

Risto Miikkulainen, James A. Bednar, Yoonsuck
Choe, and Joseph Sirosh

Computational Maps in the Visual Cortex

February 6, 2005

Springer

Berlin Heidelberg New York
Hong Kong London
Milan Paris Tokyo

Development of Maps and Connections

In this chapter, input-driven self-organization in the LISSOM model will be studied systematically, showing how features of visual input can become represented in the cortex. The focus will be on orientation, ocular dominance, and direction selectivity, where enough constraints exist from biology to constrain the models. The properties of the input patterns are first varied one dimension at a time, demonstrating that LISSOM develops maps much like those found in the cortex. In a scale-up simulation, all of these dimensions are then varied simultaneously, resulting in a map with joint representations for multiple features. Later chapters in Part II will demonstrate how the same self-organizing principles can operate in the adult, resulting in neural plasticity and visual aftereffects.

5.1 Biological Background

Computational models can be highly useful for gaining insight into biological mechanisms. Before they can be trusted, such models must be validated against biological data, to make sure their structures and processes are realistic. For maps and their connections, the biological data can be categorized in several different ways: (1) qualitative descriptions based on visual plots, vs. quantitative descriptions based on Fourier transforms, gradient calculations, and histograms (Erwin et al. 1995; Swindale 1996); (2) data on normal animals, vs. animals raised in abnormal visual environments and in sensory deprivation; and (3) data on single map features, vs. data on multiple features and their interactions. Map organization, receptive fields, and lateral connections were described qualitatively in Sections 2.1 and 2.2, focusing primarily on normal animals and single map features. This section will complement that review by describing quantitative results, studies on abnormal animals, and results on feature interactions.




Fig. 5.1. Fourier spectrum and gradient of the macaque orientation map. Plot (a) shows the two-dimensional Fourier spectrum of the map in Figure 2.4, calculated using methods described by Erwin et al. (1995) on orientation map data from Blasdel (1992b). In this and subsequent Fourier spectrum figures, the center represents the DC component and the mid-point of each edge $1/2$ of the highest possible spatial frequency of the image horizontally and vertically (i.e. the Nyquist frequency; Cover and Thomas 1991); the amplitude is represented in gray scale from white to black (low to high). As typically found in animal maps, the spectrum is ring shaped, indicating that the orientations repeat in all directions with a spatial frequency that corresponds to the radius of the ring. (b) The orientation gradient of the same map is plotted in gray scale from white to black (low to high; calculated from Blasdel 1992b as described in Appendix G.6). The high-gradient areas (dark ridges) correspond to fractures; the pinwheel centers are usually located at the ends of fractures. The gradient map makes the global arrangement of these features easy to characterize.

5.1.1 Quantitative Descriptions of Maps and Connections

Neurons in the visual cortex respond selectively to a number of input features such as location, orientation, eye of origin, and direction of movement, and preferences for these features vary systematically across the cortex. This organization can be visualized in maps, which can then be described qualitatively, as was done in Section 2.1.2. These same visualizations can also be analyzed numerically, measuring distributions of features and how they change across the map.

For example, the two-dimensional Fourier transform of an orientation map reveals how regular or periodic the map is, e.g. how often patches for each orientation are repeated across the surface. Biological maps have ring-shaped Fourier transforms (Figure 5.1a), revealing that *in all directions* map features repeat regularly, with an average periodicity corresponding to the radius of the ring. Orientation preference histograms complement Fourier transforms, measuring how many neurons prefer each orientation. As will be discussed in Chapter 9, animal maps are slightly biased toward vertical and horizontal orientations, reflecting the edge statistics of the visual environment.

The gradient of a cortical map measures how much each point in the map differs from its neighbors. Regions where map properties change sharply, such as pinwheel centers and fractures, have a large gradient. As an example, Figure 5.1b displays the

gradient of the orientation map in Figure 2.4. Fractures are seen as long ridges in this plot; pinwheel centers are often located at the ends of fractures, or appear as single dots. Although gradient plots are visually similar to selectivity plots (Figure 2.4*b*), it is important to realize that they are measuring two different properties. Low selectivity often coincides with high gradient in such plots, but this result may be an artifact of averaging the responses of several cells in a high-gradient area (Maldonado, Gödecke, Gray, and Bonhoeffer 1997). Gradient plots could thus be used to identify where selectivity measurements are unreliable.

Selectivity histograms can display useful information about gross map properties as well, but they are similarly affected by averaging. Estimates of selectivity obtained with different techniques, such as optical imaging and microelectrode recordings, differ widely (Maldonado et al. 1997). Further, as will be shown in Section 13.4, different types of training patterns result in different selectivity histograms in computational models. This measure is therefore less useful for comparing models with biological maps.

Histograms can also be used to quantify lateral connectivity patterns. As was described in Section 2.2, neurons tend to connect roughly symmetrically to their near neighbors, and to neurons farther away through patchy long-range connections. The long-range patches link primarily neurons with roughly collinear orientation preferences. One way to quantify such patterns is by measuring the angles between the orientation preferences of connected neurons. A histogram of how often various angles occur demonstrates that lateral connections indeed link neurons with similar preferences (Section 11.5.3).

5.1.2 Experimental Manipulation of Maps

As was reviewed in Sections 2.1.4 and 2.2.2, disrupting or changing the input patterns to the visual cortex during development can profoundly change the resulting maps and their connectivity. Of such manipulation studies, ocular dominance is perhaps the best known, and provides a detailed test case for computational models.

In the medical condition of strabismus (cross-eye), the eyes cannot focus on the same point in space; this condition can be induced experimentally in animals by cutting some of the eye muscles. As a result, each eye sees entirely different images, instead of the highly overlapping images in normal vision.

Under artificial strabismus, ocular dominance maps still develop, but their properties differ from normal maps (Figure 5.2). Strabismic maps have more sharply delineated ocular dominance areas, with stripes containing few connections from the opposite eye (Löwel 1994). The ocular dominance stripes are also significantly larger in the strabismic maps than in normal maps. Lateral connectivity patterns are also affected: Whereas in normal animals the lateral connections do not significantly favor one eye over the other (Bosking et al. 1997; Löwel and Singer 1992), in strabismic animals they become specific to each eye (Figure 5.2; Löwel and Singer 1992). These differences apparently result from a decrease in correlation between visual activity patterns between the two eyes.




Fig. 5.2. Normal vs. strabismic cat ocular dominance maps and lateral connections. These plots show corresponding 5 mm $\times$ 5 mm portions of the V1 ocular dominance maps from a normal cat (a) and from a cat raised with artificial strabismus (b). The maps were obtained using anatomical tracers, which result in categorical eye preferences (represented by light and dark areas instead of gray scale as in Figure 2.5). Both maps contain patches specific to each eye, but the patches are larger and more sharply delineated in the strabismic case. In (a), the green star indicates where fluorescent tracer was injected, and the green dots show where lateral retrograde transport took them. The lateral connection patterns do not significantly depend on the ocular dominance patterns. In (b), the red star and the green star (pointed by the arrow) mark two separate injection sites in right-eye columns (black). The lateral connections preferentially target neurons with the same eye preference (black patches, marked with red and green dots), and avoid neurons with the opposite eye preference (white). Each injection killed the nearby cells as a side effect, and therefore the ocular dominance and connection patterns are not visible in the areas surrounding the injections. Those areas are likely to be strongly connected to the neurons at the injection site. Detail of a figure by Löwel and Singer (1992), reprinted with permission, copyright 1992 by the American Association for the Advancement of Science.

Other input features have also been manipulated experimentally, with similar results. For instance, reducing or increasing the range of orientations seen by an animal can cause corresponding changes in the orientation map (Blakemore and Cooper 1970; Sengpiel et al. 1999). Such cases will be discussed more in detail in Section 8.1.

5.1.3 Interactions Between Multiple Maps

Although selectivity to each different input feature can be mapped independently as described above, all maps in V1 are overlaid onto the same set of neurons. Each neuron thus contributes to multiple maps, and the maps of different feature dimensions interact with each other. These interactions can be visualized in combined maps of




Fig. 5.3. Combined OR/OD map in the macaque. The orientation and ocular dominance maps shown separately in Figure 2.5 are overlaid in this plot. Color encodes OR preference (according to the key on top; example map features are outlined in white as in Figure 2.4), and the black outlines represent OD stripe boundaries, obtained as areas of high OD gradient. The features in the two maps are systematically organized. The OD boundaries intersect the OR boundaries of linear zones (long rectangle) at right angles, and rarely follow an OR boundary. Pinwheel centers (circles) are usually found well inside the OD stripes, and rarely near their boundaries. Note that, unlike the small patches seen in the cat OD maps (Figure 5.2), the OD patches in the monkey typically consist of long stripes (as also shown in Figure 2.5b). Reprinted with permission from Blasdel (1992b), copyright 1992 by the Society for Neuroscience; annotations added, OD contours replotted from data by Blasdel (1992a).

orientation, ocularity, and direction; to facilitate discussion of such maps, abbreviations OR, OD, and DR will be used for these feature dimensions in this and later chapters.

For instance, Figure 5.3 shows that the boundaries between ocular dominance stripes tend to intersect the boundaries between orientation patches at right angles (Blasdel 1992b). Orientation pinwheel centers are also typically found near the centers of ocular dominance stripes, and rarely intersect their boundaries.

Orientation and direction preferences also interact. Neurons have spatiotemporal receptive fields selective for both orientation and motion direction (Figure 5.4a; DeAngelis, Ohzawa, and Freeman 1993, 1995; Shmuel and Grinvald 1996). These RFs are formed by specific excitatory (ON) and inhibitory (OFF) subregions that vary over time, providing selectivity for both orientation and motion direction.

These preferences are arranged into maps of orientation and direction selectivity in the cortex. As an example, Figure 5.4 shows such maps in the ferret V1. Direction maps have a similar structure to orientation maps: Nearby neurons prefer similar directions, and the map has linear zones, pairs of pinwheels, saddle points, and fractures (Figure 5.4b; Weliky et al. 1996). Moreover, the direction map tends to be aligned with the orientation map, with orientation and direction preferences generally meeting at right angles (Figure 5.4c). Iso-orientation patches are also often subdivided into patches for each direction of motion.




Fig. 5.4. Spatiotemporal receptive fields, direction maps, and combined OR/DR maps in animals. In addition to orientation and eye of origin, neurons in V1 are selective for direction of motion. These spatial and motion preferences can be described as spatiotemporal RFs, representing the sequence of patterns that would most excite the neuron. (a) A sample such RF for a V1 cell from the cat, measured through microelectrode recording (DeAngelis, Ghose, Ohzawa, and Freeman 1999; reprinted with permission, copyright 1999 by the Society for Neuroscience; gray scale added). Sample RFs in the two-dimensional visual space at times 20, 60, 100, and 120 ms are shown on top, and a continuous integration of the RFs along the vertical (which is the preferred orientation of the neuron) is drawn in the bottom plane. The neuron's spatial preferences change systematically over time, giving it a spatiotemporal preference for a black vertical line moving horizontally to the right. (b) Spatial arrangement of such preferences in a $3.2 \text{ mm} \times 1.6 \text{ mm}$ area of ferret V1: Nearby neurons prefer similar directions in a manner similar to orientation maps (measured through optical imaging and displayed using the color arrow key on top; Weliky et al. 1996, reprinted with permission, copyright 1996 by Nature Publishing Group; annotations added and DR arrows removed by interpolation). Example map features are outlined in white as in Figure 2.4. (c) Interaction of direction preferences with the orientation map (Weliky et al. 1996; reprinted with permission, copyright 1996 by Nature Publishing Group; arrows changed from black to white). The $1.4 \text{ mm} \times 1.1 \text{ mm}$ subarea of V1 around the right edge of the square in (b) is colored according to orientation preference (using the color bar key above the plot). Each arrow points in the preferred direction, and its length indicates how selective the neuron is for that preference. Direction and orientation preferences tend to be perpendicular, and orientation patches are often subdivided for opposite directions of motion.

The lateral connections in normal animals follow mostly orientation preferences (Section 2.2.1). However, as described above, in strabismic animals eye preference instead becomes the dominant feature, suggesting that the connection patterns develop to reflect the most dominant features of the map. How less prominent input features, such as direction of motion, interact with the dominant patterns is not yet known.

The biological observations outlined in this section are constraints that computational models of the visual cortex must satisfy and explain. Considerable progress has indeed been made in understanding this data computationally, as will be reviewed next.

5.2 Computational Models

Computational models have been crucial for understanding how orientation preference and ocular dominance maps develop, and they have recently been used to gain insight into motion preference maps as well. Most such models are based on self-organizing maps, reviewed in Section 3.4 in the context of computational models above and below the map level. This section focuses specifically on map-level models of the visual cortex, outlining how their architectures and predictions differ from LISSOM's.

Computational map models range from more realistic to more abstract along several dimensions, including: (1) models that learn incrementally through individual image presentations, vs. models that represent long-term development as an abstract process; (2) models that include specific, patchy lateral connections, vs. models based on abstract lateral interaction functions; and (3) models that can process photographic images of natural stimuli, vs. models that work only with artificial image stimuli. Models built so far have represented some of these dimensions realistically and abstracted others; LISSOM is the first incremental model that develops specific, realistic patchy lateral connectivity from natural images. The review below is organized along the first two dimensions, separately identifying models that have been tested with natural images (see Erwin et al. 1995; Swindale 1996 for further comparisons).

5.2.1 Non-Incremental Models

Models that are based on abstract representations of developmental processes can be conceptually elegant and computationally efficient. On the other hand, it is not possible to account for the same level of detail as with models that learn from individual images. The two main non-incremental approaches are spectral models and correlation-based learning, both of which suggest that large-scale structures, such as orientation and ocular dominance patches, may arise as artifacts of neural processing rather than a principled way of representing visual input.

Spectral Models

Spectral models of maps do not include neurons or connections. The goal is simply to reproduce the patterns in biological maps, as opposed to explaining how maps contribute to visual function (Grossberg and Olson 1994; Niebur and Wörgötter 1993; Rojer and Schwartz 1990). It turns out that patterns similar to ocular dominance and orientation maps can be obtained by simply filtering two-dimensional random noise using a well-chosen function. For instance, ocular-dominance-like stripes can be produced by starting with an array of random numbers, convolving the array using a band-pass filter, and thresholding the result (Rojer and Schwartz 1990). Thus, spectral models suggest that orientation and ocular dominance patterns may not be functionally significant, arising only as an artifact of biological processes unrelated to information processing.

However, the maps derived in this way are only superficially similar to biological maps, and differ in several crucial respects (Swindale 1996). For instance, in the orientation maps produced of Rojer and Schwartz (1990) each orientation vector is perpendicular to the gradient of the map at that location, which is not seen in animal maps. It is not yet clear whether spectral models could be modified to overcome these difficulties. It is also not clear what specific biological processes could implement such highly abstracted computations, making the spectral models difficult to verify or refute experimentally.

Correlation-Based Learning Models

More closely tied to biological processes are the correlation-based learning (CBL) models (Erwin and Miller 1998; Linsker 1986a; Miller 1994; Miller et al. 1989; Tanaka 1990; Yuille, Kammen, and Cohen 1989). Models in this category rely on the assumption that the visual system is essentially linear. Under this assumption, the developmental process can be represented as a simple set of functions representing long-term correlations in the response to input patterns. This approximation speeds up the calculations considerably compared with incremental learning, and makes it possible to analyze the model mathematically.

Most CBL models focus on individual neurons or small groups of neurons, with the exception of those of Miller (1994), Miller et al. (1989), and Erwin and Miller (1998). The overall architecture of the CBL map models is similar to that of LIS-SOM, with a retina, LGN with ON-center and OFF-center neurons, and V1. Initially the afferent connections have random strength within a circular receptive field, and lateral interactions have a DoG profile. Unlike in LISSOM, however, lateral interactions are fixed and non-recurrent, and weight normalization is subtractive. Also, because the activation functions are linear, individual image presentations can be replaced with long-term averages of activity correlations. Instead of implementing Hebbian learning of individual input patterns, CBL models compute what Hebbian learning in a linear system would produce over many presentations.

Miller (1994) showed that such correlation-based learning results in maps of ocular dominance and of orientation. Monocular cells develop because subtractive normalization eventually leads to inputs from one eye becoming completely dominant,

with zero-strength connections from the other eye (Miller and MacKay 1994). These cells are grouped into ocular dominance stripes because lateral interactions cause nearby neurons to have similar responses. Orientation selectivity develops due to similar competition between ON and OFF-center inputs, but within a single cell's receptive field. With appropriate parameters and correlation functions, cells develop both ON and OFF subregions, making them selective for orientation. Lateral interactions then organize these preferences to orientation patches.

However, a number of key predictions of CBL models are inconsistent with recent biological data (Erwin et al. 1995): (1) The width of biological ocular dominance columns does not depend on the width of lateral interactions, but on the input correlations (Löwel 1994); (2) the RF shapes in CBL are typically only weakly selective for orientation, unlike in animal maps; (3) the Fourier spectrum of the CBL orientation map is concentrated around the origin, i.e. at low frequencies, instead of being shaped like a ring (Swindale 1996); (4) the subtractive normalization mechanism in the model forces synapses either to their maximum weight or zero, which is biologically unrealistic; and (5) the model cannot account for adult plasticity and dynamic reorganization (as reported by e.g. Gilbert 1998; Pettet and Gilbert 1992); once the synapses reach the extreme values, further adaptation is very difficult.

In general, CBL models suggest that large-scale features, such as orientation and ocular dominance maps, develop primarily as artifacts of neural connectivity patterns. In contrast, in the incremental learning approach they emerge as a principled way of representing visual input. As a result, incremental models are more complex and computationally intensive, but can produce more realistic results and explain a broader range of phenomena.

5.2.2 Incremental Models with Fixed Lateral Connections

A large number of incremental models of the visual cortex have been proposed in the last 30 years; almost all of them are based on fixed, isotropic lateral connectivity. As was described in Section 3.4.1, the earliest model of this type was built by von der Malsburg (1973); similar models have since then been developed using the SOM algorithm, the elastic net algorithm, and similar architectures. Von der Malsburg's model developed oriented receptive fields and pinwheels (before such patterns had been found experimentally), demonstrating the basic computational processes underlying development in the visual cortex. Later models have shown how orientation, ocular dominance, and direction maps can form, and also how receptive fields selective for each of these dimensions develop (Barrow and Bray 1992; Dong and Hopfield 1992; Durbin and Mitchison 1990; Elliott, Howarth, and Shadbolt 1996; Elliott and Shadbolt 1999; Farkas and Miikkulainen 1999; Goodhill 1993; Grossberg 1976; Grossberg and Olson 1994; Grossberg and Seitz 2003; Hyvärinen and Hoyer 2001; Miyashita, Kim, and Tanaka 1997; Miyashita and Tanaka 1992; Obermayer et al. 1990d; Obermayer, Sejnowski, and Blasdel 1995; Olson and Grossberg 1998; Osan and Ermentrout 2002; Piepenbrock and Obermayer 2002; Shouino and Kurata 2001; Shouval, Intrator, and Cooper 1997; Stetter, Müller, and Lang 1994; Swindale 1992; Wimbauer, Wensch, van Hemmen, and Miller 1997b). So far, the SOM and related

models have produced the best description of such maps, measured by analytical comparisons with experimentally observed maps (Swindale 1996).

Nearly all of these models have focused only on a single feature dimension (e.g. orientation or ocular dominance). A few models have included multiple dimensions in the same simulation, such as OR and OD (Erwin and Miller 1998; Goodhill and Cimoneriu 2000; Grossberg and Seitz 2003; Obermayer et al. 1995; Olson and Grossberg 1998; Osan and Ermentrout 2002) or OR and DR (Blais, Cooper, and Shouval 2000; Farkas and Mikkilainen 1999; Miyashita et al. 1997; Shouval and Kurata 2001; Wimbauer et al. 1997b). Simulating multiple dimensions at once ensures that parameters are not tuned for only one feature, and interactions between dimensions allow validating the model more extensively. To our knowledge, no published model has developed joint OR, OD, and DR maps; the first study of such maps will be presented in Section 5.6.

Most of these models have been tested only with artificial inputs such as oriented Gaussian patterns or pure random noise, which can be strictly controlled to obtain strong results. However, many phenomena studied in later chapters of this book depend crucially on properties of natural images. Such images have wide ranges of contrast and large areas of activation, and it is therefore necessary to include the processing steps in the ON and OFF channels of the retina and the LGN into the model. Although a few such models have been proposed (Barrow and Bray 1992; Hyvärinen and Hoyer 2001; Shouval et al. 1997; Weber 2001), most models do not include these requisite processing stages.

Crucially, it is not yet clear whether the output from the existing natural image models actually preserves the essential information in the images. So far, the models have been used only to investigate how the map forms, not how it performs visual processing. This issue is important because several of the approximations common in models with fixed lateral connectivity, such as choosing a single winning location for adaptation across the entire cortex, do not scale up to natural images with a realistic size. Development is often driven by only a few strong features in each image, but visual performance requires multiple features of different contrasts to be represented simultaneously. One important goal of the experiments in Chapter 10 is to verify that LISSOM indeed preserves the information necessary for higher levels of processing, such as face perception areas. It is not known whether other models can be used for such studies.

5.2.3 Incremental Models with Modifiable Lateral Connections

The early map models were developed before specific, patchy lateral connections were discovered in the visual cortex, and assumed that lateral interactions would have a fixed, uniform shape. Nearly all subsequent models have relied on similar assumptions, primarily because it is computationally very expensive to store and simulate individual connections. However, it is crucial to include such connections to account for a number of key experimental results (as reviewed in Section 2.2), and such connections are also important for information processing (as will be demonstrated in Chapter 14). Explaining how the lateral connections self-organize into the

characteristic patchy patterns is a crucial part of understanding the development and function of topographic maps.

After LISSOM was first introduced a decade ago, several models with modifiable lateral connections have been developed (Alexander et al. 2004; Bartsch and van Hemmen 2001; Bray and Barrow 1996; Burger and Lang 1999; Weber 2001). Of these, the models by Alexander et al. (2004) and Burger and Lang (1999) form patchy long-range lateral connections, but the rest do not. The Alexander et al. (2004) model relies on abstract binary input patterns like those of von der Malsburg (1973), and it is not clear how to extend it to support gray-scale and natural-image inputs, or to develop multi-lobed receptive fields. The Burger and Lang (1999) model is very similar to LISSOM, including ON and OFF channels, modifiable lateral connections, and support for natural images. Mathematically, their most significant difference from LISSOM is the activation function: The Burger and Lang model includes a linear function (sum of all inputs), whereas LISSOM's activation function is nonlinear, and also iteratively incorporates lateral influences. The nonlinearities make the LISSOM map more selective for input features and less dependent on contrast, which is crucial for simulating responses to realistic inputs. The Burger and Lang model is also based on a different method for normalizing the afferent weights, although both methods should lead to similar results. So far, only a small range of cortical phenomena have been investigated with the Burger and Lang model, but with certain modifications it should be possible to obtain results similar to LISSOM's. Other models would need to be extended significantly to study the same range of phenomena, e.g. by including specific lateral connections and mechanisms for processing natural images.

Although considerable progress has been made in understanding the biological data reviewed in Section 5.1, most of the models have focused on a limited set of phenomena. LISSOM is the first where they are all brought together under a single common principle of adaptation and organization. In this chapter, LISSOM is first evaluated as a model of each separate abstract input feature dimension, and then as a joint model that self-organizes to represent multiple features in natural images together.

5.3 Orientation Maps

In this section, a LISSOM model of how orientation maps and lateral connections develop based on input-driven self-organization is presented. The model described in Chapter 4 is extended to oriented input and shown to develop oriented receptive fields and a global orientation map. The resulting maps and lateral connections are then analyzed numerically. In later sections, ocular dominance and direction maps are studied separately, and combined maps of all three input feature dimensions are developed both from abstract and natural image inputs.

5.3.1 Method

The LISSOM model of orientation maps is otherwise identical to that described in Chapter 4 (Figure 4.1) except different training inputs are used. In the main sim-




Fig. 5.5. Initial V1 afferent and lateral weights. These plots show the initial weights of the V1 neuron at location (79,68) in the 142×142 V1 map. As in Figure 4.3, each set of weights is outlined in black and plotted in gray scale on the neural region from which they originate; however, the padding in the LGN is omitted so that the area of (a) corresponds to that of Figure 5.6b. The neuron itself is marked with a small white square in (b) and (c). The afferent RFs were initially random, as shown in (a); the ON and OFF channel weights were identical for each neuron. Plots (b) and (c) display the lateral weights of this neuron; initially they had a Gaussian profile. Later figures will show how these connections become selective and patchy through self-organization.

ulation, the inputs consist of images of elongated Gaussian spots, instead of the unoriented Gaussians of Equation 4.2. The activity of each retinal receptor χ_{xy} is calculated according to

$$\chi_{xy} = \max_k \exp \left(- \frac{[(x-x_{c,k}) \cos(\phi) - (y-y_{c,k}) \sin(\phi)]^2}{\sigma_a^2} - \frac{[(x-x_{c,k}) \sin(\phi) + (y-y_{c,k}) \cos(\phi)]^2}{\sigma_b^2} \right), \quad (5.1)$$

where σ_a and σ_b determine the width along the major and minor axes of the Gaussian, and ϕ its orientation, chosen randomly from the uniform distribution in the range $0 \leq \phi < \pi$. Such inputs are abstractions of elongated features in images and in spontaneous neural activity and allow demonstrating the model properties clearly. In Section 5.3.5, the model is trained with a range of other patterns, and the resulting differences in organization analyzed.

The self-organization proceeds as described in Chapter 4. The way the connections are initialized does not have a large effect, as long as they are roughly isotropic. Since self-organization of scattered RF centers and random initial lateral weights was already demonstrated in Section 4.5, to make the results easier to interpret the afferent connections are initially random around topographically ordered centers, and the lateral connections have initially a Gaussian profile (Figure 5.5). At each input presentation, multiple Gaussian spots are presented on the retina at random orientations and random, spatially separated locations (Figure 5.6a). The activation propagates through the LGN and the afferent connections of V1 and produces an initial response in the V1 network (Figure 5.6c). The initial response is typically diffuse and widespread, but by recurrent lateral excitation and inhibition, it rapidly evolves




Fig. 5.6. Example input and response. A sample input on the retina, the LGN response, and the initial and settled cortical response before and after training are shown as in Figure 4.4, except the padding in retina and LGN is omitted so that all plots represent the same retinal area. To train the orientation map, two oriented Gaussians were drawn with random orientations and random, spatially separated locations on the retina (a). As discussed in Appendix A.4, while more than two spots could be used, they are too large to be distributed uniformly on the small retina. The LGN responses are plotted in (b) by subtracting the OFF cell responses from the ON. The LGN responds strongly to the edges of the oriented Gaussians. Initially, the responses of the V1 map are similar for all orientations (c and d). After 10,000 input presentations, the V1 response extends along the orientation of the stimulus, and is patchy because neurons that prefer similar positions but different orientations do not respond (e and f). An animated demo of the map response can be seen at <http://computationalmaps.org>.

into stable, focused patches of activity (Figure 5.6d). After the activity has settled, the strength of each synaptic connection is updated. A new set of oriented Gaussians is then generated at random positions and orientations, and the process repeats for 10,000 iterations. Appendix A.5 lists the details of the simulation parameters. Small variations of the parameters result in roughly similar maps; a representative example is analyzed in detail in the sections that follow.

5.3.2 Receptive Fields and Orientation Maps

Initially, the activation patterns are very similar even for different orientations, allowing the map to develop a global retinotopic order. Gradually, oriented receptive fields start to form, and the lateral connections start to follow receptive field properties. As Figure 5.7 shows, the final RFs and lateral connections are very similar to those found in biology (Section 2.1.2; Bosking et al. 1997; Hubel and Wiesel 1965, 1968; Sincich and Blasdel 2001). Afferent RFs are Gabor-shaped with separate ON and OFF lobes, making them strongly selective for orientation. Lateral connections are patchy, and the long-range connections originate from neurons with similar orientation preferences. Across the map (Figure 5.8), the RFs have a variety of shapes; most are highly selective for inputs of a particular orientation, and others are unselective. Lateral connections tend to follow the RF shape, linking neurons with similar RFs.

The global organization of the RFs can be visualized similarly to biological orientation maps, by labeling each neuron by the preferred angle and degree of selectivity for inputs at that angle. To determine these labels, responses to sine gratings of various orientations were measured and recorded (as described in Appendices G.1.3




Fig. 5.7. Self-organized V1 afferent and lateral weights. The weights of the neuron in Figure 5.5 are shown after self-organization. In (a), the OFF weights were subtracted from the ON weights, as in Figure 4.6c. This neuron prefers a line oriented at 60° , i.e. diagonal from bottom left to top right, and responds most strongly to a white line overlapping the light portion of its RF, surrounded by black areas overlapping the dark portions. Other neurons developed similar RFs with different preferred orientations (Figure 5.8). This type of RF structure is commonly seen in biological V1 neurons (Figure 2.2d; Hubel and Wiesel 1962, 1968). In the lateral weight figures (b, c, and other later such figures), the following convention is used: The hue (i.e. color) represents the orientation preference of the source neuron, according to the key along the top. The saturation of the color (i.e. its fullness, or intensity) represents how selective the source neuron is for this orientation; unselective neurons are shown in gray. The value of the color (i.e. its brightness) indicates the strength of the connection, with nonexistent or zero-weight connections shown as white. The jagged black outline traces the original lateral connections, and a small white square (in c) identifies the neuron itself. Using such a scale, plot (b) displays the lateral excitatory weights of this neuron. All connected neurons are strongly colored blue or purple, i.e. orientations similar to the orientation preference of this neuron. The lateral inhibitory weights are plotted in (c). After self-organization and connection pruning, only connections from neurons with similar orientations remain, and they are extended along the preferred orientation of this neuron. The connection pattern is patchy, because connections from neurons with opposite preferences are weak or have been pruned away entirely. Such patchy, orientation-specific connection patterns are also seen in biological V1 neurons (Figure 2.7; Bosking et al. 1997; Sincich and Blasdel 2001).

and G.3. These values were then used to plot the orientation map shown in Figure 5.9. Initially, all the afferent weights are random. As a result, the orientation preferences of the RFs appear random and unselective (top row). As self-organization progresses and afferent weights develop oriented receptive fields, a complex orientation map develops (bottom row). Even with such abstract inputs, the map is a good match to those measured in animals, and contains structures such as linear zones, pinwheels that often occur in matched pairs, saddle points and fractures (compare Figure 5.9 with Figure 2.4). The maps can be further analyzed with numerical techniques, showing quantitatively the same structures as the primary visual cortex.




Fig. 5.8. Self-organized afferent and lateral weights across V1. This plot shows the range of afferent and lateral weights developed by the neurons in the orientation map, by plotting them for every 12th neuron horizontally and vertically, using the conventions introduced in Figure 5.7. (a) A number of two- and three-lobed receptive fields exist with strong orientation preferences. Some neurons, however, have ring-shaped RFs and respond to all directions equally. (b) These neurons receive lateral inhibitory connections from all nearby neurons, but from distant neurons only if they have similar OR preferences and are located along the preferred orientation. The lateral excitatory connections of each neuron (not shown) come from all nearby neighbors, and thus are all nearly circular.

5.3.3 Analysis of the Orientation Maps

The LISSOM orientation map was analyzed using the numerical techniques described for biological maps in Section 5.1. The histogram of orientation preferences is flat (Figure 5.9d), showing that the architecture does not have biases for any particular orientations (as it would, for instance, if it had square receptive fields). It therefore reflects the uniform distribution of orientations during training. If LISSOM was instead trained with natural images, where certain angles are over-represented, the resulting histogram would be more similar to those found in animals (Sections 5.3.5 and 9.3).

The Fourier spectrum of the orientation map has the typical ring-shaped structure of biological maps, indicating that the orientation patches repeat at regular intervals in all directions (Figure 5.10a). Because the network has fewer units than the animal optical imaging data have pixels, the plot is fuzzier, but the overall shape is similar. The map gradient is also similar to that of cortical maps (Figure 5.10b). Discontinuities are represented by high gradient: Pinwheel centers are seen as high points and fractures as more linear ridges connecting pinwheel centers. In LISSOM, high gradient generally coincides with low selectivity, because neurons self-organize to




Fig. 5.9. Self-organization of the orientation map. The orientation preference and selectivity of each neuron was computed before (top row) and after self-organization (bottom row). The preferences are color coded and selectivity represented in gray scale as in Figure 2.4. (a) The orientation preferences were initially random, but over self-organization, the network developed a smoothly varying orientation map. The map contains all the features found in animal maps, such as linear zones, pairs of pinwheels, saddle points, and fractures (outlined as in Figure 2.4). (b) Before self-organization, the neurons are unselective (i.e. dark), but nearly all of the self-organized neurons are highly selective (light). (c) Overlaying the orientation and selectivity plots (by representing selectivity with color saturation as in Figure 5.7) shows that regions of low selectivity in the self-organized map tend to occur near pinwheel centers and along fractures. (d) Histograms of the number of neurons preferring each orientation (OR H) are essentially flat because the initial weight patterns were random, the training inputs included all orientations equally, and LISSOM does not have artifacts that would bias its preferences. These plots show that LISSOM can develop biologically realistic orientation maps through self-organization based on abstract input patterns. An animated demo of the self-organizing process can be seen at <http://computationalmaps.org>.

respond together with their neighbors. Whether this is true of animal maps as well is still controversial (as was discussed in Section 5.1.1).

The orientation discontinuities also affect the retinotopic mapping from the retina to V1 (Figure 5.11). The large-scale organization corresponds to the retina: for example, neurons in the upper left of the cortex respond to activity in the upper left of the retina. On small scales, however, this mapping is distorted because the orientation map represents both position and orientation smoothly across the same surface. Such distortions also occur in cat orientation maps, where orientation gradient is found to




Fig. 5.10. Fourier spectrum and gradient of the orientation map. (a) The Fourier spectrum is ring shaped as it is for biological maps (cf. Figure 5.1a), indicating that all directions are represented at regular intervals. (b) The global arrangement of high- and low-gradient areas is similar to that in biological maps: Regions with high gradient coincide with discontinuities such as pinwheel centers and fractures, and the fractures tend to connect the pinwheel centers (cf. Figure 5.1b).

correlate with distance between RF centers (Das and Gilbert 1997). Such a correlation exists in LISSOM maps as well, although it is weaker and the relationship is more complex.

The above analytical comparisons demonstrate that the afferent structures in LISSOM essentially replicate the afferent organization in the cortex. The next section will show that the patterns of lateral connections also compare well to those observed in biology.

5.3.4 Lateral Connections

The lateral connection weights self-organize at the same time as the orientation map. Initially, the connections are spread over long distances and cover a substantial part of the network (Figure 5.5). As the lateral weights self-organize, the connections between uncorrelated regions grow weaker, and after pruning only the strongest connections remain (Figure 5.7). The surviving connections of highly orientation-selective cells, such as the one illustrated in Figure 5.12a, link areas of similar orientation preference, and avoid neurons with different orientation preferences. Furthermore, the connection patterns are elongated along the direction in the map that corresponds to the neuron's preferred stimulus orientation. This organization reflects the activity correlations caused by the elongated Gaussian input pattern: Such a stimulus activates primarily those neurons that are tuned to the same orientation as the stimulus, and located along its length. At locations such as fractures, where a cell is




Fig. 5.11. Retinotopic organization of the orientation map. The center of gravity of the afferent weights of every second neuron was calculated and plotted in the retinal space, and those of neighboring neurons connected with lines (as in Figure 4.8). The overall organization of the map is an evenly spaced grid with local distortions. These distortions result from mapping both orientation and retinal position smoothly into the same two-dimensional surface; such distortions have been found experimentally on animal maps as well (Das and Gilbert 1997).

sandwiched between two orientation patches of very different orientation preference, the lateral connections are elongated along the two directions preferred by the two adjacent patches (Figure 5.12*d*). The lateral connections of unselective cells, such as those at pinwheel centers, come from all orientations around the cell (Figure 5.12*b*). Connections at saddle points are similar to those at fractures, in that they include the two orientations of the saddle, but they also include intermediate orientations that typically match the orientation preference of the saddle neuron itself (Figure 5.12*c*). Thus, the pattern of lateral connections of each neuron closely follows the global organization of receptive fields, and represents the long-term activity correlations over large areas of the network.

These results were originally discovered independently in the LISSOM model (Sirosh et al. 1996a), and some of them have already been confirmed in recent neurobiological experiments. In the iso-orientation patches of the tree-shrew cortex, horizontal connections were found to be distributed anisotropically, extending farther and giving rise to more terminals along the preferred orientation axis of the neuron in visual space (Bosking et al. 1997; Fitzpatrick et al. 1994; Sincich and Blasdel 2001). Most of these terminals also connected to cells with the same orientation preference. The connection patterns at pinwheel centers, saddle points, and fractures have not been studied experimentally so far; the LISSOM model predicts that they will have unselective, broad unimodal, and biaxial distributions, respectively.




Fig. 5.12. Long-range lateral connections in the orientation map. The lateral inhibitory connection weights of four sample neurons from the marked regions in Figure 5.9 are shown in the top row, situated in the orientation map as shown in the bottom row. The small white square in both figures identifies the neuron; the black outline on top indicates the extent of these connections before self-organization, and the white outline on the map plot shows their extent after self-organization and pruning. On top, the color coding represents the connected neuron's orientation, selectivity, and connection strength, as in Figure 5.7; the map encodes orientation and selectivity as in Figure 5.9c. The connection histogram (CH) in the middle shows how many connections come from neurons of each orientation. For every neuron, the strongest connections originate from the neuron's nearby neighbors, as indicated by the large, bright central area in each weight plot. The long-range connection patterns differ depending on where the neuron is located in the orientation map. (a) Neurons in the middle of an OR patch receive connections from neurons with similar preferences, aligned along the orientation preference of the neuron (for this neuron, about 65° , i.e. blue). (b) At pinwheel centers, the connections come from all directions and orientations and are nearly isotropic. The histogram is nearly flat, with small peaks near orientations that happen to be overrepresented in the pinwheel. (c) Connections at saddle points extend along the two orientations of the saddle, in this case red (0°) and blue (65°). The neuron also makes connections with intermediate orientations and directions; these connections match its own OR preference (30° , purple), and result in one broad peak in the histogram. (The connections of this neuron are cut off along the bottom because it is located near the bottom of the map.) (d) Connections of neurons at fractures are also elongated along the two directions of the neighboring orientation patches. The neuron plotted in (d) is on a fracture between yellow-green (130°) and blue-purple (40°), and makes connections with both of these orientations. In contrast to saddle points, it does not connect with intermediate orientations and directions, resulting in two distinct peaks in the orientation histogram. While the connection patterns in iso-orientation patches have already been confirmed in biology, the patterns at the other map features are predictions for future experiments.

5.3.5 Effect of Input Types

The preceding analysis focused on a single LISSOM model trained with oriented Gaussian patterns, which made the results clear and unambiguous. In a series of simulations, the model was trained with other patterns that have been hypothesized to contribute to orientation selectivity, including spontaneous neural activity and natural visual images. Spontaneous activity was modeled after retinal waves and represented as noisy disk patterns (as will be described in detail in Section 9.2.1). Natural images included retina-size closeups of natural objects and landscapes closely matched with natural input (from a dataset by Shouval, Intrator, Law, and Cooper 1996, described in Figure 8.4 and Section 9.3.1). The details of these simulations are listed in Appendix A.5.

The orientation maps that result from each of these training patterns are compared in Figure 5.13. In each case, orientation maps develop, but the maps and RFs are well ordered only for patterns that contain spatial structure. The Fourier transform of these maps is ring-shaped, as it is of animal orientation maps. Oriented Gaussians are not required, as long as the inputs have edges that produce elongated activity patterns at the LGN level. For instance, realistic maps develop from both natural images and retinal waves. Importantly, the simulations with disks and noisy disks demonstrate for the first time that full-fledged orientation maps can develop from large, unoriented activity patterns, which has been believed difficult to achieve (Miller 1994).

The receptive field types that develop depend strongly on the input patterns. For instance, if all patterns are brighter than their surround (the row labeled “Gaussians”), most of the resulting RFs have an ON-center with two flanking OFF lobes. If also patterns that are darker than the surround are included, both ON and OFF-center RFs develop (“Plus/Minus”). Note that in either case, the background illumination is not important, because the LGN responds only to brightness differences, not absolute levels. With disks, most RFs have two lobes, because the input contained edges but no thin lines. A variety of RF types develop in simulations with noisy disks and natural images, reflecting the wide variety of input patterns seen during training.

Orientation maps develop even with random noise inputs, because even random patterns have local clusters that are brighter or darker than their surround. These clusters lead to patches of activity in the ON channel adjacent to patches in the OFF channel, and orientation-selective neurons develop. However, the resulting map is not well ordered: Many neurons are only weakly selective and the RFs do not have well-developed profiles. The Fourier transform is disk-shaped, indicating that the map consists of orientation patches of all sizes, instead of a largely uniform size seen in animal maps and in maps formed with other input types. The conclusion is that spatial structure is necessary for realistic orientation maps to form.

Overall, these results suggest that realistic orientation maps can develop based on a wide variety of spatially coherent stimuli, and that the choice of these stimuli more strongly affects the RFs than the maps. Thus, LISSOM predicts that the types of RFs observed in different species are at least partially due to the patterns the animals see during development.




Fig. 5.13. Effect of training patterns on orientation maps. In this and later similar figures, the rows represent different self-organization experiments. Each row typically shows a sample retinal activation, the LGN response to that activation, final receptive fields (ON–OFF) of sample neurons, their lateral inhibitory connections (LIs), the orientation preference and selectivity map, the orientation preference histogram, and the fast Fourier transform (FFT) of the orientation preferences. The RFs and LIs are drawn to a smaller scale than LGN and V1. For clarity, most OR models are based on abstract input patterns like the oriented Gaussians in the top two rows. However, OR maps develop robustly with a wide variety of input patterns, including large circular patterns (middle rows) and natural images (second row from the bottom; image from a dataset by Shouval et al. 1996, 1997). Maps develop even with random noise (bottom row), although such maps are relatively unselective and the RFs do not have realistic profiles. Spatial structure is therefore necessary in LISSOM for biologically realistic maps to form.

In Chapter 9, these results are put together into a detailed model of how the orientation maps develop prenatally and postnatally in animals, including how internally generated activity like retinal and PGO waves and natural images may each contribute. The more detailed model also predicts what RF types are likely to be found in newborns, and how those will differ from adult RFs.

5.4 Ocular Dominance Maps

In this section, a second retina will be included in the LISSOM model, and it will be self-organized using unoriented Gaussian-shaped inputs on both retinas. Under these conditions, the V1 neurons develop either binocular receptive fields, or receptive fields with preference for one of the two eyes. The preferences are arranged into global maps of ocular dominance, with stripes preferring the left eye alternating with stripes preferring the right. Simultaneously, the lateral connections organize into patterns that reflect the correlations in the input. If the inputs are uncorrelated, the model replicates biological data on strabismic animals. The model is also used to study the effect of disparity more generally on ocular dominance maps.

5.4.1 Method

The LISSOM model of ocular dominance is otherwise identical to the orientation model from the previous section, except the V1 receives inputs from two retinas (Figure 5.14). The ON and OFF channels are set up the same way for both retinas, and the V1 neurons receive afferent connections from both channels through local receptive fields and topographically ordered RF centers (see Appendix A.1 for the model equations). The initial values of the afferent weights are random and identical for both eyes, to show that self-organization is not driven by initial weight differences.

Ocular dominance in LISSOM develops based on differences in activity patterns between the two eyes. Such differences are a likely source for ocular dominance in animals as well, although it is not yet known what types of activity patterns are most important for normal OD development. Possible candidates include spontaneous retinal waves that occur independently in each eye, correlated LGN activity originating in the brainstem (e.g. during sleep), position or brightness differences of corresponding visual image features in each eye, or combinations of all these factors. As an abstraction of a variety of such factors, brightness differences between matching patterns in each eye will be used in the main LISSOM experiment. In Section 5.4.4, more indirect sources such as small position differences will be evaluated as well.

The training inputs consisted of unoriented (circular) Gaussian spots of Equation 4.2, multiplied by a brightness factor s_b . Two spots were drawn in each eye in randomly chosen spatially separate locations. In the normal case these locations were constrained to be the same in the two eyes, and in the strabismic case they were independent in each eye. The brightness s_b of each spot in the left eye was chosen randomly at each iteration from the range $[0..1]$, and the brightness of the




Fig. 5.14. LISSOM model of ocular dominance. The architecture is otherwise identical to that of the LISSOM OR network, except now there are two retinas, leading to two ON and OFF channels. Each V1 neuron receives afferents from corresponding positions on the ON and OFF LGN sheets for each eye.

corresponding spot in the right eye was computed as $1.0 - s_b$ so that the total input activation remained constant. With such input, the LGN responds to both the left and the right spot, but more strongly to the brighter one. V1 responds as it would to a bright spot in one eye, but will learn stronger connections from the eye with the brighter input. The network was trained for 10,000 iterations like the orientation map of Section 5.3. The rest of the simulation parameters are described in Appendix A.6.

The OD maps that emerged in normal and strabismic self-organization are described in the next two subsections, and the effect of input disparity, i.e. the degree to which the inputs are independent in the two eyes, is analyzed in the last subsection.

5.4.2 Normal Ocular Dominance Maps

Through self-organization, the network developed eye-specific receptive fields and responses (Figure 5.15). The RFs are not significantly selective for orientation, because the training inputs were small and unoriented. Different neurons prefer one eye over the other to different degrees, but as in animals, nearly all neurons are binocular to some degree.

The global arrangement of the eye preferences was visualized by recording the response of each neuron to patterns presented in each eye individually, as described in Appendix G.4. Figure 5.16 shows the resulting ocular dominance map, consisting of alternating stripes in irregular patterns across the network. Selectivity in that map measures how strongly the neurons favor inputs from one eye. The most strongly binocular neurons are found near the OD stripe boundaries and the most strongly monocular ones are at the center of such stripes. Overall, most neurons are binocular to some degree. Similar graded functional patterns of OD are seen e.g. in the macaque monkey using optical imaging techniques (Figure 2.5).




Fig. 5.15. Self-organization of afferent weights into OD receptive fields. (a) The afferent weights of a sample neuron, located as shown in Figure 5.17a, are plotted before self-organization (as in Figure 5.5). Initially these weights are random and identical for both eyes and both channels in each eye. (b) The final receptive fields of the same neuron are visualized for each eye by subtracting the OFF weights from the ON weights (as in Figure 5.7). Over the course of self-organization, most neurons develop a preference for one eye or the other, although they retain significant connections from both eyes. Many of this neuron's connections from the left eye are weak (indicated by medium gray), so it responds more strongly to input in the right eye. (c) On the other hand, neurons near the OD stripe boundaries, like the one in Figure 5.17b, become strongly binocular, with smooth, isotropic RFs that are nearly identical in each eye. The ocular dominance stripes shown in Figure 5.16 are based on such subtle eye preferences, as they are in animal OD maps.




Fig. 5.16. Self-organized ocular dominance map. Light areas in (a) indicate neurons that prefer the left eye, dark areas those that prefer the right eye, and medium gray indicates no net preference. The histogram (b) shows how these preferences are distributed, with left monocular neurons at the top, binocular neurons in the middle, and right monocular neurons at the bottom. Most neurons are binocular, slightly preferring one eye or another, as they do in animals (Figure 2.5). Plot (c) illustrates how selective the neurons are for ocularity, with light areas indicating monocular neurons and dark areas those that are binocular. Less selective regions fall between ocular dominance stripes, as in animal maps.

The stripes form based on the push-pull effect of the lateral connections: Local excitation ensures that nearby neurons will respond to similar stimuli and thus have correlated activity (causing eye-specific regions to develop), and long-range inhibition causes activity to be anti-correlated over larger distances (causing the eye-specific regions to alternate).




Fig. 5.17. Long-range lateral connections in the ocular dominance map. The inhibitory lateral connection strengths for the two neurons in Figure 5.15 are plotted on top in gray scale, and their local neighborhood is shown in the map below (as in Figure 5.12). In both cases the strongest connections come from the neuron's near neighbors. The connections of the partly monocular neuron (a) follow the ocular dominance map structure, with strongest connections from neurons with the same eye preference (dark). As a result, the connection histogram (middle) is biased toward the right eye (dark). In contrast, the connections of strongly binocular neurons (b) are not influenced by the OD map, and their connection histograms mirror the histogram of the OD map (Figure 5.16b). Similar patterns have been found experimentally in cats (Löwel 1994; Löwel and Singer 1992).

In the normal OD case, the lateral connections are not particularly patchy, and most neurons receive connections from neighbors of both eye preferences (Figure 5.17). Such connectivity arises because most neurons are only partly monocular: Their activation is correlated even with neighbors that prefer the opposite eye, and such connections remain active. Only the connections of the most strongly monocular neurons follow the ocular dominance stripes. Such neurons are activated predominantly by inputs from one eye, and their activity patterns are more strongly correlated with other similar neurons.

5.4.3 Strabismic Ocular Dominance Maps

As discussed in Section 5.1.2, the maps and lateral connections in strabismic animals differ significantly from those of normal animals. Similarly, a LISSOM OD map organized with uncorrelated inputs is very different from the normal case (Figure 5.18). By the end of strabismic development, nearly all neurons have become




Fig. 5.18. Ocular dominance and long-range lateral connections in the strabismic ocular dominance map. The strabismic simulation was otherwise identical to the normal case of Figures 5.15–5.17, except the inputs were presented at random positions in each eye. Compared with the normal case, the OD stripes are wider (c), and nearly all neurons are highly selective (e) and highly monocular (indicated by the wide separation of the peaks in the histogram d). As an example, the neuron shown with the small white square in (c) has the receptive fields shown in (b); the connections from the left eye are poorly organized and weak (indicated by medium gray). In (c), the white outline delineates the strongest lateral inhibitory connections to this neuron. Unlike in the normal case, these connections include only monocular neurons responding to the same eye (visible in the weight histogram in (a)), and strictly follow the ocular dominance stripes. The connections are strongest in the immediate vicinity of the neuron, but not much weaker even near the stripe boundaries (a). Overall, strabismus changes the map organization, RFs, and lateral connections much like it does experimentally in animals (Löwel 1994; Löwel and Singer 1992; Figure 5.2).

strongly monocular, and the boundaries between the stripes for each eye are sharply defined. The lateral connections also correlate strongly with the ocular dominance stripes, i.e. nearly all of the lateral connections come from neurons that prefer the same eye. Similar patterns have been found in cats using anatomical tracing techniques (Löwel 1994; Löwel and Singer 1992).

These results are due to the uncorrelated inputs in the two eyes. Corresponding receptive fields on the two eyes are rarely active at the same time, so all active neurons in a small patch of cortex are likely to be receiving input from the same eye. Through Hebbian learning, the afferent connections of all these neurons to that

eye are strengthened, as are the lateral connections between them. Over the course of self-organization, these weight changes result in sharply defined boundaries between stripes and between lateral connection patches, each selective for one eye or the other.

The OD stripes are also significantly wider in the strabismic LISSOM model than in the normal model, matching experimental findings in cats (Löwel 1994) and in previous computational models (Goodhill 1993). In LISSOM, nearby neurons learn similar preferences if they frequently respond to the same inputs. When the inputs are uncorrelated, two neurons respond together if they prefer the same eye and their RFs on that eye overlap. Since the RFs of nearby neurons overlap significantly, the OD stripes tend to be wide, alternating between large groups that prefer one eye or the other. In the normal case, however, the inputs in the two eyes are correlated. Even neurons that prefer different eyes often respond together because they represent the same retinotopic location. As correlation increases, retinotopy determines more of the response, and the OD stripes become less selective and narrower. In the limit, the inputs are fully correlated, and the OD stripes disappear altogether.

The strabismic patterns of ocular dominance and lateral connections that LISSOM develops closely match observations in biology. The LISSOM model shows how such connections and global organization can develop, based on a network that extracts structure from correlations in the visual input.

5.4.4 Effect of Input Disparity

Abstracting the possible sources of ocular dominance, the LISSOM OD model above was based on inputs that differ in brightness in the two eyes (like the models of Bauer, Brockmann, and Geisel 1997; Riesenhuber, Bauer, Brockmann, and Geisel 1998). Some previous models have shown that OD maps might instead result from retinal disparity, i.e. differences in feature position in each eye (Burger and Lang 1999; Sirosh and Miikkulainen 1994a). Disparity itself is a separate feature dimension from ocular dominance. Neurons can be fully binocular yet prefer different positions in each eye, which is important for stereo vision. Even so, disparity does reduce the activity correlations between the eyes, and could thus contribute to ocular dominance as well.

Disparity was modeled in LISSOM by randomly placing a spot at (x, y) in the left retina, and then placing a corresponding spot within a radius of $s_s R$ of (x, y) in the right retina, where R is the width of the retina. The scatter parameter $s_s \in [0..1]$ specifies the spatial correlations between spots in the two retinas. When $s_s = 0$, the inputs are in perfectly matched positions, and when $s_s = 1$ they are scattered independently (as in strabismus). By adjusting s_s , it is possible to simulate different degrees of disparity between the images in the two eyes (see Appendix A.6 for the rest of the simulation parameters).

In Figure 5.19, networks trained with brightness differences alone ($s_s = 0$), with mild disparity ($s_s = 0.2$), with moderate disparity ($s_s = 0.4$), and with perfect strabismus ($s_s = 1.0$) are compared. All these networks developed OD maps, but those with disparity are less realistic. With only mild disparity, the eye preferences




Fig. 5.19. Effect of disparity on ocular dominance maps. Each row presents a different simulation using the same network but a different degree of disparity between the inputs in the two eyes. From left to right, each row shows a retinal activation (left and right eyes), final RFs of a set of sample neurons, their lateral connections, the OD preference map, its histogram, and the OD selectivity map. For comparison, the results from brightness differences are reproduced in the top row (labeled “Dimming”), and the results from strabismic maps in the bottom row. The main result is that OD maps can be obtained from disparity differences (with no brightness differences), but the results do not match animal maps well. Small amounts of disparity (scatter $s_s = 0.2$) result in unrealistically clear boundaries between stripes even with relatively weak OD preferences (row “Mild”), as is evident in the histogram. Moderate disparity ($s_s = 0.4$; row “Moderate”) approaches the strabismic results, with strongly monocular RFs, sharp stripe boundaries, and connections only to neurons that prefer the same eye, unlike in normal animals. These results suggest that ocular dominance patterns can result from differences in either position or brightness, but brightness differences lead to maps that more closely match those found in animals.

are weak, yet there is already a clear boundary between OD stripes, instead of the smooth boundaries seen in animal OD maps. Lateral connections are also patchy and specific to one eye, unlike in animals. As the disparity is increased further, the maps form sharply defined monocular stripes, more closely matching strabismic animal OD maps than normal maps (compare the moderate case with the strabismic case).

The reason for such sharp stripes is that the neurons have sharply defined ON and OFF subregions, and are therefore highly sensitive to position. Even a small

difference in position between the eyes activates entirely different sets of neurons, and thus responses to each eye are effectively uncorrelated (or even anti-correlated). These results suggest that input disparities must be small relative to the RF size, or else strabismic-like OD maps will develop.

Interestingly, realistic maps can be obtained even with disparity-based inputs if the ON and OFF channels are bypassed (not shown; Sirosh and Miikkulainen 1994a). In such a case, any input that a V1 neuron receives is excitatory, nearby inputs always excite nearby V1 neurons, and smooth OD preferences are obtained. Such self-organization may be possible in early prenatal development, based on internally generated inputs like those postulated for orientation in Chapter 9. As discussed in more detail in Section 17.2.3, such computational constraints predict what kind of receptive fields and input pattern correlations are likely to be found in early development.

In this section, LISSOM OD maps were trained based on unoriented Gaussians only. Most other input types, like those tested in Figure 5.13, would result in an overlaid OR map as well. The effect of other input types such as natural images on OD will therefore be studied in Section 5.6 as part of combined OR/OD/DR maps.

5.5 Direction Selectivity Maps

In this section the LISSOM model is extended to moving inputs. Spatiotemporal receptive fields develop as a result, and they are organized into a map according to preferred direction of motion. These results show that activity-dependent self-organization extends to time-varying input as well, which is a crucial step toward making the model biologically realistic. The last step, organizing multiple stimulus features at once, will be discussed in the next section.

5.5.1 Method

Direction selectivity in LISSOM arises from LGN output that arrives at the cortex with a variety of delays (as in some previous models, e.g. Wimbauer, Wenisch, Miller, and van Hemmen 1997a; Wimbauer et al. 1997b). In the LGN, most cells fire soon after the retinal stimulus. However, other neurons, called lagged cells, respond only after a fixed delay (Saul and Humphrey 1992). The delay times vary between cells over a continuous range up to hundreds of milliseconds (Mastronarde, Humphrey, and Saul 1991; Wolfe and Palmer 1998). V1 neurons can use these timing differences to develop spatiotemporal receptive fields (Humphrey, Saul, and Feidler 1998).

The LISSOM architecture for direction selectivity is shown in Figure 5.20. The model is similar to those presented in previous sections, consisting of a hierarchy of two-dimensional sheets of neural units. Retinal receptors feed input to several paired sheets of ON-center and OFF-center LGN units (with a different lag for each pair), which in turn activate cortical neurons in V1 (see Appendix A.1 for the model equations). Because direction preferences are generally defined in terms of oriented




Fig. 5.20. LISSOM model of orientation and direction selectivity. The architecture is similar to the LISSOM OR network, except the ON and OFF channels consist of sheets of neurons with different lag times (from 0 to 3 in this case). Moving input patterns are drawn on the retina in discrete timesteps, like frames of a movie. At the first timestep, the ON and OFF LGN cells with time lag 3 compute their activity. At each subsequent timestep, the input pattern is moved slightly and LGN cells with lags 2, 1, and 0 each compute their activity in turn. Once all LGN cells have been activated, the initial V1 response is computed based on the responses on the eight LGN sheets. The activity then spreads laterally within V1 as usual in LISSOM.

lines, the model is self-organized to represent both direction and orientation at once, using moving oriented patterns as input.

The input consists of short movies of oriented Gaussians moving across the retina at random locations and directions. Each movie is presented as a sequence of frames. At each timestep t , the frame t is drawn on the retina, and the activity levels of all LGN cells with lag t are calculated. After all the LGN neurons have been activated, each V1 neuron computes its initial response based on the activation on all LGN sheets. After the initial response, the V1 activity settles through short-range excitatory and long-range inhibitory lateral interaction, and afferent and lateral weights are modified as described in previous sections. The model is then ready for the next input movie presentation.

In the experiments in this section, there were four 36×36 ON-center and four 36×36 OFF-center sheets, and they received input from a single 54×54 retina with a single moving oriented Gaussian pattern. Single inputs were used to avoid overlap and bias in the input distributions. The network was organized in 20,000 presentations of moving input patterns, so that the total number of inputs was the same as in OR and OD simulations. The rest of the parameters are described in Appendix A.7. The orientation and direction maps that resulted are analyzed next, followed by a comparison of maps obtained at different input speeds.




Fig. 5.21. Self-organization of afferent weights into spatiotemporal RFs. (a) The lag-0 weights for a sample neuron, located as shown in Figure 5.24a, are plotted before self-organization (as in Figure 5.5). Initially, all four lags in both channels have the same random weights; these weights are different for each neuron. (b) The final afferent weights for the same neuron are visualized by subtracting the OFF weights from the ON weights (as in Figure 5.7). Together, these plots show that the most effective stimulus for this neuron is a diagonal light bar moving diagonally down and to the right. More specifically, this neuron will be highly active at time t if there was a light bar aligned with the ON subregion in the “Lag 3” RF at time $t - 3$, a bright bar aligned with the ON subregion of the “Lag 2” RF at time $t - 2$, and so on. Visual cortex neurons in animals have similar spatiotemporal properties (Figure 5.4a; DeAngelis et al. 1995).

5.5.2 Direction Maps

Figure 5.21 shows the self-organized afferent weights for a representative neuron in V1. Nearly all neurons developed spatiotemporal receptive fields strongly selective for both direction and orientation. Each neuron responds best to a line with a particular orientation moving in a direction perpendicular to that orientation. Such receptive fields are similar to those found experimentally in the cortex (Section 5.1.3; DeAngelis et al. 1993, 1995).

The orientation preferences form a smoothly varying orientation map (Figure 5.22). As in the LISSOM OR-only model, the map contains realistic features such as iso-orientation patches, linear zones, pairs of pinwheels, saddle points, and fractures. The techniques used to analyze the orientation map in Section 5.3 (selectivity measures, Fourier transform, gradient, retinotopic mapping, and OR preference histogram), lead to essentially the same results. This outcome verifies that OR-only simulations are a valid approach to understanding orientation maps.

The same neurons are also selective for motion direction, forming a direction map (calculated as described in Appendix G.5). Direction preferences are also smoothly organized across the cortex, and contain similar features. Regions of lower direction selectivity occur near pinwheel centers and along fractures, as in the OR map. The orientation and direction preference histograms are essentially the same (i.e. flat), both reflecting the distribution of edges in the training input.

The interaction between the orientation and direction maps is illustrated in Figure 5.23, and is similar to what has been observed in animals (Section 5.1.3; Shmuel and Grinvald 1996; Weliky et al. 1996). For instance, a patch of neurons highly selective for one orientation and direction of motion will usually have an adjacent or contiguous patch selective for the same orientation but opposite direction. In LIS-




Fig. 5.22. Self-organized OR/DR map. The orientation (top row) and direction (bottom row) maps in the LISSOM OR/DR model were computed separately after self-organization. The orientation preferences are coded using the color bar key on top, and the direction preferences using the color arrow key in the middle. Selectivity is shown in gray scale in both cases, with black indicating low selectivity (as in Figure 5.9). (a) The network represents both orientation and direction in smoothly varying maps that contain all the features found in animal maps, such as linear zones, pairs of pinwheels, saddle points, and fractures (outlined as in Figure 2.4). (b) Most neurons become selective for specific orientation and direction of motion, and are therefore nearly white in the selectivity plots. (c) Overlaying the preference and selectivity plots shows that regions of low selectivity occur near pinwheel centers and along fractures in both maps. (d) The histograms are essentially flat because the training inputs were unbiased. These plots show that LISSOM can develop biologically realistic orientation and direction maps through self-organization based on abstract input patterns.

SOM, these patterns develop because neurons that prefer similar orientations but opposite directions have more similar RFs (and thus responses) than neurons that prefer different orientations. As a result, opposite direction preferences are often grouped together on the map.

Lateral connections within the map follow its global organization, primarily linking neurons with similar orientation and direction preferences (Figure 5.24). Such connection patterns reflect the activity correlations during self-organization. Neurons with similar orientation and direction preferences are often active together, and become more strongly connected. Connections of neurons that are highly selective for orientation and direction extend along the orientation preference, not direction, and avoid orthogonal orientations and opposite directions. Neurons in DR fractures




Fig. 5.23. Combined OR/DR map. Using the plotting conventions from Figure 5.4, each of the 142×142 neurons in the LISSOM OR/DR map is colored with its preferred orientation, and the direction preferences for every third neuron are plotted as arrows overlaid on the orientation map. The direction preferences are generally perpendicular to the preferred orientation, and large iso-orientation patches are often divided into two areas with opposite direction preferences. Such an organization matches experimental data well (Figure 5.4; Weliky et al. 1996). An animated version of this plot can be seen at <http://computationalmaps.org>.

connect with neurons of both directions; they all have similar orientation preferences, and the connections extend along that orientation. At DR pinwheels the connections come from all direction preferences, and at DR saddles they correspond to the direction preferences in the saddle. Both DR pinwheels and DR saddles can occur at a variety of OR map features.




Fig. 5.24. Long-range lateral connections in the combined OR/DR map. The inhibitory lateral connections of four sample neurons from the marked regions in the DR map of Figure 5.22 are shown both on the OR (top) and DR (bottom) preference and selectivity maps, as in Figure 5.12. (a) The neuron from Figure 5.21 receives connections from neurons with similar OR and DR preferences, and its connections extend along the OR axis (60°), not the DR axis (330°). (b) Neurons in pinwheels receive connections from neurons of all DR preferences; their OR histograms vary depending on their location in the OR map. (c) Saddle neurons receive connections corresponding to the DR preferences of the saddle, and their OR histograms vary. (d) At fractures, connections originate primarily from neurons of the two fracture preferences, and extend along their common OR preference. These OR- and DR-specific lateral connections are predictions that can be tested in future biological experiments.

While lateral connections are known generally to follow orientation preferences, how such connectivity is further affected by direction selectivity has not been analyzed. The LISSOM results constitute predictions for future such biological experiments.

5.5.3 Effect of Input Speed

Compared with the orientation-only simulations, the OR/DR map introduces one new parameter: the speed of the input patterns, i.e. how many retinal units the pattern moves between lags. Figure 5.25 shows how this parameter affects the formation of direction maps.

For speed zero, the inputs are stationary and no direction map develops. The RFs to each lag become identical, and are essentially the same as in the OR-only map (Section 5.3). As the speed is increased, the neurons become more selective for direction of motion. Interestingly, the spacing between DR patches becomes larger as the neurons become more selective; this trend is also visible in the ring diameter in the Fourier plots (Figure 5.25). At the same time, the ring diameter in the OR map increases slightly, and by speed 2 the features in the OR map become smaller than those in the DR map.

Thus, with fast enough inputs, the DR map becomes the largest-scale organization, with smaller patches for orientation. Although the results from speed 1 are most similar to the existing animal results (from the ferret), the results from higher speeds are predictions for maps in other species with greater motion sensitivities. LISSOM also predicts that animals raised in environments with more motion during the animals' critical period will develop direction maps whose spatial scale is larger than that of their orientation maps.

Overall, the results in this section show that LISSOM can account for spatiotemporal preferences in addition to spatial ones. Although Gaussian images were used in this section for clarity, similar results can be obtained with natural images (as will be demonstrated in the next section).

So far in this chapter we have seen how orientation selectivity, direction selectivity, and lateral connectivity each develop synergetically in the model, and the results match data from animal experiments. The next step is to show how they interact in a single unified model of feature preferences in the cortex.

5.6 Combined Maps of Multiple Features

The LISSOM OR, OD, and DR models introduced in the previous sections are all based on the same basic architecture, differing only in the number of eyes and LGN sheets simulated and the type patterns used to train the model. Such uniformity makes it possible to combine them into a comprehensive model that self-organizes all three feature preferences at once. The model is first validated against biological data on how orientation and ocular dominance interact. It is then trained with moving images, and predictions are made on how all three features self-organize together. In the final




Fig. 5.25. Effect of input speed on direction maps. From left to right, each row shows a sample retinal activation at lag 0, final receptive fields to LGN regions with lags 3, 2, 1, and 0 (left to right) of six sample neurons, the inhibitory lateral connections of those six neurons, the orientation preference and selectivity map, the Fourier transform of the OR preferences, the direction preference and selectivity map, and the Fourier transform of the DR preferences. Orientation and direction histograms are not shown because they are all nearly flat. Each row shows the result from using training inputs moving at a different speed, ranging from zero (stationary) to moving three retinal units between each group of lagged LGN cells. For the example input shown, the lag 3 input was always the one shown in the top row (labeled “Speed 0”), and by lag 0 it had moved to the position shown in each row. When the inputs were stationary (i.e. all lags had the same input patterns), no direction map or direction-selective units developed, and the “DR pref. & sel.” map is entirely dark. As the speed increases, more units become direction selective, and direction becomes the largest-scale organization in the map. This increase in feature size is visible in the DR Fourier transform plots, where a smaller spatial frequency (larger feature spacing) leads to smaller rings as speed is increased. These results are predictions for maps in animals with different retinal motion sensitivities or those raised in environments with different speeds of visual motion.




Fig. 5.26. LISSOM model of orientation, ocular dominance, and direction selectivity. The architecture is a combination of the LISSOM OD network of Figure 5.14 and the DR network of Figure 5.20. For each input movie, one retina is chosen randomly to be dimmer than the other. The frames are then drawn on the retina one timestep at a time, and propagated to the LGN sheets with the appropriate lag. The V1 neurons receive input from all 16 LGN sheets simultaneously, and settle the activation through lateral connections as usual in LISSOM.

section, the effects of organizing it with natural images instead of oriented Gaussians are considered. This stepwise approach makes it clear which aspects of the input patterns and architecture are crucial for each type of feature preference.

5.6.1 Method

The LISSOM model of simultaneous OR/OD/DR is a combination of the ocular dominance network of Section 5.4 and the direction selectivity model of Section 5.5. It contains 16 LGN sheets, consisting of four ON and four OFF sheets for each of the two eyes (Figure 5.26; see Appendix A for the model equations).

The model is trained with moving oriented Gaussian patterns at randomly chosen locations and directions. As in the DR network, one Gaussian is used per iteration, but multiple Gaussians would be used for larger retinas. As in the OD network, the brightness s_b of the Gaussian in the left eye is chosen randomly in the range $[0..1]$ for each movie, and the brightness of the corresponding spot in the right eye is computed as $1.0 - s_b$. Also as in the DR network, each of the four sheets in the four LGN channels receives a frame of the input movie with a different lag. Once all 16 LGN sheets have been activated, V1 neurons calculate their initial activation based on the total LGN activation, settle it through lateral connections, and adapt their weights as in previous models.

While there are significant biological data on how OR and OD interact in biological maps, little is known at present on how DR confounds these interactions. Moreover, although lagged LGN cells have been found in cats (Hubener, Shoham, Grinvald, and Bonhoeffer 1997; Löwel, Bischof, Leutenecker, and Singer 1988), they may not exist in monkey LGN (Saul and Humphrey 1992). Therefore, in the first experiment in this section, only the LGN sheets with lag 0 will be used, and the resulting OR/OD map will be compared with OR/OD results in monkeys. The parameters of this experiment are set as in the OD-only simulation (Appendix A.6). In the second experiment, the full OR/OD/DR model is trained for 20,000 iterations as described above, and predictions are made on how these features interact in animals like cats. The parameters for the OR/OD/DR simulation are described in Appendix A.9.

5.6.2 Combined Orientation / Ocular Dominance Maps

Through self-organization, the combined network without lagged inputs developed realistic maps for both orientation and ocular dominance (Figure 5.27). Both maps are similar to the single-feature maps of Sections 5.3 and 5.4, and contain the features typical of animal maps. The OR map contains linear zones, pairs of pinwheels, saddle points, and fractures, and its selectivity, Fourier transform, gradient, retinotopic mapping, and OR preference histogram are similar to those in previous LISSOM OR maps. The OD preference histogram and the OD selectivity map are similar to the previous LISSOM OD map. These results demonstrate that LISSOM is capable of forming realistic maps of independently varying feature dimensions simultaneously.

The maps interact in patterns similar to those seen in monkeys (compare Figure 5.27 with 5.3). Pinwheel centers and fractures rarely overlap with the OD stripe boundaries; ocular dominance boundaries tend to cross orientation patches in linear zones at right angles, and OD boundaries rarely follow an orientation boundary. Neurons that are least selective for orientation are thus most selective for the eye of origin, and vice versa. Such patterns emerge in LISSOM because the self-organizing process favors an organization where the responses are spread evenly across the cortex (Section 3.4.3). As a result, the different regions of the map become selective for different features. Overall, these patterns are consistent with biological data, suggesting that LISSOM is a valid model of how such features develop in animals like monkeys.

Interestingly, when disparity is used to create differences between the eyes (instead of dimming), the joint OR/OD map that develops is not as realistic (not shown). In this case, each eye develops independent orientation maps, and the boundaries between orientation patches often follow the OD boundaries (unlike in animals). In effect, the neurons responding to each eye develop independently, because a given neuron rarely receives activation from both eyes at once. This result provides further computational evidence that processes like dimming, not disparity, are the source of OD maps.

The receptive fields for each eye are similar to those in the LISSOM OR network. They are all binocular to some degree, with slightly stronger connections from one eye or the other. The lateral connections are strong locally, and in the longer range




Fig. 5.27. Self-organized OR/OD map. Based on oriented Gaussian patterns with different brightnesses in each eye, LISSOM develops realistic orientation (a) and ocular dominance (b) maps in the same area of cortex. The orientation map features are outlined in (a), (c) and (d) as in Figure 2.4. In (c), the orientation preferences are overlaid with the ocular dominance gradient: High gradient (black) marks the boundary between OD stripes. These boundaries rarely overlap pinwheel centers or fractures, they intersect OR boundaries in linear zones at right angles, and they rarely follow OR boundaries. These relationships are further highlighted in (d), where regions of low orientation selectivity (pinwheel centers and fractures) are plotted in dark gray, overlaid with the ocular dominance boundaries. Again, these features rarely intersect, suggesting that the map organization results from distributing selectivity for different features evenly across the cortex. Similar interaction between orientation and ocular dominance is seen in biological maps (Figure 5.3; Blasdel 1992b)

follow the orientation preferences, as they do in animals (Figure 5.28). The neurons most selective for one eye also prefer connections from that eye, but the preference is not absolute like it is in strabismic animals (Section 5.4.3). These results show that the lateral connection patterns seen in the simpler separate maps also exist in the combined maps, with preferences for multiple features overlaid onto the same set of neurons.

5.6.3 Combined Orientation / Ocular Dominance / Direction Maps

To model all three feature types in the same network, the entire model of Figure 5.26 was trained with movies of moving oriented Gaussians, randomly dimmed in one eye. The network developed realistic maps for orientation, ocular dominance, and direction selectivity (Figure 5.29). The combined map is more complex than those developed by the single-feature networks, because the network is representing three overlaid, interacting feature maps.

However, the typical structures of animal maps are still found in this more complex map. The OR, OD, and DR maps have the same features and the same quantitative measures as the individual LISSOM OR, OD, and DR maps. Orientation patches are often divided into patches selective for opposite directions of motion, and orientation pinwheels and fractures avoid boundaries of the ocular dominance stripes. The receptive fields are analogous to those in single-feature LISSOM networks, i.e. oriented and binocular, with a slight preference for one eye or the other (Figure 5.32). The lateral connections are strong in a local neighborhood, but at longer distances mostly follow the orientation and direction preferences, as was found in the individual LISSOM OR, OD, and DR networks (Figure 5.32). These results demonstrate that LISSOM can form realistic maps and lateral connections based on multiple features; on the other hand, they also show that the approach of studying each feature separately is valid and leads to insights that carry over to more complex maps.

Interestingly, some relationships that were clear in the OR/OD network are not as uniform in the OR/OD/DR network. For instance, the ocular dominance boundaries in linear zones do not intersect orientation patches at right angles as often, and pinwheel centers are not always at the center of OD stripes. Such variation corresponds to differences between maps in different animal species. The patterns observed in the OR/OD network have been characterized primarily in monkeys, and it is not known if monkeys have lagged cells at the LGN level (Saul and Humphrey 1992). Conversely, the lagged cells in the OR/OD/DR network are similar to those found in cats, where the OR/OD intersection patterns are less clear (Hubener et al. 1997; Löwel et al. 1988), similar to the OR/OD/DR LISSOM model. The two versions of the combined LISSOM model therefore suggest that the differences seen in the intersection patterns may be due to the differences in how these species represent time-varying input.

5.6.4 Effect of Input Types

The oriented Gaussians used in the above experiments allow obtaining maps where the feature interactions are most clear. The last step is to extend these results to




Fig. 5.28. Long-range lateral connections in the combined OR/OD map. The inhibitory lateral connections for four sample neurons from the regions marked in Figure 5.27 are shown situated on the OR preference and selectivity map (top) and the OD map (bottom) as in Figures 5.12 and 5.17. The connection patterns on the OR map are similar to those in the orientation-only map. The most selective iso-OR neurons (*a*) are located near the OD stripe boundaries, and receive connections equally from neurons of both eye preferences. Neurons in iso-OR regions away from the OD boundaries connect more strongly to one eye (not shown). OR fractures (*b*) and pinwheels (*d*) tend to occur near the centers of OD stripes and receive connections primarily from the same eye preference. Saddle points (*c*) can occur either in the middle or near the boundaries of OD stripes, and thus can have either monocular or binocular connection patterns; the example neuron is in a monocular OD region and connects primarily to the left eye. These connection patterns further extend the predictions from Figure 5.12, showing how the connection patterns are shared between the OD and OR maps.




Fig. 5.29. Combined OR/OD/DR map trained with Gaussians. Based on oriented, moving Gaussian patterns with different brightnesses in each eye, LISSOM develops overlaid orientation, ocular dominance, and direction maps simultaneously. This plot shows the orientation preferences in color coding, the boundaries of the OD stripes in black, and the direction preferences and selectivities as white arrows, as in Figures 5.23 and 5.27. The network develops a realistic orientation and direction map, with OR patches subdivided into areas preferring the opposite directions of motion. Ocular dominance boundaries tend to cross linear zones at right angles, rather than following the orientation map. These results are similar to the ones with individual input dimensions, complicated by the fact that multiple dimensions are being mapped at once. Similar results have been observed experimentally with the cat visual cortex (Hubener et al. 1997; Löwel et al. 1988).




Fig. 5.30. Example natural image input for training the OR/OD/DR map. The top row shows a sequence of images presented to the right eye and the bottom row the corresponding images presented to the left eye. The two sequences are identical except for brightness; in this example, the right eye was randomly chosen to be darker than the left. The sequence represents a short movie where the image is moving to the left and slightly downward. The original image is from a dataset by Shouval et al. (1996, 1997).

other types of inputs, including noisy disks and natural images. These input types will be used in Part III as internal and external training patterns, in order to obtain maps most comparable to biology. Identifying how the complexity of natural internal and external input changes the observed map structures is an important step toward understanding how biological maps develop.

The simulations were run otherwise with the same parameters as the Gaussian OR/OD/DR experiment except for the training inputs, which consisted of the noisy disk and natural image patterns of Section 5.3.5 (see Appendix A.9 for details). A single noisy disk per iteration was used in the first experiment, located randomly on the retina and moving in a random direction. In the second, an input pattern was selected randomly from the dataset of images (by Shouval et al. 1996, 1997) and swept across each eye in a random direction as shown in Figure 5.30. In both simulations, the left and right eyes had identical inputs, differing only in brightness.

The maps developed with natural image inputs (Figure 5.31) were similar to the Gaussian input case (Figure 5.29), although more variable and less smoothly organized. Again, OD boundaries tend to cross OR linear zones at right angles, and a patch of neurons highly selective for a direction often has a nearby patch selective for the same orientation but the opposite direction. Neurons highly selective in one dimension receive connections primarily from neurons with similar preferences for that dimension, while neurons with low selectivity (e.g. binocular neurons) receive connections from neurons with a wide range of preferences.




Fig. 5.31. Combined OR/OD/DR map trained with natural images. The combined LIS-SOM map from natural images is similar to the one in Figure 5.29. More of the units are selective for orientations near horizontal and vertical, so the map has more red and cyan than with artificial training inputs. Not as many neurons are highly selective for direction, but many of the selective patches are located next to other patches selective for the same orientation but opposite direction. Overall, these results are similar to those from artificial stimuli, with greater variability reflecting the more complicated feature correlations in natural images.

Natural image input affects the RFs and lateral connections more strongly than the maps (Figure 5.32). As in the single-feature simulations, the afferent RFs develop a variety of shapes, including both two-lobe and three-lobe RFs, in contrast to the uniformly three-lobed RFs of the Gaussian-trained map. Because natural images have correlations with longer range, the lateral connections are also wider and patchier, as they are in the single-feature natural image maps.




Fig. 5.32. Effect of training patterns on OR/OD/DR maps. From left to right, each row shows a sample retinal activation, final spatiotemporal receptive fields of sample neurons with lags 3, 2, 1, and 0 from left to right, their inhibitory lateral connections, the orientation preference and selectivity map, the ocular dominance map, and the direction preference and selectivity map. The top row shows a network trained with oriented Gaussians, the middle with noisy disks, and the bottom with natural images. All networks develop realistic orientation maps. Maps with Gaussian inputs and noisy disks develop only smooth, two- or three-lobed RFs, whereas networks with natural images develop a wide variety of RF types, corresponding to the wide range of patterns seen in natural images. In each case, the lateral connection patterns follow the features in the map. These results show that joint OR/OD/DR maps can develop from a variety of abstract and realistic input stimuli.

The simulations with noisy disk patterns show that joint maps of ocular dominance, orientation, and direction can also develop based on spontaneous activity patterns (Figure 5.32). Again, the RFs and lateral connections differed significantly from the Gaussians case, with primarily two-lobed RFs and more long-range lateral connections. The map was less selective for direction than with the other input types, primarily because the neurons became selective for moving curved edges whereas selectivity was measured using straight sine gratings. The ocular dominance map is more sharply delineated (as it is in strabismus) than with other inputs, primarily because the background noise was often stronger than the dim input in the other eye, reducing the overall correlation. Overall, these results suggest that if spontaneous activity patterns have enough motion and differ enough between the eyes, maps of all three preference types can develop.

Together, the results presented in this section show computationally for the first time how the OR, OD, and DR input features interact during development, and sug-

gest a simple explanation for how the combined structures observed in animal maps can emerge. When multiple maps are overlaid, the global structure becomes more difficult to interpret, which is why the single-feature maps constitute a useful abstraction. The combined results are realistic and allow matching of anatomical differences with map organization in different species (as will be discussed in more detail in Section 17.2.2). The model also suggests which complexities are due to mapping multiple features and which are a result of complexities in natural inputs.

5.7 Discussion

The LISSOM model shows how a single local and unsupervised self-organizing process can be responsible for the development of topographic maps and the lateral connection structures in the primary visual cortex. The model is the first to account for the self-organization of lateral connections together with orientation maps, ocular dominance maps, and direction selectivity maps. Starting from random-strength connections, neurons develop oriented spatiotemporal receptive fields and lateral interaction profiles cooperatively and simultaneously. When the input varies in several dimensions at once, the cells and lateral connections become selective for multiple features, using the same architecture and learning rules that apply to individual feature dimensions. Such self-organization stores long-range activity correlations between feature-selective cells in the lateral connections. As will be demonstrated in more detail in Section 14.2, these correlations can be used to eliminate redundant information during visual processing, and to make cortical cells more selective.

LISSOM makes several testable predictions about specific lateral connection patterns in the cortex. The LISSOM orientation map predicts that the long-range connections at pinwheel centers, saddle points, and fractures have unselective, broad unimodal, and biaxial distributions, respectively. The direction map predicts that long-range lateral connections primarily link neurons with similar orientation and direction preference, extend along the orientation preference, and avoid orthogonal orientations and opposite directions. At DR fractures they connect with both directions and extend along their common orientation, in DR saddles they connect with the directions of the saddle, and in DR pinwheels they connect with all directions. The OD preferences are overlaid with the OR and DR preferences: The neurons most selective for one eye prefer connections from that eye, but such a preference is absolute only in strabismic animals. These predictions can be tested experimentally by combining optical imaging and injected tracers (like Bosking et al. 1997; Löwel 1994; Löwel and Singer 1992; Malach et al. 1993; Sincich and Blasdel 2001), as will be discussed in Section 16.4.2.

For clarity, the simulations shown in this chapter were each based on a single type of input pattern, such as Gaussians or natural images. The biological visual cortex, however, may be exposed to multiple sources of activity during development, including spontaneous, internally generated patterns and visual, externally evoked inputs. Biological development is thus likely to depend on a complex combination of such patterns. As will be shown in detail in Chapter 9, modeling both prenatal

internal activity and postnatal visual images allows the model to account for both the primitive orientation maps seen at birth in animals and the more refined adult maps. Such more complex simulations are a natural extension of the single-input experiments presented in this chapter, providing a solid basis for understanding how V1 develops in animals.

LISSOM simulations serve to identify the types of input patterns that lead to normal and abnormal development of maps and connections, which can in turn help focus future biological experiments. For instance, realistic LISSOM ocular dominance maps develop from strength or brightness differences between patterns in the two eyes, but not from large position differences. This result leads to interesting predictions about the roles of internal and external input in shaping the OD maps. In newborn mammals, a rudimentary OD map exists at birth, presumably organized based on prenatal internally generated input such as retinal waves (Crair, Horton, Antonini, and Stryker 2001; Horton and Hocking 1996). If retinal waves are the only source, the LISSOM model predicts that the map would be strabismic, because the retinal waves are uncorrelated between the two eyes. This prediction is difficult to verify directly, but because normal adult maps are not strabismic, two interesting possibilities follow: (1) If the newborn maps are not strabismic, the LISSOM results suggest that the internally generated input that constructs the newborn map must be correlated between the two eyes. Since the retinal waves are not, there must be some additional source of internally generated patterns, like the PGO waves, that plays this role. (2) If newborn maps are indeed strabismic, visual experience must play a crucial role in shaping a normal adult map. In effect, the OD map is then constructed in two phases, prenatal and postnatal, like the orientation map. Possibilities for identifying such patterns and processes experimentally, and for modeling them computationally, are described in more detail in Section 17.2.3.

The LISSOM direction selectivity model is based on a simple but effective model of moving inputs, i.e. translation in a random direction within the image plane at each presentation. In the future it can be made more realistic by interleaving input frames with propagation and settling. Such a process is more complicated but should lead to similar results. More importantly, it is not yet clear what types of moving input drive the development of direction-selective neurons and direction maps in animals. Possible candidates include drifting retinal waves before birth, moving objects in the environment (relative to a stationary background), and optic flow due to eye or head movements. Each of these types of natural motion have different statistical properties, which is important to take into account in a more detailed model. As suggested above for OD and OR, multiple types of internal and external time-varying input are likely to contribute to the developmental process. Future models will depend on measuring and characterizing sources of these inputs in detail, as will be discussed in Section 17.2.3.

In addition to the lagged LGN cells modeled in LISSOM, different species may utilize other mechanisms for establishing direction selectivity and motion preferences (see Clifford and Ibbotson 2002 for a review). For instance, some connection pathways within V1 (e.g. between cortical layers in a column) may have delays long enough to serve as a memory of previous activity levels, much like the lagged af-

ferent inputs do. Once they are characterized in sufficient detail, such mechanisms can be implemented in LISSOM, and they can potentially help account for a wider variety of motion preference results within the same basic framework (as will be discussed in Section 17.1.5).

In a similar manner, the LISSOM model can be instrumental in understanding differences between species more generally. The simulations in this chapter were necessarily based on pooled data from several different species because comparable experimental data do not yet exist across species. As such data are obtained, the model can be parameterized to account for a specific species in detail. As will be discussed in Section 17.2.2, such computational experiments should allow us to understand which differences are significant and what their origins are.

The LISSOM results in this chapter focused on orientation, ocular dominance, and direction maps, which are experimentally the best understood feature dimensions. Future models may also be trained with input that vary in spatial frequency, color, and disparity. Such simulations should result in preferences and maps for these additional features without requiring significant changes to the LISSOM model itself. Because little is known about these dimensions in biology, the simulations could be used as a guideline for further experimental studies. Such simulations will be discussed in more detail in Section 17.2.1.

5.8 Conclusion

The results in this chapter demonstrate that a single local and unsupervised self-organizing process can develop both the afferent and lateral connection structures in the primary visual cortex. This model is the most complete computational simulation of V1 to date. The model suggests that the afferent connections represent visual features with the highest variance, such as topography, orientation, ocularity, and direction. The lateral connections represent correlations between such features, and implement an efficient coding of visual information.

The same self-organizing process may continue to operate in the adult, serving a different role: It may be responsible for plasticity and adaptation after damage, as will be described in the next chapter.