Prenatal development of ocular dominance and orientation maps in a self-organizing model of V1

Stefanie Jegelka [a] James A. Bednar [b] Risto Miikkulainen [c]

[a] *Wilhelm Schickard Institute for Computer Sciences, University of Tübingen, Germany*

[b] *Institute for Adaptive and Neural Computation, The University of Edinburgh, UK*

[c] *Department of Computer Sciences, The University of Texas at Austin, USA*

Abstract

How orientation and ocular-dominance maps develop before visual experience begins is controversial. Possible influences include molecular signals and spontaneous activity, but their contributions remain unclear. This paper presents LISSOM simulations suggesting that previsual spontaneous activity alone is sufficient for realistic OR and OD maps to develop. Individual maps develop robustly with various previsual patterns, and are aided by background noise. However, joint OR/OD maps depend crucially on how correlated the patterns are between eyes, even over brief initial periods. Therefore, future biological experiments should account for multiple activity sources, and should measure map interactions rather than maps of single features.

Key words: Ocular dominance maps, Orientation maps, Development, Visual cortex, Spontaneous activity

1 Introduction and related work

At the time newborn cats, ferrets, and monkeys open their eyes, their primary visual cortex (V1) is already organized into maps of ocular dominance (OD) and orientation (OR) [3, 9]. The mechanisms by which these maps initially develop remain controversial. Changing or blocking the patterns of spontaneous neural activity before birth can alter the maps, suggesting that such activity is involved in normal development (reviewed in [18]). Yet some researchers have argued that previsual OD map development depends primarily on molecular signals or other activity-independent processes, because maps develop even after removal of retinal activity by enucleation [4]. It is also not clear whether the known sources of previsual activity have the properties necessary to drive the development of maps that have matching OR in both eyes [7] and specific interactions between OR and OD maps

Email addresses: jegelka@informatik.uni-tuebingen.de (Stefanie Jegelka), jbednar@inf.ed.ac.uk (James A. Bednar), risto@cs.utexas.edu (Risto Miikkulainen).

[1, 2]. On the other hand, concerns have been raised that the enucleation experiments may not have entirely eliminated the earliest traces of retinal activity [14]. In addition, the retina is not necessarily the only relevant source of spontaneous activity [14]. Thus it has been difficult to ascertain the role of spontaneous previsual activity in initial map development.

Computational models allow a much wider variety of scenarios to be tested than do current experimental techniques. Hence they can be used to gain insight into the complex process of cortical map development, and to guide future experiments. Many models of OR and OD development have achieved maps with realistic features ([5, 8, 12, 13, 15, 16], review in [17]), but none of them models previsual development using activity patterns that have a clear biological interpretation. The first experiment described in this paper demonstrates that a balance of correlated and uncorrelated spontaneous activities from different areas can generate initial maps with realistic interactions. The second experiment further demonstrates how random noise in such activity makes this process more robust.

2 The LISSOM model

The simulations are based on the LISSOM model of the primary visual cortex [11], which is outlined in Fig. 1. The model includes three hierarchical levels of two-dimensional sheets of neural units. The lowest-level sheets simulate the photoreceptors in the two retinas. The middle level models the LGN ON-center/OFF-surround and OFF-center/ON-surround cells, one ON/OFF-center pair of sheets for each eye. Each LGN unit receives input from a circular receptive field on one of the photoreceptor sheets. At the top level, a sheet of cortical units represents V1. Because the focus is on the two-dimensional organization of V1, each cortical unit corresponds to a vertical column of cells through the six anatomical layers of the cortex. V1 units interact through short-range excitatory and long-range inhibitory connections which model the net effect of lateral connections for the high-contrast inputs that drive Hebbian learning [11].

Input to the model consists of a series of activity patterns on the two retinas, such as grayscale photographic bitmaps. Each LGN unit computes its response by applying a piecewise linear sigmoid activation function σ to the scalar product of a fixed Difference-of-Gaussians weight vector w and its receptive fields:

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

where $X_{\rho ab}$ denotes the activation of unit (a, b) in receptive field ρ, $w_{ij,\rho ab}$ denotes the weight from that unit to LGN unit (i, j), and γ_{ρ} is a constant scaling factor. In this simulation, each LGN unit has one receptive field, on either the left or right eye.

V1 units compute an initial response analogously, from their receptive fields on the

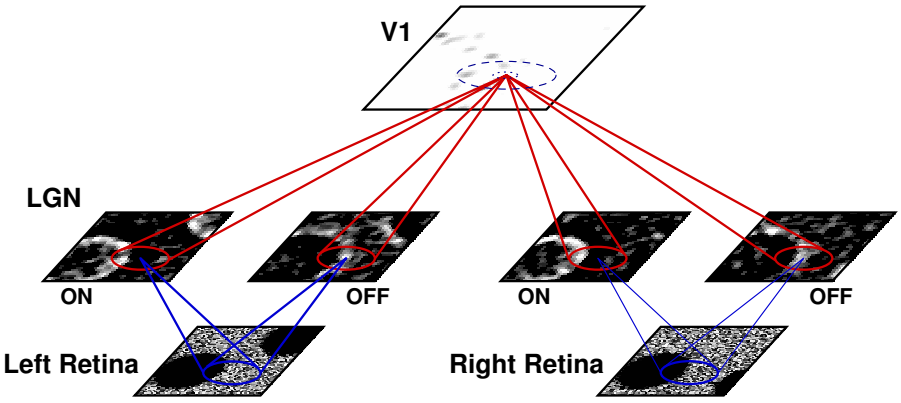

Fig. 1. LISSOM model. Each sheet of LGN units receives input from its circular receptive fields (RF) on the retina. Sample connections are shown for one unit in each sheet. V1 units have circular RFs on LGN sheets of both types. Initial lateral excitatory and inhibitory inputs to a V1 unit are illustrated by the small dotted and large dashed circle, respectively.

LGN sheets. The activity $\eta_{\mathrm{V1},ij}$ of each V1 unit (i, j) then settles through short-range excitatory and long-range inhibitory lateral interactions:

$$\eta_{\mathrm{V1},ij}(t) = \sigma\left(\sum_{\rho}\gamma_{\rho}\sum_{ab}X_{\rho ab}(t-1)w_{ij,\rho ab}\right), \tag{2}$$

where ρ denotes one of the six receptive fields of V1 neurons (four afferent LGN RFs and two lateral), γ_{ρ} a constant scaling factor (negative for lateral inhibitory connections), $X_{\rho ab}(t-1)$ the activation of unit (a, b) (on either the LGN or V1) in the previous settling step, and $w_{ij,\rho ab}$ is the weight assigned to that unit. The V1 activity pattern starts out diffuse, but within a few iterations of Eq. 2, it converges into a small number of stable focused patches of activity, or activity bubbles. Throughout settling, the LGN activity remains constant.

After the activity has settled, the connection weights of each V1 neuron are modified. Initially, the afferent weights for V1 units are random and identical for both eyes, with locations centered around the corresponding location on the retina. All V1 weights adapt according to a normalized Hebb rule with learning rate α_{ρ} for each connection type ρ

$$w'_{ij,\rho ab} = \frac{w_{ij,\rho ab} + \alpha_{\rho}\eta_{\mathrm{V1},ij}X_{\rho ab}}{\sum_{ruv}(w_{ij,ruv} + \alpha_{r}\eta_{\mathrm{V1},ij})}, \tag{3}$$

where $w'_{ij,\rho ab}$ is the new weight value. The afferent, lateral excitatory, and lateral inhibitory connections are normalized separately. All weights are defined to be positive, including the inhibitory weights, whose inhibitory effect is due to a negative multiplicative factor γ in Eq. 2. At long distances, very few neurons have correlated activity and therefore most long-range connections eventually become weak. The weak connections are eliminated periodically, resulting in patchy lateral connectivity similar to that observed in the visual cortex. The radius of the lateral excitatory interactions starts out large, but is decreased to the nearest neighbors according to

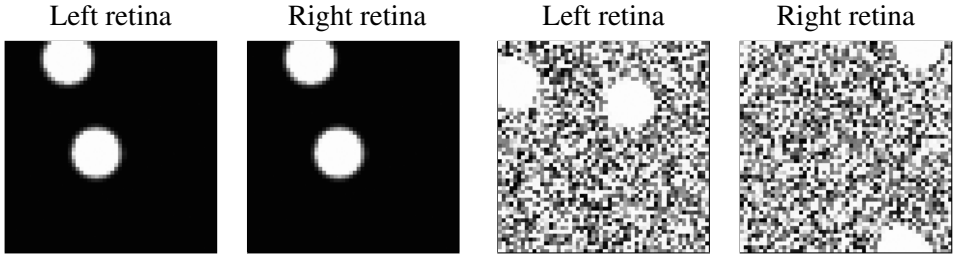

(a) Correlated disks (b) Uncorrelated disks with noise

Fig. 2. Sample input patterns from two conditions: (a) correlated disks without noise, and (b) uncorrelated disks with uncorrelated random noise.

a preset schedule during self-organization.

In the experiments reported in the paper, the four 36×36 LGN ON/OFF-center sheets received input from two 54×54 retina sheets. The V1 sheet was 64×64. The learning parameters were the same as in earlier simulations [11], scaled for this cortex size. The afferent receptive fields had a radius of 6.5, and the initial excitatory and inhibitory connections a radius of 6.4 and 15, respectively. The number of settling steps was initially 13, and (to save computation time) was gradually reduced to 9 as neurons became selective. By 10,000 iterations, the maps had become stable, changing little with further training.

3 Experiments

These LISSOM experiments concentrated on the developmental phase prior to visual experience. Spontaneous activity was modeled as disk-shaped regions of correlated activity combined with an overall level of uniform random noise over the entire retina (Fig. 2, [11]). In each training iteration, a new input pattern was randomly generated using the specified input components. Similar results were obtained with stationary disks presented over several iterations, with a lower learning rate. The fuzzy disks correspond to local patches of highly responding units over a subcortical or cortical area. Such patterns can arise in various parts in the visual pathway. For simplicity, all such activity was presented to LISSOM through the retina, modeling activity arising anywhere in the input pathway to V1; activity arising within V1 itself and in other sources will be considered in greater detail in future work.

The properties of these input patterns were varied to model different types and sources of spontaneous activity. Binocularly correlated activity, such as that arising from ponto-geniculo-occipital waves in the brainstem during sleep [10] and possibly also spontaneous cortical activity [19], was modeled by presenting patterns at the same random location in both retinas (Fig. 2). The positions for uncorrelated inputs, simulating e.g. spontaneous waves of activity in the developing retina [6], were determined independently for the left and right retina.

3.1 OR and OD maps and their interactions

Using this model of spontaneous activity, we performed simulations to account for four primary qualitative properties of previsual maps: (i) smooth, binocularly matching OR maps [7], (ii) an OD map with significant left or right eye preferences and smooth transitions between them, (iii) orthogonal intersections of OR and OD maps [1], and (iv) OR pinwheel centers located within OD columns [1]. Each simulation varied the number of disks, their correlation, and/or the amount of noise. For each simulation, maps were developed by presenting 10,000 random inputs. The results suggest that a balance between correlated and uncorrelated inputs determines the joint development of OR and OD maps.

In the extreme case of very strong correlation between patterns in the two eyes, matching OR maps develop, but no OD map develops (Fig. 3a). Conversely, in the extreme case of very weak correlation between the disk patterns in each eye (even with perfect correlation of the noise), strong OD maps develop but the OR maps are completely different in each eye (Fig. 3b). In the latter (uncorrelated) case, the OR maps appear disorganized because neurons become nearly monocular, leaving only the left or right-eye areas in each OR map well-organized. Thus spatial correlation leads to matching OR maps, yet some degree of uncorrelation is required for OD maps to develop, making it difficult to satisfy properties (i-iv).

In contrast to these extreme cases, joint OR and OD maps satisfying all the properties (i-iv) were obtained if one disk, modeled by a circular Gaussian with $\sigma = 1.5$, was presented in each retina on a background of uncorrelated uniform random noise of brightness 1.0. The disks' brightness was varied randomly ($\mu = 1.2$, $\sigma = 0.5$). When correlation was decreased (from ratio 1:1 to 1:2) after 2000 iterations, ocular selectivity became clearer (Fig. 4c) (than keeping the original ratio). The OR maps still matched on both eyes (Fig. 4a,b), as they do in animals at eye-opening [7]. An overlay of OR and OD maps (Fig. 4d) shows that pinwheel centers tend to be located within OD columns, and orientation boundaries intersect the boundaries of OD stripes (indicated by white lines) at right angles. Similar patterns have been observed in macaque monkeys [1].

By comparing Fig. 4 with Fig. 3, it can be seen that the maps resulting from a balance between correlated and uncorrelated patterns are qualitatively different from maps trained on nearly uncorrelated patterns, where the OR maps in each eye are entirely unrelated, and from maps trained on nearly correlated patterns, where the OD map is weak or absent. In future work, we will develop techniques to make it computationally feasible to quantify these qualitative differences. However, the results above already indicate that realistic maps of single features develop over a wide range of input pattern types, whereas the joint development of realistically interacting maps depends on the balance and possibly the schedule of correlation in the input patterns.

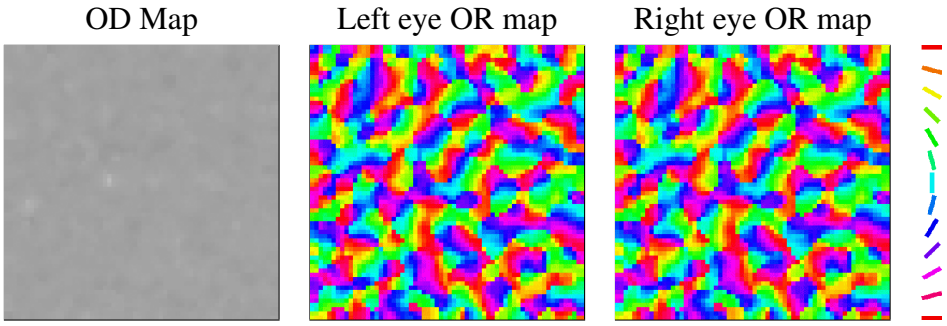

(a) Strong correlation, strong noise

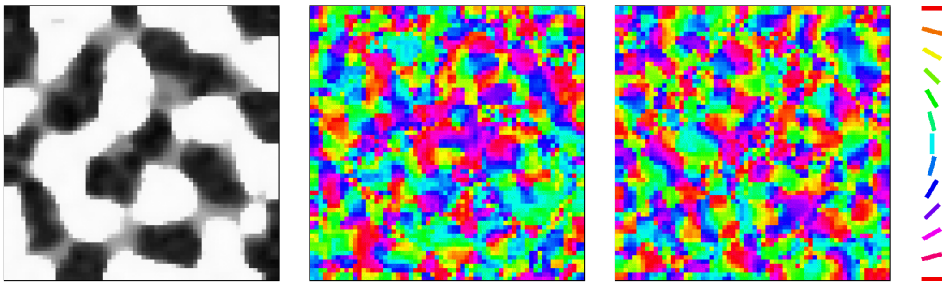

(b) Weak correlation

Fig. 3. Sample maps developed at two extreme levels of correlation in the input. In the OD map, eye preference from left to right is encoded in grayscale from white to black; in the OR map, the colors indicate orientation preference according to the key on the right. (a) The input pattern consisted of two pairs of disks (two disks in each eye), one uncorrelated with varying random brightness ($\mu = 0.0, \sigma = 1.0$), and one correlated with brightness level 1.0. In addition, there was correlated random noise with brightness level 1.0. Such strong correlation and strong noise yield matching, strong OR maps but no OD map. The weak uncorrelated input component is not enough to drive the development of the OD map. (b) The input consisted of one uncorrelated pair of disks (one disk in each eye) with varying, random brightness ($\mu = 0.5, \sigma = 1.0$) together with correlated noise that had random brightness ($\mu = 0.5, \sigma = 1.0$). Such a weak correlation yields very strong ocular dominance, but patchy and non-matching OR maps.

3.2 The effect of noise on selectivity in OD map development

In prior LISSOM experiments, we found that noiseless patterns that were uncorrelated between the eyes led to extremely strong selectivity in OD maps, similar to adult maps in strabismic animals [11]. Those results suggested that uncorrelated patterns like retinal waves could lead to the earliest maps being strabismic even in normal animals, which is possible but has not yet been established experimentally. However, the experiments in this section show that if the input patterns are made more realistic by including an overall pattern of noise, normal OD maps can develop even when the activity patterns are uncorrelated between the eyes.

In pairs of parallel experiments, during the first 1000 iterations the model was presented with uncorrelated disks of diameter 9 without noise, or with uncorrelated disks on a noisy background. In the following 9000 iterations, the input only con-

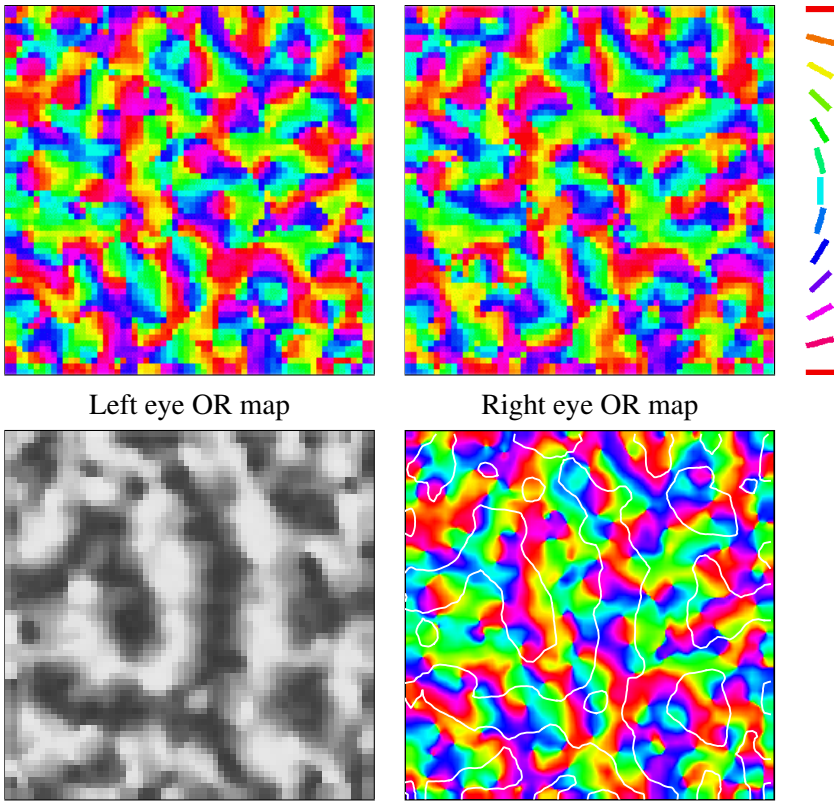

Fig. 4. Maps developed jointly with correlated and uncorrelated disks and uncorrelated random noise input patterns. Similar to biological maps, the OR maps match in both eyes, pinwheel centers tend to be located within OD columns, and OR and OD bands intersect orthogonally. The maps in the Joint OR/OD plot have been smoothed so that the intersection trends are easier to see.

sisted of disks without noise. The level of correlation was varied in the second phase throughout different pairs of experiments. In the first phase, the level of uncorrelation of the disks was always 1.0. The resulting OD maps suggest that noise, i.e. spontaneous spatially uncorrelated activity, prevents extreme ocular selectivity from forming (Fig. 5). The noise was only present during the first tenth of the development, yet the equalizing effect persists long after it has disappeared. For weaker noise, the effect was less pronounced. Even though the noise is also uncorrelated between the eyes, similar random activity patches occur in each eye by chance, which apparently provides enough coincident activity to prevent extreme values of ocular selectivity from developing. Since noise is likely to be part of the natural input patterns of developing neurons, these results suggest that even spatially uncorrelated patterns could lead to normal OD maps. In future research, similar techniques will be applied to study the role of noise in joint OR/OD maps like those in section 3.1. As part of this work, the parameter space will be explored more thoroughly in a simplified, computationally less expensive model.

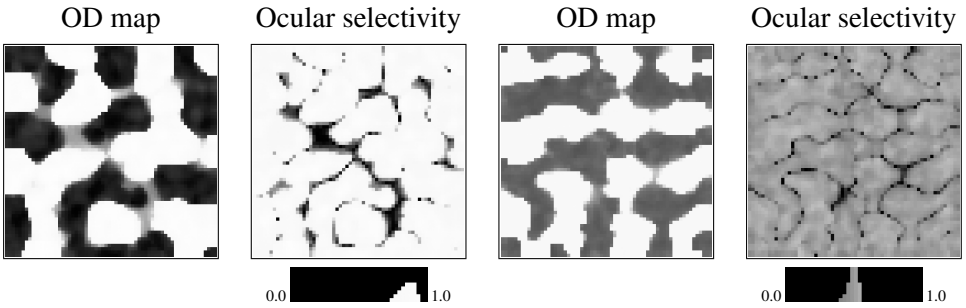

(a) Without noise in phase 1 (b) With noise in phase 1

Fig. 5. OD map and ocular selectivity after 10,000 iterations. The brightness of the selectivity map represents how strongly the neurons prefer their dominant eye, with white denoting maximum and black minimum. Below the selectivity map, a histogram illustrates the number of units at a particular level of selectivity. In a series of simulations, the input during the first 1,000 iterations consisted of disks with an uncorrelation level of 1.0, (a) without and (b) with uncorrelated random noise. In the following 9,000 iterations, noiseless disks were presented with an uncorrelation level between 0.4 and 1.0 (0.4 for the map shown in this figure). Without noise in the first 1,000 iterations, neurons become much more selective (indicated by brighter and darker colors in the OD map, and lighter colors in the selectivity map). Higher uncorrelation in the second phase does not change the maps significantly, indicating that the effect is robust.

4 Discussion and Future Work

The computational results suggest that care must be taken in interpreting biological experiments where only a portion of the total neural activity reaching the developing visual cortex is removed or modified. The long-lasting effects that were found for even brief periods of noisy inputs in the model support the concern of Stellwagen et al. [14] that the Crowley and Katz [4] enucleation experiments might not have removed all retinal activity. In the model, the removal of one type of activity can cause development to be driven by other, usually less dominant inputs. For instance, equally realistic joint OR/OD maps can develop from inputs with different spatial positions in each eye, as shown in this paper, or with matching positions but different strengths [11]. These latter patterns might correspond to residual LGN activities present after enucleation. Hence it is crucial to consider the entire ensemble of neural activities when determining the role of activity in development.

Similarly, the model suggests that it is important to measure multiple maps at once rather than maps of single types. In simulations, many types of activity patterns lead to realistic OD or OR maps alone [11], but a much smaller range yields realistic joint maps. For instance, small input position differences in each eye result in realistic individual maps, but unrealistic joint maps, with OR and OD boundaries aligned [11]. Shifting the balance of activities might have significant effects on the interactions that are not visible in the single-feature maps. Hence the interactions observed between previsual maps may provide a clue to what factors are driving their development in nature, and can guide future biological experiments.

As mentioned above, future work will include a thorough exploration of the parameter space comprising quantitative measures in a computationally less expensive model. Moreover, we will study the effects of differences in the timing of activity patterns in more detail. For instance, Crair et al. [3] found that orientation maps develop sooner for the contralateral eye than for the ipsilateral. This condition can be modeled using inputs presented later in one eye than the other. As with the previous work described above, it will be crucial to consider all the possible sources of activity involved, including potential shared cortical inputs prior to the arrival of ipsilateral projections.

5 Conclusion

The simulations suggest that realistic initial cortical maps can develop from patterns of spontaneous activity likely to occur prenatally in animals. However, details of the maps and the interactions between them depend critically on the balance of correlated and uncorrelated patterns and the amount of noise, even over brief periods. Future work will be needed to determine which sources of activity are the most crucial, focusing on differences between correlated and uncorrelated patterns.

Acknowledgments

This research was supported in part by the NSF under grant IIS-9811478 and by the NIMH under Human Brain Project grant 1R01-MH66991.

References

[1] G. G. Blasdel, Differential imaging of ocular dominance columns and orientation selectivity in monkey striate cortex, *Journal of Neuroscience* (1992), 12:3115–3138.

[2] G. G. Blasdel and G. Salama, Voltage-sensitive dyes reveal a modular organization in monkey striate cortex, *Nature* (1986), 321:579–585.

[3] M. C. Crair, D. C. Gillespie, and M. P. Stryker, The role of visual experience in the development of columns in cat visual cortex, *Science* (1998), 279:566–570.

[4] J. C. Crowley and L. C. Katz, Early development of ocular dominance columns, *Science* (2000), 290:1321–1324.

[5] E. Erwin and K. D. Miller, Correlation-based development of ocularly matched orientation and ocular dominance maps: Determination of required input activities, *Journal of Neuroscience* (1998), 18:9870–9895.

[6] M. B. Feller, D. P. Wellis, D. Stellwagen, F. S. Werblin, and C. J. Shatz, Requirement for cholinergic synaptic transmission in the propagation of spontaneous retinal waves, *Science* (1996), 272:1182–1187.

[7] I. Gödecke and T. Bonhoeffer, Development of identical orientation maps for two eyes without common visual experience, *Nature* (1996), 379:251–254.

[8] G. J. Goodhill and A. Cimponeriu, Analysis of the elastic net model applied to the formation of ocular dominance and orientation columns, *Network* (2000), 11:153–168.

[9] J. C. Horton and D. R. Hocking, An adult-like pattern of ocular dominance columns in striate cortex of newborn monkeys prior to visual experience, *Journal of Neuroscience* (1996), 16:1791–1807.

[10] G. A. Marks, H. P. Roffwarg, and J. P. Shaffery, Neuronal activity in the lateral geniculate nucleus associated with ponto-geniculo-occipital waves lacks lamina specificity, *Brain Research* (1999), 815:21–28.

[11] R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh, *Computational Maps in the Visual Cortex* (Springer, 2005), in press.

[12] C. Piepenbrock, H. Ritter, and K. Obermayer, The joint development of orientation and ocular dominance: Role of constraints, *Neural Computation* (1997), 9:959–970.

[13] H. Shouval, N. Intrator, and L. Cooper, Bcm network develops orientation selectivity and ocular dominance in natural scene environment, *Vision Research* (1997), 37:3339–3342.

[14] D. Stellwagen and C. J. Shatz, An instructive role for retinal waves in the development of retinogeniculate connectivity, *Neuron* (2002), 33:357–367.

[15] M. Stetter, E. W. Lang, and K. Obermayer, Synapse clustering can drive simultaneous on-off and ocular-dominance segregation in a model of area 17, in: *ICANN-97* (1997).

[16] N. V. Swindale, A model for the coordinated development of columnar systems in primate striate cortex, *Biological Cybernetics* (1992), 66:217–230.

[17] N. V. Swindale, The development of topography in the visual cortex: A review of models, *Network: Computation in Neural Systems* (1996), 7:161–247.

[18] R. O. L. Wong, Retinal waves and visual system development, *Annual Review of Neuroscience* (1999), 22:29–47.

[19] R. Yuste, D. A. Nelson, W. W. Rubin, and L. C. Katz, Neuronal domains in developing neocortex: Mechanisms of coactivation, *Neuron* (1995), 14:7–17.