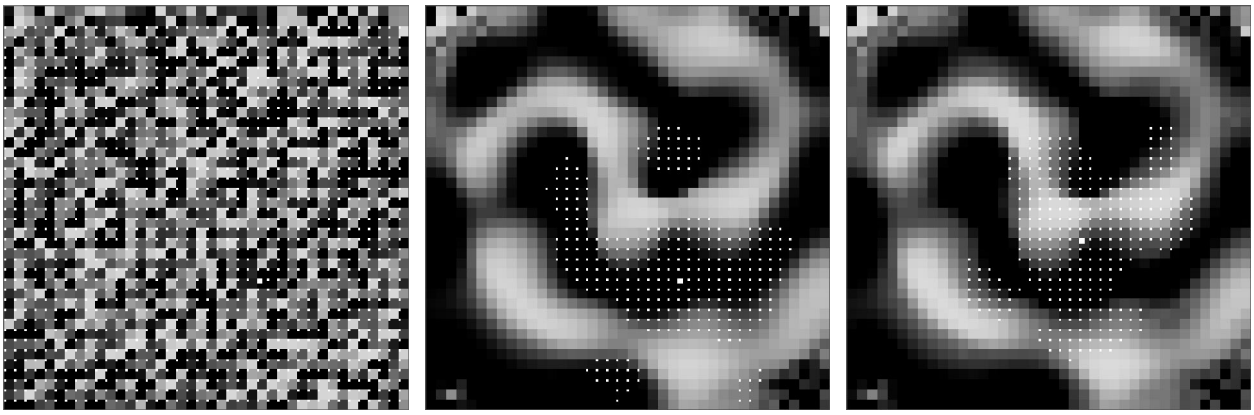

(a) Initial: Random ocular dominance

(b) Final: Mature ocular dominance & connections

(c) Final: Connections of a binocular neuron

Figure 5: **Emergence of ocular dominance and corresponding lateral connection patterns.** In this experiment, the primary visual cortex was modeled by a 40×40 network of neurons. Each neuron is labeled with a grey-scale value corresponding to its ocular dominance. Brightness varying from *black* to *white* represents continuously changing eye preference from exclusive left through binocular to exclusive right. Small white dots indicate the surviving lateral input connections to the neuron marked with a big white dot. (a) The initial state of the network. Neurons have randomly distributed ocular dominance and random-weight lateral connections covering a wide area. (b) After the critical period, weak lateral inhibitory connections die away and the surviving connections predominantly link areas of the same ocular dominance. As simulation proceeds, both the ocular dominance patterns and the lateral connection patterns refine and eventually stabilize. (c) The lateral connections of a binocular neuron come from both eye regions. Note that in (b) and (c), neurons that are physical neighbors do not necessarily influence each other because of the nontopographically distributed lateral connections.

how lateral interaction affects input activity in the cortex as well. As a first step, a model for the development of ocular dominance columns and lateral connections in the primary visual cortex was built.

The inputs to the LISSOM network consisted of three values (x,y,z), (x,y) representing the retinotopic position of the current visual input and $z \in [-Z, +Z]$ representing the ocular dominance. The values $-Z$ and $+Z$ indicate the extremes of left and right dominance, and $z = 0$ that the input comes equally from both eyes. Each neuron had three afferent weights corresponding to the three input values. Initially, the afferent weights were set up with a random ocular dominance $z \in [-Z, +Z]$ and with roughly retinotopic x and y coordinates within $[0, 1]$. The network was trained using input vectors with uniformly distributed random values for x, y and z within these same ranges.

Figure 5 demonstrates the evolution of ocular dominance in the self-organizing process. At the start of the simulation, each input initially results in a rough pattern of activity, which focuses into a single localized activity bubble. As the self-organization progresses, the initial activity patterns begin to split into multiple bubbles. This happens because the map self-organizes like a crumpled sheet in the (original) 3D input space, and cannot completely cover it. Therefore, several areas of the network respond equally well to the inputs falling between folds. Through lateral excitation and inhibition, each of these areas forms a localized bubble. These areas are likely to have similar ocular dominance values. As a result, activity is highly correlated between similar ocular dominance patches and less correlated between those with different ocular dominance. Lateral connections learn to represent such correlations. As afferent connections change and ocular dominance patterns refine, weak lateral connections die away and the surviving connections become clustered around the ocular dominance areas for the same eye. The process continues until a stable equilibrium is reached. The final network has (1) retinotopic afferent weights, (2) well-defined ocular dominance columns, and (3) lateral connections that link areas with the same ocular dominance.

Very similar development was observed by Lowel and Singer [8] in cats. They found that if a newborn kitten is raised with divergent squint-eyes (strabismus), which decorrelates the visual input to the cortex from the two eyes, lateral connections develop preferentially between ocular dominance columns of the same eye. However, binocular neurons at the border of two ocular dominance columns had connections to neurons in both the right and left eye areas. As can be seen from figure 5b and 5c, the LISSOM simulation results

match both observations.

Even based on this simple model of ocular dominance formation, interesting functional predictions can be made. For example, divergent squinters cannot achieve binocular vision even if the visual stimulus is completely binocular [7, 21]. Such impairment is puzzling because most squinting people still have some neurons responsive to input from both eyes, and these neurons should be able to perform binocular integration. The LISSOM model suggests an explanation: If only a few binocular neurons exist, the lateral interactions may override the binocular response. Activity bubbles that initially form in binocular areas tend to break up and move into nearby monocular areas. The resulting activity patterns are disjoint and spatially separated, and therefore cannot be recognized as a single and coherent binocular percept.

Standard feature maps represent the topology of the input space by the network's grid-like layout, so that units close by in the map represent nearby input vectors. In the LISSOM model, the neighborhood relations are mediated by the lateral connections, and the model is not limited to strictly 2-dimensional topology. Units that are far apart on the map grid may still belong to the same neighborhood if they are strongly connected laterally as in figure 5. Such topologies are automatically learned as part of the self-organizing process. If several areas of the map are simultaneously active, long-range connections between them remain strong enough to survive connection pruning. These connections cause the units to behave as if they were neighbors on the map.

This property of LISSOM is potentially very significant in representing complex high-dimensional input spaces. While low-level sensory representation in the brain seems to be physically organized in 2-dimensional maps (such as retinotopic maps), it is possible that the functional representations make use of long-range lateral connections to represent more complex similarity relationships [19, 16]. LISSOM provides computational evidence that such maps can form by input-driven self-organization.

Conclusion

The LISSOM model demonstrates how lateral interaction and topological organization of cortical maps can be learned simultaneously from correlations in the input information. The model is biologically motivated, and its predictions agree well with experimental observations on cortical development. LISSOM is potentially capable of explaining various aspects of lateral and afferent connection development in the cortex, as well as the nature of lateral interactions in the cortex. The model delineates a functional role for connection death and demonstrates how it can result in nontopographically organized lateral connectivity. These results suggest that a single, general self-organizing process might underlie the development of most aspects of the structure and function of the neocortex. In future research, the model of the visual cortex shall be extended to the development of orientation columns, and used to investigate how lateral connections assist feature grouping and segmentation during perception.

Acknowledgments

This research was supported in part by National Science Foundation under grant #IRI-9309273. The simulations were performed on a Cray Y-MP 8/864 at the University of Texas Center for High-Performance Computing.

References

[1] E. M. Callaway and L. C. Katz, "Effects of binocular deprivation on the development of clustered horizontal connections in cat striate cortex," *Proceedings of the National Academy of Sciences, USA*, vol. 88, pp. 745–749, 1991.

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

[3] C. D. Gilbert, J. A. Hirsch, and T. N. Wiesel, "Lateral interactions in visual cortex," in *Cold Spring Harbor Symposia on Quantitative Biology, Volume LV*, pp. 663–677, Cold Spring Harbor Laboratory Press, 1990.

[4] D. H. Hubel and T. N. Wiesel, "Receptive fields and functional architecture in two nonstriate visual areas (18 and 19) of the cat," *Journal of Neurophysiology*, vol. 28, pp. 229–289, 1965.

[5] L. C. Katz and E. M. Callaway, "Development of local circuits in mammalian visual cortex," *Annual Review of Neuroscience*, vol. 15, pp. 31–56, 1992.

[6] T. Kohonen, "The self-organizing map," *Proceedings of the IEEE*, vol. 78, pp. 1464–1480, 1990.

[7] P. König, A. K. Engel, S. Löwel, and W. Singer, "Squint affects synchronization of oscillatory responses in cat visual cortex," *European Journal of Neuroscience*, vol. 5, pp. 501–508, 1993.

[8] S. Löwel and W. Singer, "Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity," *Science*, vol. 255, pp. 209–212, 1992.

[9] J. S. McCasland, K. L. Bernardo, K. L. Probst, and T. A. Woolsey, "Cortical local circuit axons do not mature after early deafferentation," *Proceedings of the National Academy of Sciences, USA*, vol. 89, pp. 1832–1836, March 1992.

[10] R. Miikkulainen, "Self-organizing process based on lateral inhibition and synaptic resource redistribution," in *Proceedings of the International Conference on Artificial Neural Networks* (Espoo, Finland), (Amsterdam; New York), pp. 415–420, North-Holland, 1991.

[11] K. D. Miller, "Development of orientation columns via competition between on- and off-center inputs," *NeuroReport*, vol. 3, pp. 73–76, 1992.

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

[13] 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*, vol. 45, pp. 7568–7589, 1992.

[14] 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*, vol. 87, pp. 8345–8349, 1990.

[15] H. Schwark and E. Jones, "The distribution of intrinsic cortical axons in area 3b of cat primary somatosensory cortex," *Experimental Brain Research*, vol. 78, pp. 501–513, 1989.

[16] W. Singer, C. Gray, A. Engel, P. König, A. Artola, and S. Bröcher, "Formation of cortical cell assemblies," in *Cold Spring Harbor Symposia on Quantitative Biology, Vol. LV*, (Cold Spring Harbor, NY), pp. 939–952, Cold Spring Harbor Laboratory, 1990.

[17] J. Sirosh and R. Miikkulainen, "How lateral interaction develops in a self-organizing feature map," in *Proceedings of the IEEE International Conference on Neural Networks* (San Francisco, CA), (Piscataway, NJ), IEEE, 1993.

[18] M. Stryker, J. Allman, C. Blakemore, J. Gruel, J. Kaas, M. Merzenich, P. Rakic, W. Singer, G. Stent, H. van der Loos, and T. Wiesel, "Group report. Principles of cortical self-organization," in *Neurobiology of Neocortex* (P. Rakic and W. Singer, eds.), (New York), pp. 115–136, Wiley, 1988.

[19] C. von der Malsburg and W. Singer, "Principles of cortical network organization," in *Neurobiology of Neocortex* (P. Rakic and W. Singer, eds.), pp. 69–99, New York: Wiley, 1988.

[20] C. von der Malsburg, "Self-organization of orientation-sensitive cells in the striate cortex," *Kybernetik*, vol. 15, pp. 85–100, 1973.

[21] G. K. von Noorden, *Binocular Vision and Ocular Motility. Theory and Management of Strabismus.* Baltimore: Mosby, 1990.