

A Self-Organizing Neural Network for Contour Integration through Synchronized Firing

Yoonsuck Choe and Risto Miikkulainen

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

Abstract

Contour integration in low-level vision is believed to occur based on lateral interaction between neurons with similar orientation tuning. The exact neural mechanisms underlying such interactions, and their developmental origins, are not well understood. This paper suggests through computational simulations that synchronized firing of neurons mediated by patchy lateral connections, formed through input-driven self-organization, can serve as such a mechanism. Furthermore, we argue that different degree of such patchy connections established during development may explain why different areas of the visual field show different degrees of contour integration in psychophysical experiments.

Introduction

Contour integration in low-level vision means forming a single coherent percept (i.e. a continuous contour) from a discontinuous sequence of line segments. Humans are very good at contour integration; understanding the underlying mechanisms can give us insights into how perceptual grouping can be implemented in general.

Psychophysical experiments (Field, Hayes, & Hess, 1993; Pettet, McKee, & Grzywacz, 1998; Geisler *et al.*, 1999) and computational models (Yen & Finkel, 1997; Li, 1998; Yen & Finkel, 1998; Geisler *et al.*, 1999) suggest that contour integration in the visual cortex may be due to interaction of neurons with similar orientation tuning. The models perform integration based on predetermined set of constraints such as relative orientation difference, distance, curvature, and change in curvature along with specific association rules, and they match experimental data quite well. However, how such constraints are implemented in the brain and how they emerge during early life of the animal remains an open question. Furthermore, the models do not explain why contour integration performance widely differs among different areas of the visual field. For example, contour integration is absent in the peripheral vision (Hess & Dakin, 1997), and subjective contours are harder to detect in the upper vs. the lower hemifield (Rubin, Nakayama, & Shapley, 1996).

This paper explores the possibility that a single mechanism of synchronized firing of neurons mediated by self-organized lateral connections may be responsible for these

phenomena. Previous computational studies on cortical self-organization have shown that strong lateral connections develop between neurons with similar orientation preference, and these connections can serve as a foundation for segmentation and binding (RF-SLISSOM, or Receptive Field Spiking Laterally Interconnected Synergetically Self-Organizing Map; Choe & Miikkulainen, 1998; Sirosh, 1995; Sirosh, Miikkulainen, & Choe, 1996; Miikkulainen *et al.*, 1997). This paper shows that (1) such lateral connections can also establish contour integration, (2) the specific lateral connection patterns necessary for contour integration can be learned in an input-driven self-organizing process like the rest of the orientation map, and (3) the weaker integration in peripheral vs. foveal and upper vs. lower hemifield may result from weaker lateral connectivity patterns, which in turn may originate from attentional modulation of the frequency and quality of input. In order to study these phenomena, RF-SLISSOM is extended to include long-range excitatory lateral connections.

Extending RF-SLISSOM for Contour Integration

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

The RF-SLISSOM model shows how self-organization and segmentation can be achieved in a single unified model. The long-range lateral interactions play a crucial role in both behaviors: they establish competition that drives self-organization, and they establish desynchronization that drives segmentation. The model does not include any long-

Copyright © 2000, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

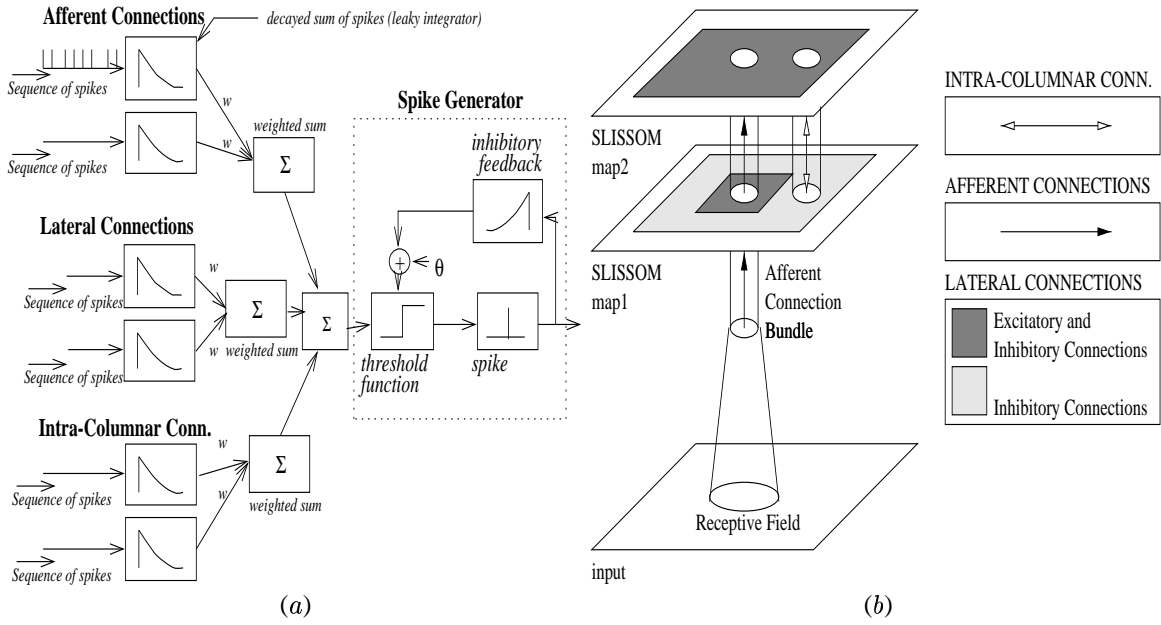


Figure 1: **The Stacked RF-LISSOM Architecture.** (a) The model neuron. Leaky integrators at each synapse perform decayed summation of incoming spikes, and the outgoing spikes establish a dynamic spiking threshold. (b) The overall organization of the Stacked RF-SLISSOM network. The cortical network consists of two sub-maps: MAP1 has short-range excitation and long-range inhibition, and drives the self-organization of the cortex. In MAP2, both excitation and inhibition are long range, establishing segmentation and binding. The two maps are connected with intra-columnar connections so that both the self-organization and segmentation dynamics are established in both maps.

range excitatory connections because they were not found necessary to model the above behaviors.

However, it turns out that such a parsimonious model cannot account for filling-in phenomena such as contour integration. The network has to be able to bind together representations that are separated by gaps: that is, it has to have long-range excitatory connections that link together the representations of the different segments of a fragmented contour.

The RF-SLISSOM model is extended in this paper with such long-range excitatory connections (figure 1). For conceptual clarity, the cortical network is divided into two separate components: (1) MAP1, which is similar to the RF-SLISSOM cortex with short-range excitatory and long-range inhibitory connections. This map has the task of driving the self-organization of the network into an ordered map. (2) MAP2, which has the task of establishing long-range segmentation and binding, with long-range excitatory connections that allow contour integration, and long-range inhibitory connections that allow segmentation of separate objects.

The two maps are assumed to be overlaid in one cortical network. In other words, the model predicts that some of the neurons in each hypercolumn are involved in establishing and maintaining organization, whereas others perform visual segmentation and binding. The details of the architecture, referred to as Stacked RF-LISSOM, are described next.

The Network Architecture

The details of the neuron model are shown in figure 1a. Each connection is a leaky integrator that performs exponentially

decayed summation of incoming spikes, thereby establishing not only spatial summation, but also temporal summation of activity. 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 an exponentially decayed sum of past spikes, formed by a similar leaky integrator as in the input synapses (Eckhorn *et al.*, 1990; Reitboeck, Stoecker, & Hahn, 1993).

The overall organization of the Stacked RF-SLISSOM model is shown in figure 1b. The net input $\sigma_{i,j}$ to the spike generator of the cortical neuron (in each map) at location (i, j) at time t consists of the input from a fixed-size receptive field in the retina, centered at the location corresponding to the neuron's location in the cortical network, from neurons around the same location in the other map, and from neurons around it in the same map:

$$\sigma_{i,j}(t) = \gamma_a \sum_{r_1, r_2} \xi_{r_1, r_2} \mu_{ij, r_1 r_2} + \gamma_c \sum_{p_1, p_2} \zeta_{p_1, p_2} \nu_{ij, p_1 p_2} + \gamma_e \sum_{k, l} \eta_{kl}(t-1) E_{ij, kl} - \gamma_i \sum_{k, l} \eta_{kl}(t-1) I_{ij, kl}, \quad (1)$$

where $\gamma_a, \gamma_c, \gamma_e$, and γ_i are the relative strengths of the afferent, intra-columnar, excitatory, and inhibitory contributions, ξ_{r_1, r_2} is the decayed sum of spikes of the retinal neuron (r_1, r_2) , $\mu_{ij, r_1 r_2}$ is the corresponding afferent connection weight, ζ_{p_1, p_2} is the decayed sum of spikes of the cortical neuron (p_1, p_2) of the other cortical map, $\nu_{ij, p_1 p_2}$ is the corresponding intra-columnar connection weight, $\eta_{kl}(t-1)$ is the decayed sum of spikes from the map neuron (k, l) at time $t-1$, and $E_{ij, kl}$ is the corresponding excitatory and




Figure 2: **The self-organized orientation map and patchy lateral connections.** (a) The orientation preferences of the neurons in MAP1 are shown; both maps (MAP1 and MAP2) developed similar organizations. The gray scale black \rightarrow white \rightarrow black represents preferences from 0 to 180 degrees as also indicated by the oriented line segments. The map has organized into orientation columns similar to those found in the visual cortex. (b) The strength of excitatory lateral connections of neuron at (41,29) of MAP2 is (i.e., at the location indicated by the arrow in (a)) plotted in gray scale (white \rightarrow black represents weak \rightarrow strong connection weight). This neuron is sensitive to 140 degree orientation, and its lateral connections link it to other neurons with the same orientation preference along the 140 degree direction across the map.

$I_{ij,kl}$ the inhibitory lateral connection weight.

The input is kept constant while the cortical response settles through the lateral connections, forming a concentrated, redundancy-reduced activation pattern. 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. The afferent, lateral and intra-columnar weights are modified according to the Hebbian principle:

$$w_{ij,mn}(t) = \frac{w_{ij,mn}(t-1) + \alpha V_{ij} X_{mn}}{\sum_{mn} [w_{ij,mn}(t-1) + \alpha V_{ij} X_{mn}]}, \quad (2)$$

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_c for intra-columnar, α_e for excitatory, and α_i for inhibitory connections), V_{ij} and X_{mn} are the average spiking rates of the neurons. Those connections that become less than 0.001 in this process are killed, modeling death of unused connections.

This process of weight adaptation is repeated for a number of input patterns (e.g. oriented Gaussians), and the neurons become gradually sensitive to particular orientations at particular locations, and the map forms a global retinotopic orientation map similar to that in the visual cortex (Blasdel, 1992; Blasdel & Salama, 1986). The self-organized map will then synchronize and desynchronize the firing of neurons to indicate binding and segmentation of visual input to different objects.

Experiments and Results

A Stacked RF-SLISSOM network with a 44×44 retina and a 72×72 cortex was trained for 30,000 iterations with elongated Gaussian bars at random locations in the retina (such input approximates the local features of natural visual stimuli after the edge detection and enhancement mechanisms in the retina). Excitatory lateral connections in MAP1 had an initial radius of 3 and gradually reduced to 1, and inhibitory lateral connections had a fixed radius of 7. In MAP2, both types of lateral connections had a radius of 19. Afferent connections to the retina had a radius of 7 in both maps, and intra-columnar connections a radius of 1. During each training presentation, the network was allowed to settle for 13 time steps (through equation 1) and all connections except the inhibitory lateral connections were updated according to equation 2. The fixed inhibition provides a baseline global inhibition similar to other cortical models (Eckhorn *et al.*, 1988; von der Malsburg & Buhmann, 1992), used because it is simple and sufficient to establish segmentation. The simulations are not particularly sensitive to these parameter values as long as they are qualitatively similar.

A well-formed orientation map emerged in this process (figure 2a). Because of the intra-columnar connections, similar activity patterns formed on both maps during self-organization, and they developed almost identical organizations (only MAP1 is shown in figure 2). The lateral connections with weights less than 0.001 were killed in the end, thus leaving a patchy connection profile (figure 2b). Only those connections that link areas with highly-correlated ac-





Figure 3: **Contour integration with Stacked RF-SLISSOM.** (a) A smooth input contour. (b) A randomly oriented contour. (d,e) The MUAs of the areas responding to each line segment in (a) and (b). Time is on the x -axis and the three rows from top to bottom represent MUA of area 1, area 2 and area 3 by gray-scale coding. (c) A comparison of the correlation coefficients between the MUAs. In each experiment (smooth vs. random contours), the MUAs from three areas in MAP2 were gathered over 600 iterations and the correlations between the MUAs were calculated. The x -axis is the index for MUA pairs. For example, 0, 1, ... represents MUA pairs (1, 2), (1, 3), The y -axis is the correlation coefficient r as calculated by equation 3. When the contour elements are smooth, the MUAs are highly correlated, but the correlation is very low when the contour elements are randomly oriented relative to each other. This indicates that the network binds smooth contours together through synchronized firing of neurons but perceives random contours as multiple objects. This result is consistent with experimental data and can be explained by the existence of patchy excitatory lateral connections developed through self-organization.

tivity, such as those along a continuous contour, remain in the end.

The patchy lateral connections therefore form the foundation for feature binding and contour integration in the model. The network should be able to bind together separated line segments if they are collinear and therefore likely to be part of the same contour. If their directions do not agree, they should not be bound together because they are likely to belong to different objects.

To test this hypothesis, three elongated Gaussian bars arranged in two different configurations, i.e. smooth and random contours (figure 3a,b), were presented to the network. For each input bar, the number of spikes generated by the

area of the cortex that responded to the bar was counted at each time step (figure 3d,e). This quantity is called the Multi-Unit Activity of the response, or MUA, and it can be used to identify which area of the cortex is the most active at each time step.

In order to measure synchronization between two areas, linear correlation coefficient r between their MUA sequences can be calculated as follows:

$$r = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2} \sqrt{\sum_i (y_i - \bar{y})^2}} \quad (3)$$

where x_i and y_i , $i = 1, \dots, N$ are the MUA values at time i for the two areas representing the two different objects in

the scene, and \bar{x} and \bar{y} are the mean of each sequence.

Using this measure, the contour integration capability of the network was tested. The degree of synchrony in the two different cases was compared: in the first case, the segments lined up to form a smooth contour, and in the second, their orientation was random (figure 3).

During self-organization, patchy lateral connections formed between neurons that represent collinear line segments (figure 2). In the first case these connections synchronize the responses of the three areas. No such connections exist for line segments that are not collinear, and the MUAs remain desynchronized in the second case. This way the network perceives the collinear segments as one object, but the random line segments as three separate objects.

Although the example in figure 3 shows only three bars, the effect is robust. The network may integrate long chains of near-collinear line segments this way, whereas a perpendicular segment in the middle will break the integration. This way the network suggests a possible mechanism for contour integration, and demonstrates how the necessary connections emerge automatically as a side effect of Hebbian self-organization of lateral connections.

Discussion

Whether contour integration occurs or not in the model depends on whether the cortical areas are connected with excitatory lateral connections or not. The model therefore suggests an explanation for the different contour integration capability of the different visual areas: integration is possible only if focused (i.e. patchy) lateral connections exist linking collinear neurons with similar orientation preferences. If peripheral areas and the upper visual field do not receive dense enough visual input for such connections to form during development, the connections become diffuse, resulting in weaker integration.

Statistics of images projected on the retina indeed support this hypothesis. Reinagel & Zador (1999) showed that human gaze most often falls upon areas with high contrast and low pixel correlation than other areas. As a result, sharp images project more often on the fovea than the periphery, allowing more specific connections to form. A similar method can be used to find out if there's a difference in statistical distribution of image features in the lower vs. upper hemisphere. Such statistical difference together with Hebbian self-organization would result in different contour integration capability in different visual areas.

Another way of verifying this hypothesis would be to rear an animal with eye glasses that flip the input to the upper and lower hemifield. After the critical period, the animal's performance on contour detection task could be measured, and the connectivity patterns formed in the upper and lower hemifield analyzed. The prediction is that high connectivity and good integration would occur in the upper hemifield, instead of the lower hemifield as in normal animals.

The Stacked RF-SLISSOM network can be tested further in more advanced contour integration and gestalt perception tasks such as subjective contour detection and occluded object recognition. The excitatory lateral connections should

be able to mediate these phenomena as well. Moreover, orientation maps for subjective contours are known to exist in V2 (Sheth *et al.*, 1996); the Stacked RF-SLISSOM could be extended to model V2 as well, thereby extending our understanding of the visual self-organization and function at higher levels.

Conclusion

Input-driven self-organization of afferent, intra-columnar, and lateral connections was shown to give rise to patchy connectivity patterns that can facilitate contour integration in the visual cortex. The model also suggests that different statistics of input presentations, and the resulting patchy lateral connection patterns, may be the cause for the different degrees of contour integration observed in the different visual areas. It should be possible to account for other low-level gestalt phenomena and also subjective contour effects with similar computational principles.

Acknowledgments.

This research was supported in part by National Science Foundation under grant #IRI-9811478.

References

- Blasdel, G. G., and Salama, G. 1986. Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature* 321:579–585.
- Blasdel, G. G. 1992. Orientation selectivity, preference, and continuity in monkey striate cortex. *Journal of Neuroscience* 12:3139–3161.
- Choe, Y., and Miikkulainen, R. 1998. Self-organization and segmentation in a laterally connected orientation map of spiking neurons. *Neurocomputing* 21:139–157.
- Eckhorn, R.; Bauer, R.; Jordan, W.; Kruse, M.; Munk, W.; and Reitboeck, H. J. 1988. Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics* 60:121–130.
- Eckhorn, R.; Reitboeck, H. J.; Arndt, M.; and Dicke, P. 1990. Feature linking via synchronization among distributed assemblies: Simulations of results from cat visual cortex. *Neural Computation* 2:293–307.
- Field, D. J.; Hayes, A.; and Hess, R. F. 1993. Contour integration by the human visual system: Evidence for a local association field. *Vision Research* 33:173–193.
- Geisler, W. S.; Thornton, T.; Gallogly, D. P.; and Perry, J. S. 1999. Image structure models of texture and contour visibility. In *Proceeding of the NATO Workshop on Search and Target Acquisition*.
- Hess, R. F., and Dakin, S. C. 1997. Absence of contour linking in peripheral vision. *Nature* 390:602–604.
- Li, Z. 1998. A neural model of contour integration in the primary visual cortex. *Neural Computation* 10:903–940.
- Miikkulainen, R.; Bednar, J. A.; Choe, Y.; and Sirosh, J. 1997. Self-organization, plasticity, and low-level visual phenomena in a laterally connected map model of

- the primary visual cortex. In Goldstone, R. L.; Schyns, P. G.; and Medin, D. L., eds., *Perceptual Learning*, volume 36 of *Psychology of Learning and Motivation*. San Diego, CA: Academic Press. 257–308.
- Pettet, M. W.; McKee, S. P.; and Grzywacz, N. M. 1998. Constraints on long range interactions mediating contour detection. *Vision Research* 38:865–879.
- Reinagel, P., and Zador, A. M. 1999. Natural scene statistics at the center of gaze. *Network: Computation in Neural Systems* 10:1–10.
- Reitboeck, H.; Stoecker, M.; and Hahn, C. 1993. Object separation in dynamic neural networks. In *Proceedings of the IEEE International Conference on Neural Networks* (San Francisco, CA), volume 2, 638–641. Piscataway, NJ: IEEE.
- Rubin, N.; Nakayama, K.; and Shapley, R. 1996. Enhanced perception of illusory contours in the lower versus upper visual hemifields. *Science* 271:651–653.
- Sheth, B. R.; Sharma, J.; Rao, S. C.; and Sur, M. 1996. Orientation maps of subjective contours in visual cortex. *Science* 274:2110–2115.
- Sirosh, J.; Miikkulainen, R.; and Choe, Y., eds. 1996. *Lateral Interactions in the Cortex: Structure and Function*. Austin, TX: The UTCS Neural Networks Research Group. Electronic book, ISBN 0-9647060-0-8, <http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96>.
- Sirosh, J. 1995. *A Self-Organizing Neural Network Model of the Primary Visual Cortex*. Ph.D. Dissertation, Department of Computer Sciences, The University of Texas at Austin, Austin, TX. Technical Report AI95-237.
- von der Malsburg, C., and Buhmann, J. 1992. Sensory segmentation with coupled neural oscillators. *Biological Cybernetics* 67:233–242.
- Yen, C.-C., and Finkel, L. H. 1997. Identification of salient contours in cluttered images. In *Computer Vision and Pattern Recognition*, 273–279.
- Yen, S.-C., and Finkel, L. 1998. Extraction of perceptually salient contours by striate cortical networks. *Vision Research* 38:719–741.