

Joint maps for orientation, eye, and direction preference in a self-organizing model of V1

James A. Bednar^a Risto Miikkulainen^b

^a*Institute for Adaptive and Neural Computation, The University of Edinburgh, UK*

^b*Department of Computer Sciences, The University of Texas at Austin, USA*

Abstract

Primary visual cortex (V1) contains overlaid feature maps for orientation (OR), motion direction selectivity (DR), and ocular dominance (OD). Neurons in these maps are connected laterally in patchy, long-range patterns that follow the feature preferences. Using the LISSOM model, we show for the first time how realistic laterally connected joint OR/OD/DR maps can self-organize from Hebbian learning of moving natural images. The model predicts that lateral connections will link neurons of either eye preference and with similar DR and OR preferences. These results suggest that a single self-organizing system may underlie the development of spatiotemporal feature preferences and lateral connectivity.

Key words: Ocular dominance maps, Orientation maps, Direction maps, Development, Visual cortex

1 Introduction

Most simple cells in the primary visual cortex (V1) in mammals are selective for the direction and orientation of a moving stimulus, and respond to inputs at corresponding locations in each eye. Recent measurement techniques have made it possible to plot the neurons' full spatiotemporal receptive fields, which include specific excitatory (ON) and inhibitory (OFF) subregions that vary over time [5]. The functional properties of these cells form a mosaic across V1, with patches of nearby cells preferring the left or right eye (or both), and similar directions and orientations [3, 6, 9].

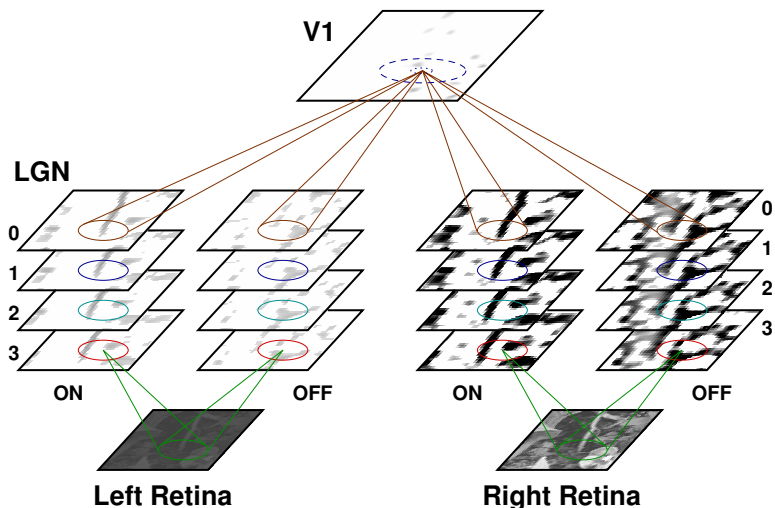
In addition to their afferent input from the LGN, neurons in these maps are connected intracortically through specific long-range lateral connections. The lateral connections have been found to link cells with similar orientation preferences [4], which can allow the connections to suppress redundancy in the input and improve the cells' ability to detect changes in a stimulus [7]. However, the role of these connections in the development and adult function of directional selectivity is not yet clear.

Several computational models have demonstrated that directional selectivity and interleaved orientation (OR) and direction selectivity (DR) maps can develop through activity-dependent self-organization [8, 11]. However, these simulations have only modeled afferent connection

Email addresses: jbednar@inf.ed.ac.uk (James A. Bednar), risto@cs.utexas.edu (Risto Miikkulainen).

Fig. 1. **LISSOM OR/OD/DR model.**

Each V1 unit receives inputs from RFs on sixteen LGN sheets modeling different response types (ON- or OFF-center) and different connection lags (3, 2, 1, or 0). LGN RFs for one V1 unit are shown as solid circles; V1 units also have lateral excitatory (small dotted circle) and lateral inhibitory (large dashed circle) connections to their neighbors. Each LGN unit receives input from either the left or the right eye; sample RFs are shown for two LGN units for each eye. Moving input patterns are drawn on the retinal sheets in four discrete timesteps, like frames of a movie. Matching patterns are presented in each eye, with the relative brightness of each eye chosen randomly. At the first timestep, the ON and OFF LGN units with time lag 3 compute their activity. At each subsequent timestep, the input pattern is moved slightly and LGN units with lags 2, 1, and 0 compute their activity in turn. Once all LGN units have been activated, initial V1 activity is computed from the LGN responses, and the activity then spreads laterally within V1.



learning, and have treated lateral connections as fixed. Also, although ocular dominance (OD) has been modeled in conjunction with orientation, joint OR/OD/DR models have not yet been published. Including OD is potentially crucial, because it is known to interact with the orientation map [3]. Finally, most prior models are based on abstract input patterns, and it is difficult to extend such models to process realistic images.

In prior work with the LISSOM self-organizing model (Laterally Interconnected Synergetically Self-Organizing Map), we have demonstrated how a Hebbian learning process can develop orientation and direction maps, with patterned lateral connections between them [2]. We have also shown separately how individual maps, including those for ocular dominance, can form from natural images or spontaneous neural activity [1, 7]. The model suggests that self-organized maps and lateral connections function in adult visual perception to segment and bind coherent objects and reduce redundancy in the input, and that visual illusions and aftereffects arise through this process [7]. In this paper we extend the combined OR/DR model to include two eyes, to model ocular dominance, and to use natural image stimuli as inputs. Together, these results show that activity-dependent self-organization can explain the development of much of the map-level organization of the primary visual cortex.

2 The LISSOM model

The model architecture is shown in figure 1, and will be briefly reviewed below (see [1] or [7] for more details). The model consists of a hierarchy of two-dimensional sheets of neural units modeling different areas of the visual system: two sheets of retinal photoreceptors representing the left and right eyes, several paired sheets of ON-center/OFF-surround and OFF-center/ON-surround LGN units, and a sheet of cortical units (“neurons”) representing V1. The ON/OFF units are also called LGN units for simplicity, although they represent the entire pathway be-

tween the retinal photoreceptors and V1, including the retinal ganglion cells and connection pathways. Because the focus is on the two-dimensional organization of V1, each cortical neuron corresponds to a vertical column of cells through the six anatomical layers of the cortex.

Compared to simpler OR-only LISSOM networks, the model introduced in this paper includes two eyes (to model OD) and multiple LGN sheets with different time-delayed copies of previous input patterns. The time delays model the “lagged” cells recently found in cat LGN that respond to retinal inputs only after a fixed delay. The delay time in these lagged cells varies over a continuous range up to hundreds of milliseconds [12]. We will show below that V1 neurons can use these timing differences to develop spatiotemporal receptive fields. Similarly, neurons will develop eye preferences based on differences in the patterns presented to each eye. These extensions allow the model V1 to develop simultaneous maps for orientation, ocular dominance, and motion direction. The extensions result in sixteen LGN sheets in total, one ON-center and one OFF-center sheet at each of four different lags for each eye, compared to two (one ON and one OFF) sheets in the OR-only LISSOM model.

The training patterns consisted of randomly chosen patches of natural images, moving in random directions. A sample input is drawn in the photoreceptor sheets of figure 1. The strength of the image in each eye was chosen randomly, but the total strength between the eyes was kept constant. For each image presentation, the activity level of each LGN unit is calculated from the activity level its receptive fields. Specifically, each LGN unit (i, j) with lag t computes its response η_{ij} as a scalar product of a fixed Difference of Gaussians weight vector and the activity in its receptive fields on a photoreceptor sheet at time t :

$$\eta_{ij} = \sigma \left(\sum_{\rho} \gamma_{\rho} \sum_{ab} X_{\rho ab} w_{ij, \rho ab} \right), \quad (1)$$

where ρ iterates over all RFs of this unit (for LGN units in this simulation, a single RF on either the left or right retina), σ is a piecewise linear sigmoid activation function, γ_{ρ} is a constant scaling factor, $X_{\rho ab}$ is the activation of input unit (a, b) on sheet ρ at timestep t , and $w_{ij, \rho ab}$ is the corresponding weight value. Each V1 neuron computes its initial response like an LGN unit, except that ρ iterates over all of the sixteen ON or OFF LGN sheets for which the V1 unit has RFs. After the initial V1 response, the LGN activity remains constant while V1 activity settles through short-range excitatory and long-range inhibitory lateral interaction:

$$\eta_{ij}(t) = \sigma \left(\sum_{\rho} \gamma_{\rho} \sum_{ab} X_{\rho ab}(t-1) w_{ij, \rho ab} \right), \quad (2)$$

where ρ identifies either an RF on the LGN sheet, or the lateral excitatory or inhibitory weights to V1, γ_{ρ} is a constant scaling factor for each ρ (negative for inhibitory lateral weights), and $X_{\rho ab}(t-1)$ is the activation of input unit (a, b) during the previous settling step. The V1 activity pattern starts out diffuse, but within a few iterations of equation 2, converges into a small number of stable focused patches of activity, or activity bubbles.


After the activity has settled, the connection weights of each RF of every V1 neuron are modified according to a normalized Hebb rule. For a given V1 unit (i, j) and connection to any unit (a, b) in a specific RF ρ , the new connection weight is:

$$w'_{ij, \rho ab} = \frac{w_{ij, \rho ab} + \alpha_{\rho} \eta_{ij} X_{\rho ab}}{\sum_{ruv} [w_{ij, ruv} + \alpha_r \eta_{ij} X_{ruv}]}, \quad (3)$$

where η_{ij} stands for the activity of neuron (i, j) in the final activity bubble, $w_{ij, \rho ab}$ is the con-

Fig. 2. Self-organized spatiotemporal receptive field.

These images show the OFF LGN weights subtracted from the ON for each lag, for a typical V1 neuron. Light shades indicate a net ON subregion of the retina, i.e. an area that will excite the neuron if it is brighter than its surround. Dark shades indicate OFF regions, and medium gray indicates no preference. For this neuron, the most effective stimulus would be a binocular vertical bar or grating moving to the right. More specifically, the neuron will be highly active at time t if a bright bar is aligned with the ON region of both eyes in (a) at time $t - 3$, with the ON region of (b) at time $t - 2$, and so on until at the present time t it is aligned with the ON region of (d). Visual cortex neurons in animals have similar spatiotemporal properties [5].



nection weight from the previous fixation, α is the learning rate for each type of connection, and $X_{\rho ab}$ is the presynaptic activity. For afferent connections, r in the denominator iterates over all afferent receptive fields (and (u, v) over all units in each such RF), whereas the weights for lateral connections are normalized over only one type, either excitatory or inhibitory. The larger the product of the pre- and post-synaptic activity $\eta_{ij} X_{\rho ab}$, the larger the weight change. At long distances, few neurons have correlated activity and most long-range connections eventually become weak. The weakest connections are eliminated periodically, resulting in patchy lateral connectivity similar to that observed in the visual cortex.

For the experiments reported in this paper, eight 36×36 ON-center and eight 36×36 OFF-center unit sheets received input from a 54×54 photoreceptor sheet. Initially, the afferent weights of the 142×142 V1 neurons were random, and the lateral weights had a smooth circular Gaussian profile. The learning parameters are the same as those for our previous OR-only model [7].

3 Results

Figure 2 shows the self-organized afferent weights for a representative neuron after 20,000 image fixations. Nearly all neurons developed spatiotemporal receptive fields selective for orientation, and most were also selective for direction and either the right or left eye. These RFs are similar to those found experimentally in the cortex [5]. As in V1 cells, the model RFs can only rarely be expressed as a function of a fixed spatial RF and a time-dependent scaling factor. The RFs are thus *space-time inseparable* [5]. Also as in the cortex, the orientation and direction preferences of a neuron were generally perpendicular.

Figure 3 shows the self-organized global DR, OD, and OR maps, which are also similar to those found in animals [3, 6, 9]. For instance, a patch of neurons highly selective for one orientation and direction of motion will often have an adjacent or contiguous patch selective for the same orientation, but the opposite direction. Ocular dominance stripe boundaries also tend to intersect orientation patch boundaries at right angles, as in macaque monkeys [3]. Long-range lateral connections within the map follow this global organization, connecting neurons similar in both orientation and direction preference, but to both eyes (figure 3). These connection patterns represent the correlations between units over the course of self-organization. These connections constitute predictions of the model, and match preliminary findings reported for ferrets [10]. As will be discussed below, the self-organized lateral weight patterns are likely to play a crucial role in adult visual perception as well as development.

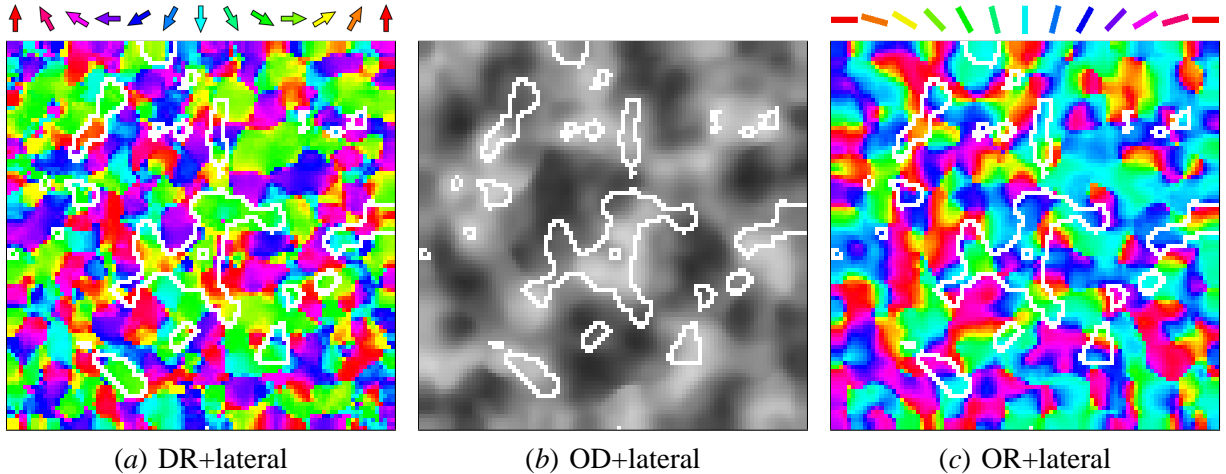


Fig. 3. Maps and lateral connections. These plots show the separate maps for (a) motion direction, (b) ocular dominance, and (c) orientation that emerged through self-organization. The DR and OR maps use the color keys shown, while the OD map shows left eye preferences as dark areas and right eye preferences as light. The OD gray level indicates how strong the preference is; fully black or fully white neurons would be completely monocular. On each map, the self-organized long-range lateral inhibitory connections for the neuron in figure 2 are outlined in white. This neuron, colored in green in (a) and blue in (c), and outlined with a small white square in the center of the plot, prefers nearly vertical stimuli moving to the right and slightly downwards. Its connections are patchy and link to similar DR preferences (other green neurons in the DR map) and OR preferences (other blue and cyan neurons in the OR map). The neuron is binocular, and connects to neurons selective for both eyes, although it has stronger connections to one eye. Overall, the long-range connections target similar orientation preferences, direction preferences, and both eyes, and avoid orthogonal orientations and opposite directions. Color versions of these plots are available at www.nn.cs.utexas.edu/keyword?bednar:cns05.

4 Discussion and Future Work

The results demonstrate that a single Hebbian learning algorithm can explain how topographic maps, spatiotemporal receptive fields, and lateral connections synergetically self-organize from moving stimuli. The model predicts that long-range lateral connections in V1 will be found to connect neurons with similar direction preference as well as orientation preference, as suggested in preliminary reports [10]. In future work we will simulate larger cortical areas and lateral connection radii, to determine whether the lateral connections extend further along the preferred orientation, as found in tree shrew [4]. We will also validate the model maps against other experimental results from specific species and study how the orientation, direction preference, and lateral connectivity interact with other stimulus dimensions, such as spatial frequency.

The preferences for each of the features represented in the model can develop through a variety of mechanisms in addition to those presented above. For instance, realistic LISSOM orientation maps can develop from spontaneous neural activity before eye opening, without requiring natural images [1, 7]. Eye preferences can also develop from differences in object position between the two eyes, apart from the intensity differences modeled here. Similarly, motion preferences can arise from any source of different delays in neural activity propagation, in addition to the LGN lags modeled here. Thus the map development process is highly robust, suggesting that many types and sources of neural activity may contribute to map development in animals.

Future work will investigate how the OR/OD/DR map and its connections operate in the adult. The self-organized connection patterns should allow moving objects to be bound together and segmented from other objects, as previously shown for static patterns [7]. Short-term direction-specific adaptation of the lateral connections should replicate the motion aftereffect, also called the waterfall illusion, as we found for the tilt aftereffect [7]. Simulating both eyes should allow the interocular transfer of such effects to be tested in the model. This work will help unify explanations of development and adult visual function into a coherent theory of the visual cortex.

5 Conclusion

Hebbian learning of moving, oriented patterns in a self-organizing model can explain how orientation selectivity, direction selectivity, eye preference, and lateral connectivity develop in V1. The model provides specific predictions for the role of the lateral connections in the development and function of the direction map. This study and future work will help strengthen our understanding of the visual cortex as a continuously adaptive, self-organizing system.

Acknowledgments

Supported in part by the National Science Foundation under grant IIS-9811478 and by the National Institutes of Mental Health under Human Brain Project grant 1R01-MH66991.

References

- [1] J. A. Bednar, *Learning to See: Genetic and Environmental Influences on Visual Development*, Ph.D. thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX (2002), technical Report AI-TR-02-294.
- [2] J. A. Bednar and R. Miikkulainen, Self-organization of spatiotemporal receptive fields and laterally connected direction and orientation maps, *Neurocomputing* (2003), 52–54:473–480.
- [3] G. G. Blasdel, Differential imaging of ocular dominance columns and orientation selectivity in monkey striate cortex, *The Journal of Neuroscience* (August 1992), 12:3115–3138.
- [4] W. H. Bosking, Y. Zhang, B. R. Schofield, and D. Fitzpatrick, Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex, *J. Neurosci* (1997), 17:2112–2127.
- [5] G. C. DeAngelis, I. Ohzawa, and R. D. Freeman, Receptive-field dynamics in the central visual pathways, *Trends in Neurosciences* (1995), 18:451–458.
- [6] S. Löwel, Ocular dominance column development: Strabismus changes the spacing of adjacent columns in cat visual cortex, *The Journal of Neuroscience* (1994), 14:7451–7468.
- [7] R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh, *Computational Maps in the Visual Cortex* (Springer, Berlin, 2005), in press.
- [8] H. Shouno and K. Kurata, Formation of a direction map by projection learning using Kohonen’s self-organization map, *Biological Cybernetics* (2001), 85:241–246.
- [9] M. Weliky, W. H. Bosking, and D. Fitzpatrick, A systematic map of direction preference in primary visual cortex, *Nature* (1996), 379:725–728.
- [10] L. E. White, W. H. Bosking, M. Weliky, and D. Fitzpatrick, Direction selectivity and horizontal connections in layers 2/3 of ferret primary visual cortex (V1), in: *Society for Neuroscience Abstracts* (Washington, DC: Society for Neuroscience, 1996), volume 22, 1610.
- [11] S. Wimbauer, O. G. Wenisch, J. L. van Hemmen, and K. D. Miller, Development of spatiotemporal receptive fields of simple cells: II. Simulation and analysis, *Biol. Cybernetics* (1997), 77:463–477.
- [12] J. Wolfe and L. A. Palmer, Temporal diversity in the lateral geniculate nucleus of cat, *Visual Neuroscience* (1998), 15:653–675.