

# Constructing visual function through prenatal and postnatal learning

**James A. Bednar and Risto Miikkulainen**

{jbednar,risto}@cs.utexas.edu

Department of Computer Sciences

The University of Texas at Austin

1 University Station C0500 Austin, TX 78712 USA

February 16, 2004

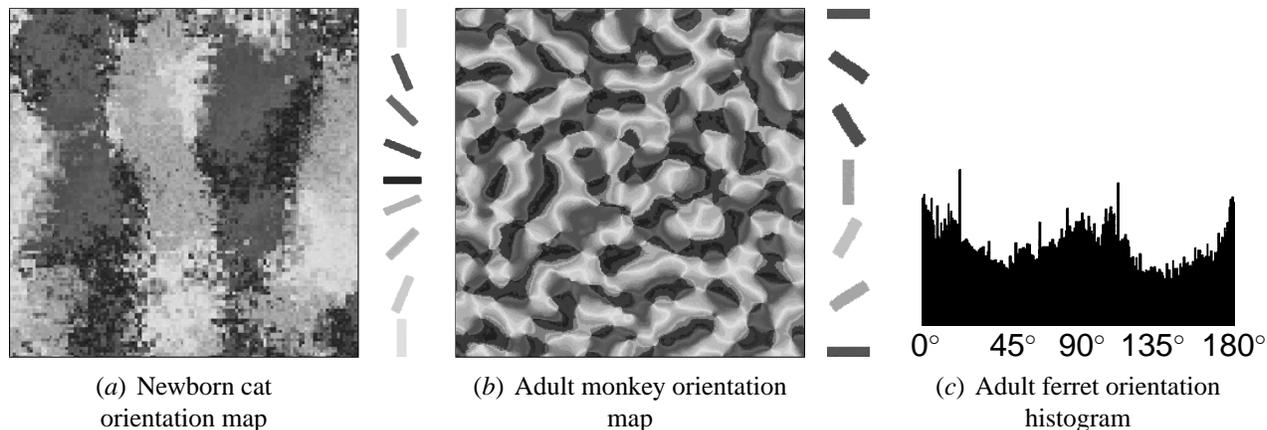
## Abstract

Although environment-driven learning can explain much of postnatal neural development, substantial organization and functional ability is present even at birth. Recent experimental discoveries of widespread spontaneous neural activity suggest that prenatal development may utilize very similar mechanisms and principles as postnatal learning, driven by internally generated sources instead of the environment. This chapter shows how this idea can explain features of the organization and function of the primary visual cortex (V1) and higher level face-processing areas. Specifically, we simulate how neural preferences for contour orientation and human faces can develop prenatally from internally generated activity and postnatally from natural image stimuli. These simulations are based on HLISSOM, a hierarchical self-organizing model of the development of topographic neural maps. The results match experimental neuroimaging and psychophysical data from newborn and older animals and humans, and provide concrete predictions about infant behavior and neural activity for future experiments. They also suggest that combining internally generated activity with a learning algorithm is an efficient way to develop complex neural machinery.

## 1 Introduction

Specific regions in the visual cortex of adult humans and other adult mammals selectively respond to different visual features. For instance, the primary visual cortex (V1) is arranged into an orientation map, consisting of an array of neurons that respond to visual contours with certain orientations, but not to others (Figure 1; Blasdel 1992; Hubel and Wiesel 1962, 1968). Similarly, visual areas in the fusiform gyrus respond more strongly to human faces than to other complex stimuli (Kanwisher, McDermott, and Chun 1997).

How these selective responses develop is not yet clear. On the one hand, experimental and computational results (reviewed later) suggest that much of the structure and function of the visual



**Figure 1: Newborn and adult animal orientation maps.** Plot (a) shows a  $1.9 \times 1.9$  mm area of newborn cat V1, and (b) shows a  $5 \times 5$  mm area of adult monkey V1, both measured by optical imaging techniques. (a reprinted with permission from Crair et al. (1998), *Science*, 279:566–70, copyright 1998 American Association for the Advancement of Science; b reprinted with permission from Blasdel (1992), *J. Neurosci.*, 12:3139–61, copyright 1992 by the Society for Neuroscience). Each neuron in the maps is colored in grayscale according to the orientation it prefers, using the key adjacent to it. Nearby neurons in the map generally prefer similar orientations, forming groups of the same color. Other qualitative features are also found, such as pinwheels, i.e. points around which orientation preference changes continuously. Newborn maps have similar qualitative features, although they are noisier. (c) The distribution of orientation preferences has been measured in ferrets, and has a bias towards vertical ( $90^\circ$ ) and horizontal ( $0^\circ$  and  $180^\circ$ ) contours, which reflects the statistics of the environment. (c reprinted with permission from Coppola et al. 1998, *PNAS* 95:2621–3, copyright 1998 National Academy of Sciences, U.S.A.).

system is constructed by a general-purpose learning process, driven by inputs from the environment. On the other hand, there is also considerable evidence that many aspects of the visual system are hardwired, i.e., constructed from a specific blueprint encoded in the genome. The conflict between these two positions is generally known as the Nature–Nurture debate, which has been raging for centuries in various forms (Diamond 1974).

The idea of a specific blueprint does seem to apply to the largest scale organization of the visual system, at the level of visual areas and their interconnections. These patterns are largely similar across individuals of the same species, and their development does not generally depend on neural activity, visually evoked or otherwise (Miyashita-Lin, Hevner, Wassarman, Martinez, and Rubenstein 1999; Rakic 1988; Shatz 1996). But at smaller scales, i.e., neurons and connections within visual areas, there is considerable evidence for both environmental and internally controlled development. Thus, debates center on how this seemingly conflicting evidence can be reconciled.

In this chapter, we will briefly give an overview of the relevant areas of the visual system, and summarize the evidence for environmental and genetic influences on orientation and face processing. We will then propose a simple hypothesis about how the cortex can combine both influences, through prenatal learning of spontaneous activity patterns followed by postnatal learning of visual stimuli. Using a computational model of orientation and face processing, we will show that this hypothesis can explain a large body of data on how low-level (i.e. orientation) and higher level (i.e. face) perception develop. The orientation map simulations are strongly constrained by existing ex-

perimental data, providing a rigorous way to test and calibrate the model, while the face-processing simulations will provide a set of concrete predictions for future experiments.

## **2 Biological and psychological background**

The human visual system, like that of most mammals, consists of a series of topographically organized visual areas (for review, see e.g. Kandel, Schwartz, and Jessell 1991). When light falls on the retina, photoreceptor cells convert the light patterns into patterns of electrical activity. Neurons in the retina called retinal ganglion cells have receptive fields (RFs) covering a local patch of photoreceptors, and are active when the light patterns are a good match to their RFs. The retinal ganglion RFs tend to have a center-surround organization, responding to light dots surrounded by dark (ON retinal ganglion cells), or dark patches surrounded by light (OFF retinal ganglion cells). Retinal ganglion cell activity is relayed through the optic nerve to ON and OFF neurons in the LGN (lateral geniculate nucleus of the thalamus), which have RFs very similar to retinal ganglion cells. The LGN cells relay the activity patterns to the primary visual cortex (V1), which in the adult has RFs selective for oriented patterns, instead of the circular RFs typical of earlier regions. As mentioned earlier, these RFs form a topographically organized orientation map, representing the orientation of each small patch of the input image. The activity levels of neurons in the orientation map are relayed to neurons elsewhere in the visual cortex, which have RFs selective for more complex objects such as human faces.

The development of the eye, LGN, and V1 have been studied extensively, in part because these areas are very similar across many species. However, relatively little is known about how higher cortical areas develop. As described in the next subsection, experiments with V1 suggest that orientation maps develop from both environmental (i.e., external) and internal cues. Face processing is much more difficult to study, requiring higher primate or human subjects. Thus most of the available face-processing data is from behavioral studies of human infants, rather than detailed measurements of individual neurons and neural regions. These data suggest that face processing may develop similarly to orientation processing, driven by both environmental and internal sources (as described in the second subsection). A third subsection describes how internally generated activity could contribute to the development of these capabilities, based on recent experimental discoveries of widespread spontaneous activity.

### **2.1 Orientation map development**

Experiments in monkeys, ferrets, and cats suggest that orientation-selective responses in adults could be learned from exposure to oriented patterns (as reviewed by Movshon and van Sluyters 1981). For instance, Blakemore and Cooper (1970) found that if kittens are raised in environments consisting of only vertical contours, most of their V1 neurons become responsive to vertical orientations. Similarly, orientation maps from kittens with such rearing devote a larger area to the orientation that was overrepresented during development (Sengpiel, Stawinski, and Bonhoeffer 1999). Conversely, kittens raised without patterned visual experience at all, e.g. by suturing their eyelids shut, have few orientation-selective neurons in V1 as an adult (Blakemore and van Sluyters

1975; Crair et al. 1998). Even in normal adult animals, the distribution of orientation preferences is slightly biased towards horizontal and vertical contours (Chapman and Bonhoeffer 1998; Coppola et al. 1998). Such a bias would be expected if the neurons learned orientation selectivity from typical environments, which have more horizontal and vertical contours than oblique ones (e.g. trees and horizons; Switkes, Mayer, and Sloan 1978).

In ferrets, it is even possible to reroute the connections from the eye that normally go to V1, so that instead they reach the auditory cortex (as reviewed in Sur, Angelucci, and Sharma 1999; Sur and Leamey 2001). As a result, the auditory cortex develops orientation-selective neurons and orientation maps, although these structures are not as pronounced as in normal maps. Furthermore, the ferret can use the rewired neurons to make visual distinctions, such as discriminating between two grating stimuli (von Melchner, Pallas, and Sur 2000). These experiments show that the input to a cortical area can profoundly affect its structure and function.

However, there is also significant evidence suggesting that V1's structure is genetically determined. For example, it has long been known that individual orientation-selective cells exist in newborn kittens and ferrets even before they open their eyes (Blakemore and van Sluyters 1975; Chapman and Stryker 1993). Psychological studies further suggest that human newborns can already discriminate between patterns based on orientation (Slater and Johnson 1998; Slater, Morison, and Somers 1988). Recent advances in experimental imaging technologies have made it possible to measure the full map of orientation preferences in young animals. Such experiments show that large-scale orientation maps exist prior to visual experience, and that these maps have many of the same features as found in adults (Chapman, Stryker, and Bonhoeffer 1996; Crair et al. 1998; Gödecke, Kim, Bonhoeffer, and Singer 1997). Furthermore, the global patterns of orientation-selective patches in the maps appear to change very little with normal visual experience, even as the individual neurons gradually become more selective for orientation (Chapman and Stryker 1993; Crair et al. 1998; Gödecke et al. 1997). Thus, despite the clear influence of environmental input on visual cortex structure, normal visual experience appears primarily to preserve and fine-tune the existing structures, rather than drive their development.

Taken together, the evidence indicates that both genetic and environmental influences interact to produce the adult orientation map. However, important questions remain. How does this interaction actually occur? Could adult-like maps develop from environmental or genetic cues alone, or are both necessary? These questions are difficult to answer through biological experiments. Computational modeling, however, can lead to valuable insights, because it is easy to separate environmental and genetic influences in computational experiments. Existing models have been used to simulate how orientation maps can develop from visual input alone (e.g. natural images; Burger and Lang 1999) or genetic factors alone, such as spontaneous neural activity (e.g. noise; Linsker 1986; for review of existing models of each type, see Swindale 1996.) However, prior models have not shown how V1 can have an initial map at birth that becomes smoother and more selective due to postnatal visual experience, while retaining the original map shape. Previous models driven by visual or internal activity alone also differ in many ways besides the source of activity, and thus it has been difficult to determine whether the activity patterns alone account for any differences between the results. This chapter will present results that show clearly how orientation maps can develop both before and after birth in the same model.

## 2.2 Development of face detection

Face perception appears to develop similarly to orientation perception, with evidence for both postnatal learning and function at birth. A number of studies have found that newborns prefer face-like patterns (Goren, Sarty, and Wu 1975; Johnson, Dziurawiec, Ellis, and Morton 1991; Simion, Valenza, and Umiltà 1998), and others show that newborns can discriminate between specific faces and learn to prefer a particular face within the first few hours and days after birth (Bushnell 2001; Pascalis, de Schonen, Morton, Deruelle, and Fabre-Grenet 1995). When shown moving schematic faces in the visual periphery, newborns and one-month-olds will follow them further than other similar patterns (Goren et al. 1975; Johnson et al. 1991; see example schematics in Figure 6*b-e*). However, these abilities change significantly over early development. Older infants do not significantly prefer schematic faces in the visual periphery (Johnson et al. 1991), but between one and two months they begin to respond to facial features in central vision (Maurer and Barrera 1981). Full face-processing abilities take several years to develop (de Haan 2001). Thus, face perception appears to be substantially organized at birth, yet only develops fully through postnatal experience.

Low-level image processing in V1 appears to be able to explain some of the newborn face preferences that have been found, as shown recently in a computational model (Acerra, Burnod, and de Schonen 2002), but many of them appear to require more face-specific circuitry (Bednar 2002). To explain these more specific face preferences, previous theoretical models invoke separate visual processing mechanisms in newborn and older infants. For instance, Johnson and Morton (1991) proposed that infants are born with a basic subcortical system (CONSPEC) that serves only to detect and direct attention to face-like patterns in the periphery, perhaps using a simple three-dot template (two dots for the eyes and a third for the nose and mouth). A separate cortical system (CONLERN) begins to control behavior after one month, and gradually develops more sophisticated face processing, through learning in central vision.

However, models like CONSPEC/CONLERN and that of Acerra et al. (2002) do not account for neonatal face learning. For example, an infant only a few days old prefers to look at its mother's face, relative to the face of a stranger (Bushnell 2001; Pascalis et al. 1995). Presumably this preference results from postnatal visual experience of the mother's face (Bushnell 2001). Unlike the facial feature learning of CONLERN, this early learning has been thought to rely only on the external outline of the face, because the preference disappears when the face outline is masked (Pascalis et al. 1995).

Accordingly, Johnson and Morton (1991) and subsequent authors have proposed adding face-outline learning to CONSPEC, or adding a third, separate subsystem for learning face outlines at birth (de Schonen, Mancini, and Liegeois 1998; Simion et al. 1998). However, recent studies suggest that newborns can also learn internal features (Slater, Bremner, Johnson, Sherwood, Hayes, and Brown 2000). Such learning would seem to require a fourth subsystem, i.e. a learning system similar to CONLERN, but for the periphery, and operational at birth. Alternatively, CONLERN itself might be functional at birth, but in that case the model would need an additional mechanism to explain why schematic faces are no longer preferred in the periphery by two months of age.

Our computational results will show that such increasingly complex models of early face processing are unnecessary. A single cortical learning system present at birth is sufficient to explain the experimental data, if the system is exposed to internally generated patterns broadly similar to

faces.

## 2.3 Internally generated patterns

Our approach relies on the learning of internally generated patterns. Such spontaneous neural activity has recently been documented in many cortical and subcortical areas as they develop, including the retina, the visual cortex, the auditory system, and the spinal cord (reviewed in O'Donovan 1999; Wong 1999). In this work, we focus on two common types of spontaneous activity: retinal waves and ponto-geniculo-occipital (PGO) waves. Retinal waves are found in the developing retina before the eyes first open, and consist of large patches of coherent activity traveling across the retina, surrounded by inactive areas (Feller, Wellis, Stellwagen, Werblin, and Shatz 1996; Maffei and Galli-Resta 1990; Meister, Wong, Baylor, and Shatz 1991). Retinal waves have been shown to be responsible for segregating the LGN into eye-specific layers before birth, indicating that internally generated activity patterns are crucial for the visual system to develop normally (Shatz 1996; Stellwagen and Shatz 2002).

PGO waves are activity patterns that are generated during rapid-eye-movement (REM) sleep. Developing embryos spend a large percentage of their time in a primitive form of REM sleep, which suggests that this state has a major role in development (Roffwarg, Muzio, and Dement 1966). During REM sleep, PGO waves originate in the brain stem, then travel to the LGN, visual cortex, and many other brain areas (see Callaway, Lydic, Baghdoyan, and Hobson 1987 for a review). PGO waves are strongly correlated with eye movements and with vivid visual imagery in dreams, suggesting that they activate the visual system as if they were visual inputs (Marks, Shaffery, Oksenberg, Speciale, and Roffwarg 1995). PGO waves also elicit different distributions of activity in different species, and interrupting them has been shown to increase the influence of the environment on visual system development (Marks et al. 1995).

Together, these findings suggest that spontaneous activity like retinal waves and PGO waves may act as training patterns for development (Jouvet 1998). Retinal waves have been measured directly using optical imaging techniques, but it is not yet known what spatial patterns the PGO waves form (Rector, Poe, Redgrave, and Harper 1997). How closely neonatal REM sleep and PGO waves are related to those of adults is also controversial (Dugovic and Turek 2001; Jouvet 1999). In any case, it is clear that there is substantial activity in the visual areas of newborns during sleep, and that it comes both from the eyes and from the brainstem. These assumptions about the spontaneous activity will be crucial for the experiments in this chapter.

Specifically, the orientation simulations model the effect of retinal waves on the primary visual cortex, based on the hypothesis that such patterns could drive the orientation selectivity and maps seen at birth. The face detection simulations explore how brainstem-generated patterns with higher level structure, such as groups of three similar dots, could drive the development of face detection and face preferences. The simulations show that much of what is known about how orientation maps develop and how newborns detect and learn faces can be explained by a single, general-purpose learning system, which learns from both internally generated patterns of activity and from the visual environment. We propose that this mechanism may be an efficient way to develop a robust system from a small amount of genetic information.

### 3 HLISSOM model

The simulations are based on the HLISSOM model (Bednar 2002). HLISSOM is a version of the LISSOM architecture (Sirosh and Miikkulainen 1994) extended to process spontaneous activity, natural images, and faces. As shown in Figure 2, the model consists of a hierarchy of two-dimensional sheets of neural units modeling different areas of the visual system: a sheet of retinal photoreceptors, a sheet representing internally generated sources of activity (the PGO sheet), a pair of sheets of LGN units (ON-center and OFF-center), and a sheet of cortical units (“neurons”). As in the cortex, the overall arrangement of areas is fixed in HLISSOM, but the smaller scale organization of cortical maps and connections is learned from input activity.

The photoreceptor and PGO input sheets act like grayscale bitmap images, i.e., arrays of numeric values (as shown in Figure 2). The ON-center LGN units respond to patches on the input sheets that are brighter than their surrounds, and the OFF-center units respond to patches darker than their surrounds. Thus the LGN sheets model the edge detection performed by the LGN and the retinal ganglia, highlighting boundaries and edges in the input. The cortical area corresponds to human V1 for the orientation simulations and a higher level area called the face-selective area (FSA) for the face-processing simulations. The FSA represents the first region in the ventral processing pathway that has receptive fields large enough to span a human face at close range. For simplicity, the simulations in this paper use only a single cortical area, either V1 or the FSA, but results are similar with a more complex model where the FSA receives input through a lower level V1 region (Bednar 2002). Each cortical neuron in V1 or the FSA corresponds to a vertical column of cells through the six anatomical layers of the cortex. These neurons learn to represent the important features of the LGN input, based on experience with a series of input patterns. After such self-organization, the cortical neurons respond to particular orientations (for the orientation simulations) or objects (for the face simulations).

The rest of this section reviews the process of self-organization in more detail; see Bednar and Miikkulainen 2003 for a complete description. Each LGN unit computes its response as a product between its weight vector and its receptive fields on the photoreceptors and PGO layers (see Equation 1 in Box 1). The LGN weight vectors are fixed radial Difference-of-Gaussian (DoG) functions. These sombrero-shaped RFs ensure that the ON LGN units respond to light areas surrounded by dark (by subtracting a broad Gaussian from a narrow one) and that OFF LGN units respond to dark areas surrounded by light (by subtracting a narrow Gaussian from a broad one).

Each cortical neuron computes its initial response from its weight vector and its input like an LGN unit does, except that the cortical RFs are on the ON and OFF LGN sheets. After the initial response, the cortical activity settles through short-range excitatory and long-range inhibitory lateral interaction (Equation 2 in Box 1). The contribution from the lateral connections is computed as a similar product between a weight vector and a receptive field, except that the activity of the units in the RF is taken from the preceding settling step. The cortical activity pattern starts out diffuse, but within a few settling iterations, it converges into a small number of stable focused patches of activity, or activity bubbles (as in Figure 2).

After the activity has settled, the connection weights of each cortical neuron are modified. All cortical weights (including the lateral weights) adapt according to the Hebb rule, normalized so that the sum of the weights from each sheet is constant for each neuron (Equation 3 in Box 1). The

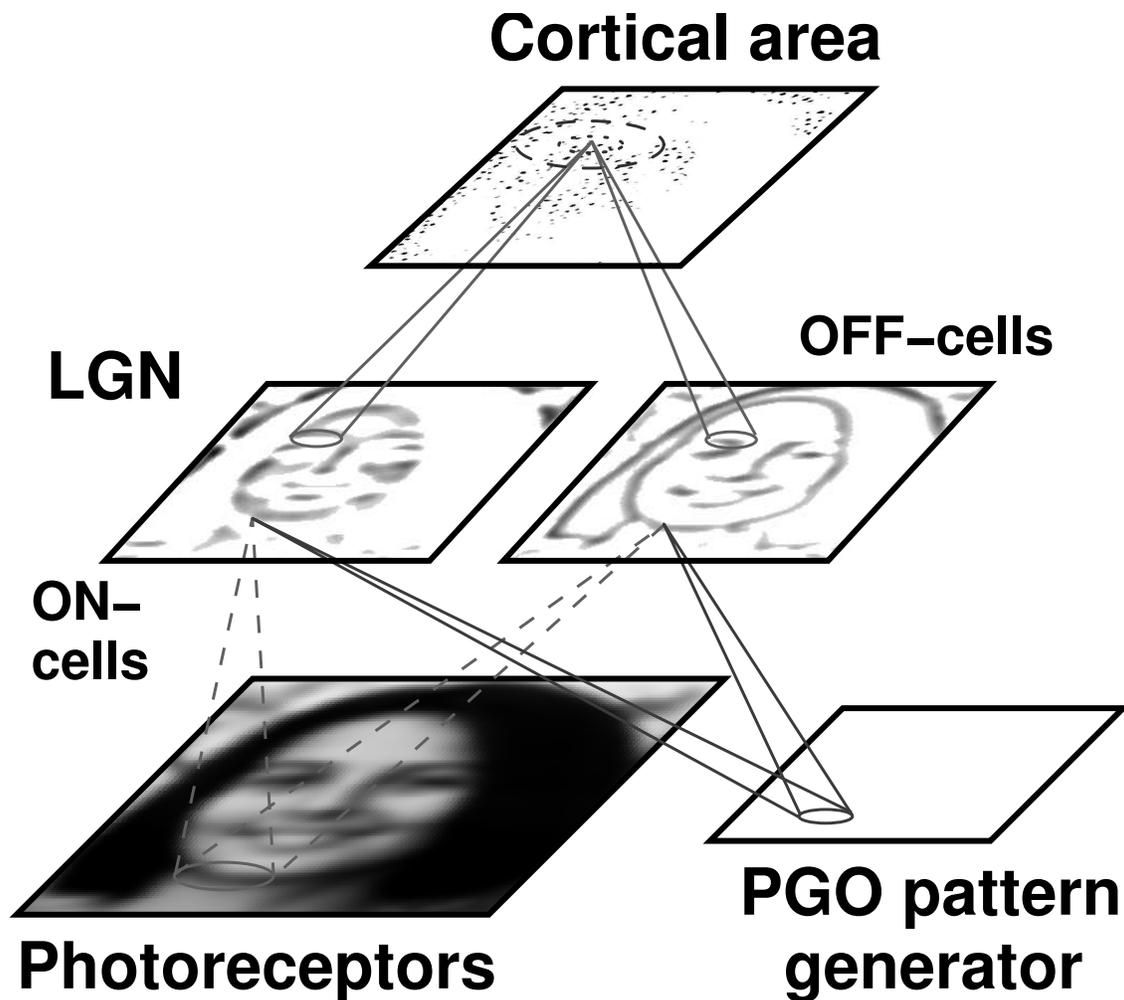


Figure 2: **HLISSOM model.** The model consists of a hierarchy of sheets of neural units, modeling the infant visual pathway. Input such as grayscale images or spontaneous activity patterns is drawn on the photoreceptor array or the PGO array, which represents sources of spontaneous activity. Each higher level unit computes its response from the activity in its receptive fields (RFs); the RFs of one unit in each map are shown. The cortical area corresponds to either V1, for orientation map simulations, or the FSA, for face processing simulations. Cortical units have receptive fields both on previous levels (afferent connections) and on the same level (lateral connections). V1 simulations use a cortical area with small receptive fields to model low-level orientation-selective cells, while FSA simulations use large receptive fields to model face-selective cells. V1 and the FSA can also be arranged hierarchically into the same model, which gives similar results but is more complex to simulate and analyze (Bednar 2002). Photograph copyright 1998, James A. Bednar.

## Box 1: HLISSOM Model Equations

### LGN activation:

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

where:

$(i, j)$  specify an LGN unit

$\eta_{ij}$  is the response of that unit

$\rho$  iterates over the RFs on the input sheets (the photoreceptors and the PGO layer)

$\sigma$  is a piecewise linear sigmoid activation function

$\gamma_{\rho}$  is a constant scaling factor

$X_{\rho ab}$  is the activation of input unit  $(a, b)$  on sheet  $\rho$

$w_{ij, \rho ab}$  is the corresponding weight value

The initial (afferent) response of cortical (V1 or FSA) units is computed using the same equation, but with  $\rho$  iterating over the ON and OFF LGN sheets.

### Cortical activation:

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

where:

$\rho$  iterates over the four RFs to cortical unit  $(i, j)$

(the ON and OFF LGN RFs and the lateral excitatory and inhibitory weights)

$\gamma_{\rho}$  is a constant scaling factor for each  $\rho$  (negative for inhibitory lateral weights)

$X_{\rho ab}(s-1)$  is the activation of input unit  $(a, b)$  during the previous settling step.

### Cortical learning:

$$w_{ij, \rho ab}(f+1) = \frac{w_{ij, \rho ab}(f) + \alpha_{\rho} \eta_{ij} X_{\rho ab}}{\sum_{ab} [w_{ij, \rho ab}(f) + \alpha_{\rho} \eta_{ij} X_{\rho ab}]}, \quad (3)$$

where:

$\eta_{ij}$  is the activity of neuron  $(i, j)$  in the final activity bubble

$w_{ij, \rho ab}(f)$  is the connection weight from the previous iteration

$\alpha$  is the learning rate for each type of connection

$X_{\rho ab}$  is the presynaptic activity.

larger the product of the pre- and post-synaptic activity, the larger the weight change. The normalization ensures that weights do not grow without bound, and approximates homeostatic regulation processes that keep the overall excitability of a neuron constant (reviewed in Turrigiano 1999). At long distances, few neurons have correlated activity and therefore most long-range connections eventually become weak. The weakest connections are eliminated periodically, resulting in patchy lateral connectivity similar to that observed in the visual cortex.

For the orientation map experiments reported in this paper, the ON-center and OFF-center LGN sheets had  $36 \times 36$  units, and received input from a  $54 \times 54$  array. For the face-processing experiments, the ON and OFF sheets were  $74 \times 74$ , and the input array was  $170 \times 170$  units. The V1 orientation map contained  $96 \times 96$  neurons, and the FSA contained  $24 \times 24$  neurons. Initially, the afferent weights of the cortical neurons were random, and the lateral weights had a smooth circular Gaussian profile. The afferent RFs were 13 units wide in V1, and 25 units wide in the FSA (larger because higher level cortical areas have larger RFs). The rest of the parameters are described in Bednar (2002).

## 4 Orientation map experiments

To study how prenatal and postnatal learning together produce the adult orientation map, development was simulated in two phases. In the prenatal phase of 1000 input presentations, input patterns consisted of uniform random noise combined with a large “disk” of activity representing a local patch of highly responding neurons. Figure 3*b* shows an example of this type of pattern. These patterns were chosen to match retinal waves, which are the best-characterized source of spontaneous activity in early development. However, they can also represent any other spontaneous activity that includes large patches that are spatially coherent. The 9000 iterations constituting the postnatal phase used images of natural objects and scenes, modeling visual experience (see Figures 3*d-f*).

This network first develops an orientation map through spontaneous activity, then gradually refines the map based on experience with natural images, without changing the overall shape of the map (Figure 4). The final orientation preference distribution is also biased towards horizontal and vertical, as seen in ferret visual cortex (Coppola et al. 1998). These results are the main novel component of the orientation map simulations. To our knowledge, only one other computational model has been tested in this way (Burger and Lang 1999). Unlike in animals or in HLISSOM, postnatal training in that model entirely rearranged the map, perhaps because their prenatal training patterns did not contain oriented edges like their postnatal patterns did. If the Burger and Lang model is tested with patterns like retinal waves, it may also show the smooth progression from prenatal to postnatal organizations seen in HLISSOM.

Further HLISSOM experiments explored how different pattern types can change the map organization (Bednar 2002). These three experiments act as controls for the main results, helping to determine what features of the model are crucial. First, different postnatal training conditions were tested, to see how different environments can change the outcome. The maps indeed reflected the skewed data sets. For the Landscapes dataset (Figure 3*e*), containing many horizontal contours, the orientation map and histogram became biased towards horizontal (Figure 5*b*). For the Faces

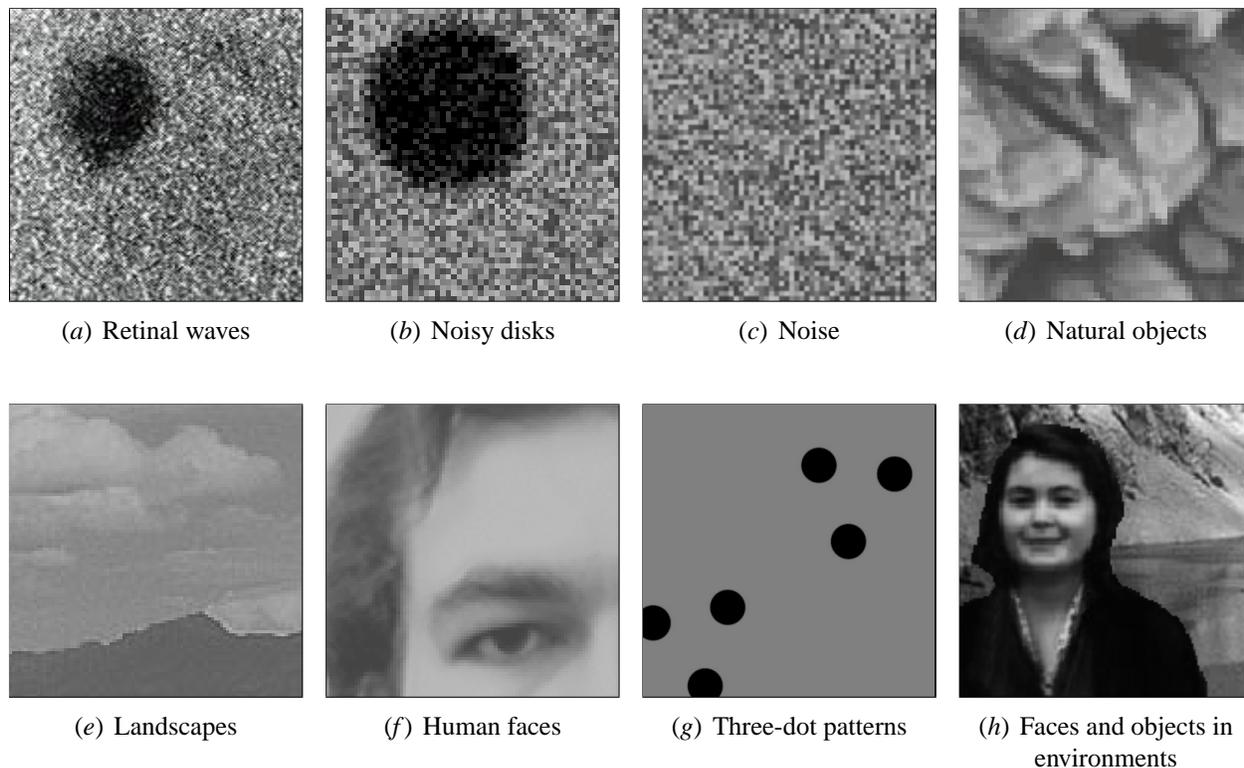


Figure 3: **Training patterns.** Plot (a) shows an example of a retinal wave in the ferret, with active areas appearing dark. (a reprinted with permission from Feller et al. (1996), *Science*, 272:1182–7, copyright 1996 American Association for the Advancement of Science.) The prenatal orientation map simulations are based on an approximation to these patterns (b). As a control, we also tested whether initial orientation maps can develop from uniform random noise alone (c). Postnatal orientation map simulations use natural images (d) or specific types of stimuli, such as natural landscapes (e) or closeups of human faces (f). Prenatal face processing simulations are similar but use more structured inputs, usually without noise to reduce computational requirements (g). Postnatal face processing simulations use photographs of people or objects presented at random locations in front of photographs of natural scenes.

dataset (Figure 3f), more neurons became tuned to vertical orientations (Figure 5c). These findings replicate experimental results from animals raised in biased environments (e.g. Sengpiel et al. 1999), and suggest that the postnatal component of training is important because it allows the map to adapt to the actual environment of the animal (Bednar 2002).

Second, different patterns of prenatal training activity were tested, to determine what features of the noisy disks were crucial. Both the disks alone and random noise alone (Figure 3c) resulted in orientation maps with locally coherent patches of orientation-selective neurons. However, the receptive fields in the disks case were much more selective than in typical newborn neurons. Thus including noise was crucial for obtaining realistic receptive fields. Yet for uniform random noise alone, the receptive fields were poorly selective, and primarily included receptive field types rarely reported in experimental studies of V1 (e.g. crosses; Bednar 2002). The orientation map was also less well organized than typical newborn maps. These results suggest that both spatially uncorrelated noise and spatially correlated activity are necessary to develop a normal prenatal

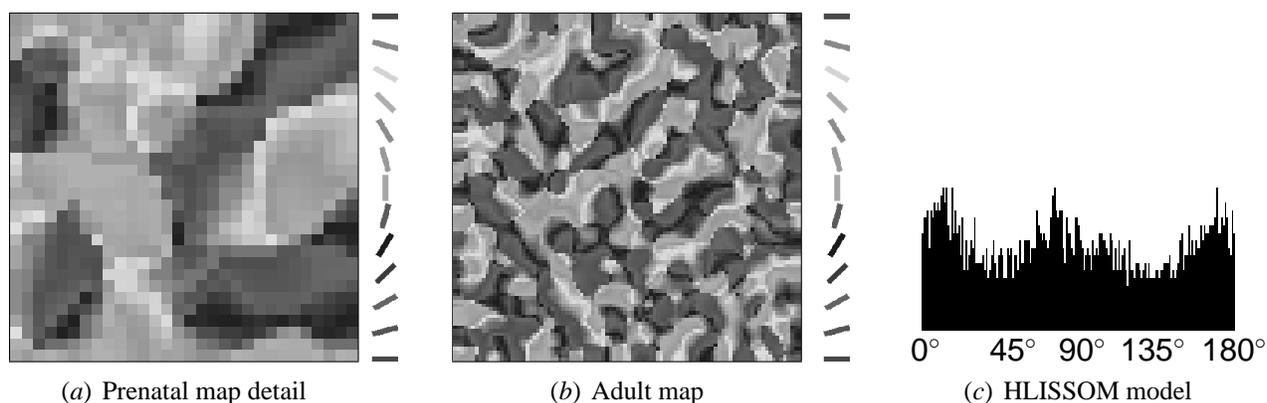


Figure 4: **HLISSOM model orientation maps.** By iteration 1000, the end of prenatal training, the HLISSOM map was a good match to newborn maps (compare the central area plotted in *a* with Figure 1*a*). Further training on 9000 natural image patches results in an adult-like orientation map (compare *b* with Figure 1*b*). The statistics of the orientation preferences of the adult map reflect the horizontal and vertical biases of the training images, causing the map to devote more area to horizontal and vertical (compare *c* with Figure 1*c*).

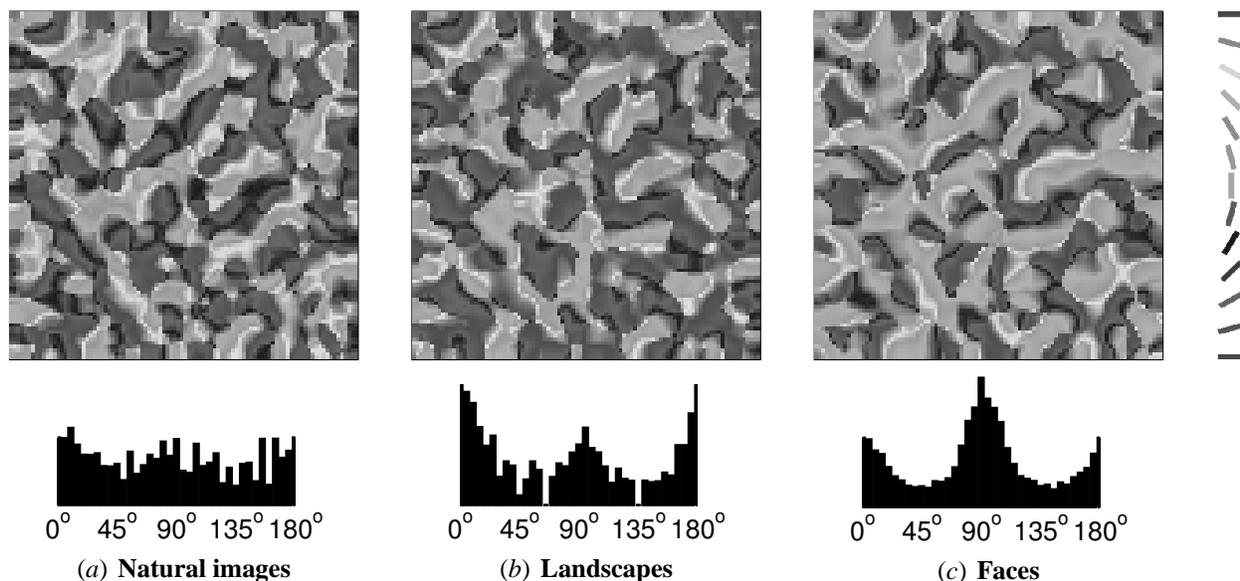


Figure 5: **Effect of different postnatal training environments.** Three copies of the same prenatally organized map were trained postnatally with different image datasets: **Nature** (as shown in Figure 4*c*), **Landscapes** (a collection of nature scenes including many strong horizontal contours), and **Faces** (a collection of frontal views of faces). Although the resulting maps all have similar large-scale organization determined by the prenatal training, the number of neurons selective for each orientation reflects the distribution of postnatal training patterns. Thus prenatal training provides the basic structure, and postnatal training helps the map represent the most common patterns in the environment.

map.

Third, an initially random map was trained on natural images alone, to test whether the prenatal training was crucial for developing orientation maps and selectivity. Although the details of the final map organization differed from the normal case, the overall properties were very similar (Bednar 2002). Thus prenatal training is not required for a functional orientation map to develop from natural image input. Overall, these simulations show how internally generated patterns can explain orientation map function at birth, while the same system later learns from real visual images.

## 5 Face preference experiments

The face preference simulations were similar to the orientation map simulations, but with more structured training patterns. For simplicity, the simulations bypassed V1, but including it leads to similar results (Bednar 2002; Bednar and Miikkulainen 2003). As for orientation, a variety of prenatal training conditions were simulated, to determine how different training patterns can lead to face preferences. Most simulations used the three-dot patterns shown in Figure 3g, presented at randomly chosen positions and random near-vertical orientations. For control purposes, networks were also trained using other internally generated patterns, and an untrained (naïve) network was also tested. Postnatally, the training regime represented gradual learning of specific individuals and objects seen against a variety of different backgrounds. For each postnatal iteration, one of six face images or six object images was chosen randomly, and presented in front of a randomly chosen natural scene background (Bednar 2002; see Figure 3h). The main goals of these simulations were to test whether prenatal learning could account for human performance at birth, and whether this learning shapes postnatal learning and the final self-organized architecture.

After prenatal training, and at intervals throughout postnatal training, the network was tested using schematic images previously tested with newborns, and with photographs of faces. In order to compare the neural activity levels in the model to the patterns preferred by babies, we assumed that an infant pays more attention to the stimuli that are most effective at activating its visual processing system, focusing on the highest level activated (Bednar 2002). Patterns activating the FSA will be preferred over those activating only lower areas, and two patterns that both activate the FSA will be ranked by their total FSA activity. Under these assumptions, it is possible to compare the model results with psychological data from human infants.

When trained on three-dot patterns at various locations and orientations, FSA neurons developed receptive fields preferring upright, triangular arrangements of three dots (Figure 6a, row RF 10000). These RFs cause the network to respond to facelike schematic images (Figure 6), with total activity levels that rank the patterns in the order preferred by infants (Goren et al. 1975; Johnson and Morton 1991). When tested on 18 schematic patterns from newborn studies (Goren et al. 1975; Johnson and Morton 1991; Simion et al. 1998; Valenza, Simion, Cassia, and Umiltà 1996), the full version of the model (including both V1 and the FSA) ranked them in the same preference order for all 22 of the statistically significant preferences found in newborns (Bednar 2002; Bednar and Miikkulainen 2003). These results demonstrate that a network exposed to three-dot patterns is sufficient to explain the experimental results with newborns tested with schematic stimuli.

Although schematic patterns are usually used in newborn studies for practical reasons, it is

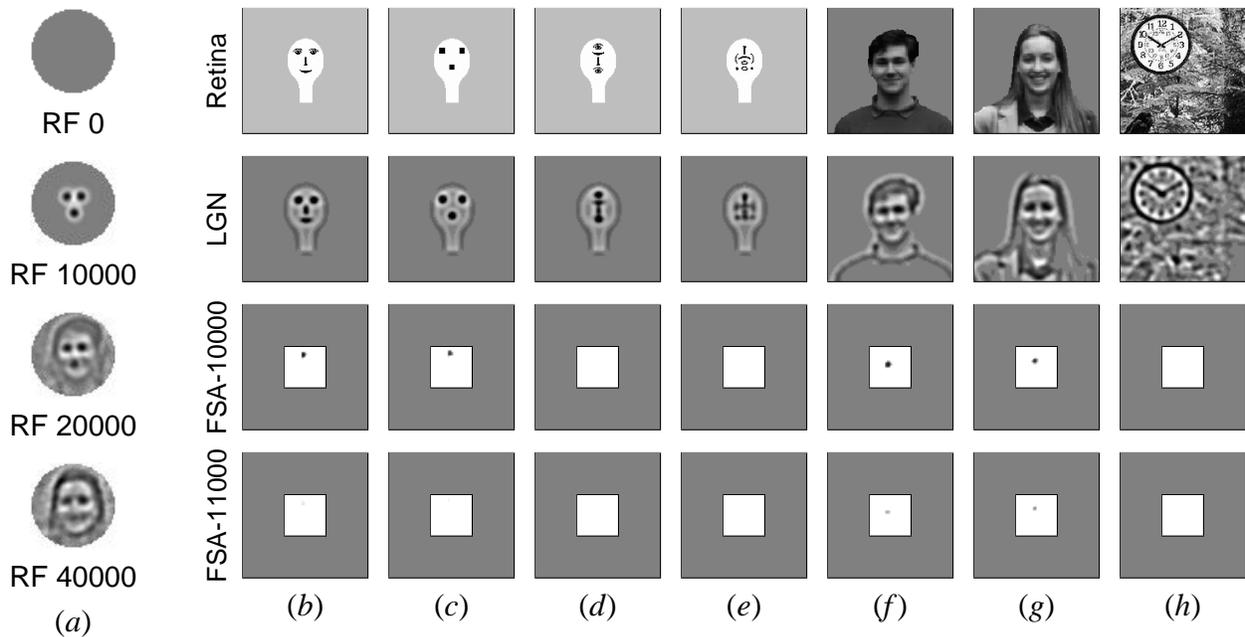


Figure 6: **Model RFs and responses.** The plots in (a) show the RF for a typical neuron in the FSA at different stages of training; the pattern shown is the LGN activity that would most excite this neuron. At first the neuron is unselective (RF 0), but after prenatal exposure to many three-dot patterns it comes to prefer them (RF 10000). The RF then gradually becomes more selective for real faces through postnatal experience (RF 20000 and RF 40000). The remaining plots show how such neurons respond to test images. The top row shows four schematic patterns (Johnson and Morton 1991) and two photographic stimuli (Rowley et al. 1998). Below each is the LGN response to that image (second row), the FSA response after prenatal training on three-dot patterns (third row), and the FSA response after postnatal training on real faces (bottom row). The “neonatal” FSA (row FSA-10000) responds in the location of face-like schematic and real stimuli, and rarely responds to other images. As in two-month-olds, the response to schematic patterns eventually disappears when the network is trained on real images (row FSA-11000), but the responses to real faces remain (f-g). For images without faces, such as the training pattern shown in (h), the network responds only when there are accidental three-dot patterns.

important to show that the HLISSOM model also exhibits preferences for real faces compared to other photographs. Given a database of 150 top-lit images of adult males, the model responded in the correct location (the center of the face) to 88% of the images (Bednar and Miikkulainen 2003). Conversely, it responded to only 4% of the images in a database of natural scenes. Overall, the FSA responded to most top-lit human faces, signaling their location in the visual field. It did not respond to most other stimuli, except when they contained accidental three-dot patterns. The model predicts that human newborns will have a similar pattern of responses in the face-selective cortical regions.

How crucial is the specific training pattern for the prenatal preferences? In addition to the three-dot pattern proposed by Johnson and Morton (1991), we tested a variety of other prenatal training pattern shapes (Bednar 2002; Bednar and Miikkulainen 2003). For all training patterns, test faces that matched the training pattern size gave higher responses than did other natural images. However, responses to the schematic patterns were very different for networks trained with different patterns. Of all training patterns tested, the three-dot pattern provided the best selectivity

for faces, and a better fit to the experimental data with schematics.

Importantly, not all patterns result in face preferences. For example, if only a single low-level feature like eye size is matched between the training and testing images, the HLISSOM model does not prefer faces. These results show that the pattern shape is important, but that a variety of shapes lead to qualitatively similar results.

What effect does the prenatal training have on postnatal learning? Trained with real faces in real images, most units in both the prenatally trained and naïve networks developed RFs that are averages (i.e. prototypes) of faces and hair outlines (Bednar and Miikkulainen 2002). However, the time course of learning and the final result differed. RFs in the prenatally trained network become gradually more selective for faces. Postnatal self-organization in the naïve network is less regular, and the final RFs are usually less selective. Overall, prenatal training appears to bias postnatal learning towards biologically relevant stimuli, i.e. faces (as suggested by Johnson and Morton 1991).

The bottom row of Figure 6 shows that HLISSOM gradually responds less to schematic faces, just as infants no longer prefer such faces in the periphery after one month (Johnson et al. 1991). In HLISSOM, the loss of face preferences is caused by afferent weight normalization. As the FSA neurons in HLISSOM learn the hair and face outlines typically associated with real faces, the connections to the internal features necessarily become weaker. As a result, the response to face-like schematic patterns decreases, because those patterns match only the internal features. Eventually, the response to the schematic patterns drops below the fixed activation threshold. At that point, the model does not respond more strongly to schematic faces than to other schematics. In a sense, the FSA has learned that real faces typically have both inner *and* outer features, and does not respond when either type of feature is absent or a poor match to real faces.

Interestingly, the FSA neurons continue to respond to real faces (as opposed to schematics) throughout postnatal learning (Figure 6f-g). Thus the model provides a clear prediction that even if infants lose their preference for schematics over time, they will continue to prefer realistic faces compared with other similar images.

When one face (i.e. the mother) appears most often, the FSA response to that face becomes stronger than to a similar stranger (Figure 7; Bednar and Miikkulainen (2002)). This result replicates the mother preference found in infants a few days old (Bushnell 2001; Pascalis et al. 1995). Interestingly, the model no longer prefers the mother when her hair outline is masked, which is consistent with Pascalis et al.'s claim that newborns learn outlines only. However, Pascalis et al. (1995) did not test the crucial converse condition, i.e. whether newborns respond when the facial features are masked, leaving only the outlines. As for the internal features alone, the model does not respond to the head and hair outline alone. Thus the HLISSOM model's learning of faces is clearly *not* restricted to the outline.

In the model, the response is lower with either type of masking because of holistic learning of *all* of the features of real faces. As real faces are learned, the afferent weight normalization ensures that neurons respond only to patterns that are a good overall match to all of the weights, not simply matching on a few features. Many authors have argued that adults learn faces holistically (e.g. Farah, Wilson, Drain, and Tanaka 1998). The simulation results suggest that newborns may learn faces in the same way, and thus that they will not prefer their mother when her internal facial features are masked. Overall, the face preference experiments provide a detailed computational

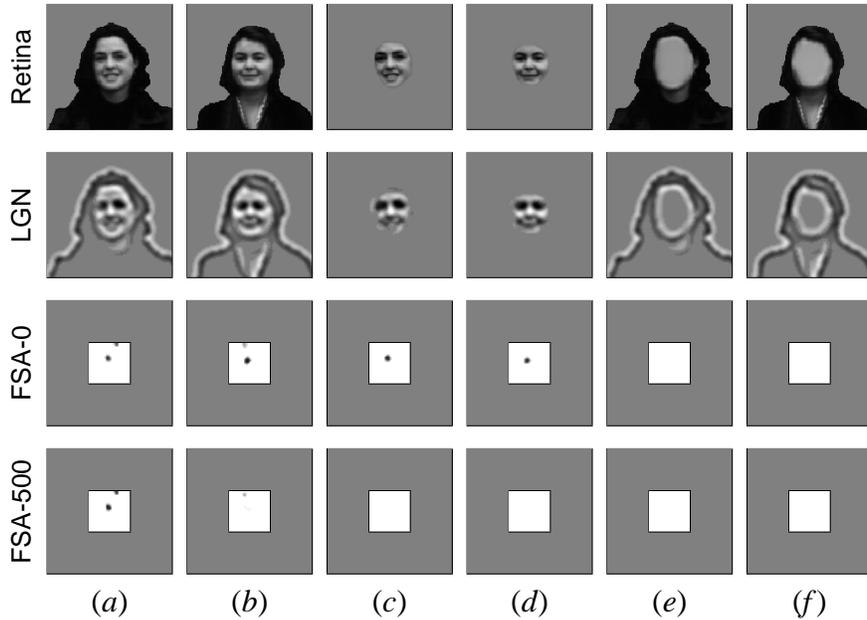


Figure 7: **Mother preferences depend on both internal and external features.** Initially, the prenatally trained FSA responds well to both women above (*a-b*; FSA-0), with no significant difference ( $p = 0.28$ ). The response is primarily due to the internal facial features (*c-d*; FSA-0), although there are some spurious three-dot responses due to alignment of the hair with the eyes (*a-b*; top of FSA-0). Designating image (*a*) as the mother, we presented it in 25% of the postnatal learning iterations. (This ratio is taken from Bushnell 2001, who found that newborns look at their mother’s face for an average of about one-fourth their time awake over the first few days.) Image (*b*), the stranger, was not presented at all during training. After 500 postnatal iterations, the response to the mother is significantly greater than to face (*b*) ( $p = 0.001$ ). This result replicates the mother preference found by Pascalis et al. (1995) in infants 3–9 days old. The same results are found in the counterbalancing condition — when trained on face (*b*) as the mother, (*b*) becomes preferred ( $p = 0.002$ ; not shown). After training with real faces, there is no longer any FSA response to the facial features alone (*c-d*), which replicates Pascalis et al.’s (1995) finding that newborns no longer preferred their mother when her face outline was covered. Yet contra Pascalis et al. (1995), we cannot conclude that what has been learned “has to do with the outer rather than the inner features of the face”, because no preference is found for the face outline alone either (*e-f*). These results suggest that newborn face learning is holistic, and includes all features associated with faces.

simulation of how preferences can be constructed prenatally and refined postnatally, and provide concrete predictions for future experiments.

## 6 Discussion and future work

The results presented in this chapter show how environmental and genetic factors can interact to produce selectivity for both high-level and low-level visual features. Both environmental and genetic factors are crucial for explaining the experimental data, and both are incorporated using the same general-purpose learning algorithm. In each case, the prenatal organization provides structure at birth, while postnatal learning allows these initial preferences to be refined to match

the actual visual patterns in the environment.

The orientation simulations show that the development of low-level features like orientation maps is robust, i.e. maps can develop from a wide range of possible input patterns. Spontaneous activity with actual oriented edges, like retinal waves, may help ensure that maps develop normally in a wider range of possible environments. Prenatal training may also allow V1 to mature more quickly, which could be important for organizing higher levels of the visual system.

The face preference simulations show that prenatal learning of spontaneous activity and postnatal learning of real faces can together account for newborn face preferences, neonatal face learning, and longer term development of face detection. The results suggest simple but novel explanations for why newborns prefer schematic faces, why newborn learning appears to depend on the face outline, and why the response to schematic faces decreases over time. They also suggest that internally generated patterns allow the genome to steer development towards biologically relevant processing, making subsequent learning more robust.

The orientation and face preference results depend primarily on having a model that learns representations for common patterns on the input, plus having prenatal patterns that share some structure with those experienced postnatally. Thus we expect that similar results could be found with other realistic models of visual processing, such as those of Burger and Lang (1999; for orientation) and Rolls and Milward (2000; for face recognition). In each case, the crucial assumptions are (1) the type of patterns that are present before visual experience, (2) the overall arrangement of visual areas, and (3) the “hardwired” processing in the early visual areas.

In future work, it will be important to measure the actual activity patterns in developing animals and humans, and understand how these patterns are generated in the brainstem and other areas. Brain imaging in young infants can help determine whether face-selective responses in newborns are subcortical, as in the CONSPEC/CONLERN model, or cortical, as in the HLISSOM model. Even before newborn imaging is practical, the predictions of the face processing model can be tested behaviorally with infant experiments. For example, HLISSOM predicts that over the first two months the response to real faces in the periphery should continue, even as response to schematics diminishes, and the mother preference of newborns should disappear when the facial features are masked. These experiments will help clarify the roles of internally generated activity in the development of the cortex.

## 7 Conclusion

The HLISSOM model shows how the interaction between environmental and genetic factors can explain many features of orientation map and face processing development. In the model, the genetic factors take several forms: (1) constraints on the overall arrangement of visual areas and connectivity between areas, (2) prespecification of processing in early visual areas such as the retina and the LGN, and (3) genetically determined activity patterns generated during early development. Given these constraints and initial patterns, the model shows how a functional visual system can develop reliably, even without detailed genetic specification of neural connectivity patterns or of specific cortical representations.

At the same time, the specific representations that do develop depend strongly on interaction

with the environment. Without environmental input, the model representations remain crude and relatively unselective, unlike neurons in adult cortical areas. Moreover, the specific properties of the environmental input are crucial: a preponderance of certain orientations or objects will result in representations biased towards common stimuli, such as horizontal contours or the mother's face. This adaptation helps ensure that appropriate representations will develop.

In the model, the environmental interactions were simulated using static photographic images, but in general an organism will develop through interaction with a much richer environment, including moving patterns, associations between different sensory and motor modalities, and associations between sensory patterns and rewards or pain. We expect that these additional consequences of embodiment or situatedness will be crucial for explaining the full range of cortical development. For instance, experience with faces from a wide range of orientations, in different lighting conditions, and with different hairstyles is likely to be necessary for the development of adult-like face recognition that is invariant to these factors. Future work will explore how these more realistic types of embodiment can be studied computationally.

In summary, the specific simulations described above demonstrate how environmental and genetic factors can together explain much of how orientation and face processing develop. In both cases, the overall approach is the same, with initial orientation and face selectivity developing prenatally and being refined postnatally to match the visual environment. The modeling results provide clear predictions for future animal and infant experiments, and they act as examples of a new way of understanding how sensory systems can be constructed.

## Acknowledgments

This research was supported in part by the National Science Foundation under grants IRI-9309273 and IIS-9811478, and by the National Institute of Mental Health under Human Brain Project grant R01-MH66991-01. Software, demonstrations, and related publications are available at <http://nn.cs.utexas.edu>.

## Glossary

**Piecewise-linear sigmoid activation function:** A neural network unit generates a real-valued output (its *activation*) as a function of the values of its inputs. A sigmoid function is an S-shaped curve that produces little activation for weak inputs, nearly linear activation for intermediate values, and saturates quickly for high values. A piecewise-linear sigmoid is an approximation to the smooth sigmoid curve, using linear functions for the low, medium, and high regimes, which is more computationally efficient to simulate.

**topographic neural maps:** Cortical sensory areas are organized in terms of topographic maps of neurons. For instance, most visual areas are retinotopic, mapping from input organized spatially in two dimensions like the retina, into another retinotopic representation. Other sensory and motor areas also have two-dimensional topographic organization.

**ventral processing pathway:** Ungerleider and Mishkin (1982) proposed that the organization of the visual system could be understood in terms of a *ventral stream* and a *dorsal stream*, also called the *what* and the *where* pathways. Areas in the ventral stream, such as IT (inferotemporal cortex), are thought to be concerned with object identity. Those in the dorsal stream, such as posterior parietal cortex, are thought to be concerned primarily with precise object location and the detection of motion. The two streams share early cortical areas such as V1 and V2.

**Hebb rule:** Donald Hebb (1949) proposed a neural network learning rule that is now called the Hebb rule or the Hebbian rule. He suggested that neurons that often activate at the same time should become more strongly connected, an idea that has since received wide experimental and theoretical support. A variety of mathematical equations can be characterized as Hebbian, as long as they result in connections being strengthened between coactive neurons.

## References

- Acerra F, Burnod Y, and de Schonen S (2002). Modelling aspects of face processing in early infancy. *Developmental Science*, 5(1), 98–117.
- Bednar JA (2002). *Learning to See: Genetic and Environmental Influences on Visual Development*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin. Technical Report AI-TR-02-294.
- Bednar JA and Miikkulainen R (2002). Neonatal learning of faces: Interactions between genetic and environmental inputs. In Gray W and Schunn C, editors, *Proceedings of the 24th Annual Conference of the Cognitive Science Society*, pp. 107–112. Erlbaum, Mahwah, NJ.
- Bednar JA and Miikkulainen R (2003). Learning innate face preferences. *Neural Computation*, 15(7), 1525–1557.
- Blakemore C and Cooper GF (1970). Development of the brain depends on the visual environment. *Nature*, 228, 477–478.
- Blakemore C and van Sluoyers RC (1975). Innate and environmental factors in the development of the kitten's visual cortex. *Journal of Physiology (London)*, 248, 663–716.
- Blasdel GG (1992). Orientation selectivity, preference, and continuity in monkey striate cortex. *Journal of Neuroscience*, 12, 3139–3161.
- Burger T and Lang EW (1999). An incremental Hebbian learning model of the primary visual cortex with lateral plasticity and real input patterns. *Zeitschrift für Naturforschung C—A Journal of Biosciences*, 54, 128–140.
- Bushnell IWR (2001). Mother's face recognition in newborn infants: Learning and memory. *Infant and Child Development*, 10(1/2), 67–74.
- Callaway CW, Lydic R, Baghdoyan HA, and Hobson JA (1987). Pontogeniculooccipital waves: Spontaneous visual system activity during rapid eye movement sleep. *Cellular and Molecular Neurobiology*, 7(2), 105–49.
- Chapman B and Bonhoeffer T (1998). Overrepresentation of horizontal and vertical orientation

- preferences in developing ferret area 17. *Proceedings of the National Academy of Sciences, USA*, 95, 2609–2614.
- Chapman B and Stryker MP (1993). Development of orientation selectivity in ferret primary visual cortex and effects of deprivation. *Journal of Neuroscience*, 13(12), 5251–5262.
- Chapman B, Stryker MP, and Bonhoeffer T (1996). Development of orientation preference maps in ferret primary visual cortex. *Journal of Neuroscience*, 16(20), 6443–6453.
- Coppola DM, White LE, Fitzpatrick D, and Purves D (1998). Unequal representation of cardinal and oblique contours in ferret visual cortex. *Proceedings of the National Academy of Sciences, USA*, 95(5), 2621–2623.
- Crair MC, Gillespie DC, and Stryker MP (1998). The role of visual experience in the development of columns in cat visual cortex. *Science*, 279, 566–570.
- de Haan M (2001). The neuropsychology of face processing during infancy and childhood. In Nelson CA and Luciana M, editors, *Handbook of Developmental Cognitive Neuroscience*. MIT Press.
- de Schonen S, Mancini J, and Liegeois F (1998). About functional cortical specialization: The development of face recognition. In (Simion and Butterworth 1998), pp. 103–120.
- Diamond S (1974). Four hundred years of instinct controversy. *Behavior Genetics*, 4, 237–252.
- Dugovic C and Turek FW (2001). Similar genetic mechanisms may underlie sleep-wake states in neonatal and adult rats. *Neuroreport*, 12(14), 3085–3089.
- Farah MJ, Wilson KD, Drain M, and Tanaka JN (1998). What is “special” about face perception? *Psychological Review*, 105(3), 482–498.
- Feller MB, Wellis DP, Stellwagen D, Werblin FS, and Shatz CJ (1996). Requirement for cholinergic synaptic transmission in the propagation of spontaneous retinal waves. *Science*, 272, 1182–1187.
- Gödecke I, Kim DS, Bonhoeffer T, and Singer W (1997). Development of orientation preference maps in area 18 of kitten visual cortex. *European Journal of Neuroscience*, 9(8), 1754–1762.
- Goren CC, Sarty M, and Wu PY (1975). Visual following and pattern discrimination of face-like stimuli by newborn infants. *Pediatrics*, 56(4), 544–549.
- Hebb DO (1949). *The Organization of Behavior: A Neuropsychological Theory*. Wiley, New York.
- Hubel DH and Wiesel TN (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *Journal of Physiology (London)*, 160, 106–154.
- Hubel DH and Wiesel TN (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London)*, 195, 215–243.
- Johnson MH, Dziurawiec S, Ellis H, and Morton J (1991). Newborns’ preferential tracking of face-like stimuli and its subsequent decline. *Cognition*, 40, 1–19.
- Johnson MH and Morton J (1991). *Biology and Cognitive Development: The Case of Face Recognition*. Blackwell, Oxford, UK; New York.
- Jouvet M (1998). Paradoxical sleep as a programming system. *Journal of Sleep Research*, 7(Suppl 1), 1–5.
- Jouvet M (1999). *The Paradox of Sleep: The Story of Dreaming*. MIT Press, Cambridge, MA.
- Kandel ER, Schwartz JH, and Jessell TM (1991). *Principles of Neural Science*. Elsevier, New

York. Third edition.

- Kanwisher N, McDermott J, and Chun MM (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17(11), 4302–4311.
- Linsker R (1986). From basic network principles to neural architecture: Emergence of orientation columns. *Proceedings of the National Academy of Sciences, USA*, 83, 8779–8783.
- Maffei L and Galli-Resta L (1990). Correlation in the discharges of neighboring rat retinal ganglion cells during prenatal life. *Proceedings of the National Academy of Sciences, USA*, 87, 2861–2864.
- Marks GA, Shaffery JP, Oksenberg A, Speciale SG, and Roffwarg HP (1995). A functional role for REM sleep in brain maturation. *Behavioural Brain Research*, 69, 1–11.
- Maurer D and Barrera M (1981). Infants' perception of natural and distorted arrangements of a schematic face. *Child Development*, 52(1), 196–202.
- Meister M, Wong ROL, Baylor DA, and Shatz CJ (1991). Synchronous bursts of action-potentials in the ganglion cells of the developing mammalian retina. *Science*, 252, 939–943.
- Miyashita-Lin EM, Hevner R, Wassarman KM, Martinez S, and Rubenstein JL (1999). Early neocortical regionalization in the absence of thalamic innervation. *Science*, 285(5429), 906–909.
- Movshon JA and van Sluylers RC (1981). Visual neural development. *Annual Review of Psychology*, 32, 477–522.
- O'Donovan MJ (1999). The origin of spontaneous activity in developing networks of the vertebrate nervous system. *Current Opinion in Neurobiology*, 9, 94–104.
- Pascalis O, de Schonen S, Morton J, Deruelle C, and Fabre-Grenet M (1995). Mother's face recognition by neonates: A replication and an extension. *Infant Behavior and Development*, 18, 79–85.
- Rakic P (1988). Specification of cerebral cortical areas. *Science*, 241, 170–176.
- Rector DM, Poe GR, Redgrave P, and Harper RM (1997). A miniature CCD video camera for high-sensitivity light measurements in freely behaving animals. *Journal of Neuroscience Methods*, 78(1-2), 85–91.
- Roffwarg HP, Muzio JN, and Dement WC (1966). Ontogenetic development of the human sleep-dream cycle. *Science*, 152, 604–619.
- Rolls ET and Milward T (2000). A model of invariant object recognition in the visual system: Learning rules, activation functions, lateral inhibition, and information-based performance measures. *Neural Computation*, 12(11), 2547–2572.
- Rowley HA, Baluja S, and Kanade T (1998). Neural network-based face detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(1), 23–38.
- Sengpiel F, Stawinski P, and Bonhoeffer T (1999). Influence of experience on orientation maps in cat visual cortex. *Nature Neuroscience*, 2(8), 727–732.
- Shatz CJ (1996). Emergence of order in visual system development. *Proceedings of the National Academy of Sciences, USA*, 93, 602–608.
- Simion F and Butterworth G, editors (1998). *The Development of Sensory, Motor and Cognitive*

- Capacities in Early Infancy: From Perception to Cognition*. Psychology Press, East Sussex, UK.
- Simion F, Valenza E, and Umiltà C (1998). Mechanisms underlying face preference at birth. In (Simion and Butterworth 1998), pp. 87–102.
- Sirosh J and Miikkulainen R (1994). Cooperative self-organization of afferent and lateral connections in cortical maps. *Biological Cybernetics*, 71, 66–78.
- Slater A, Bremner G, Johnson SP, Sherwood P, Hayes R, and Brown E (2000). Newborn infants' preference for attractive faces: The role of internal and external facial features. *Infancy*, 1(2), 265–274.
- Slater A and Johnson SP (1998). Visual sensory and perceptual abilities of the newborn: Beyond the blooming, buzzing confusion. In (Simion and Butterworth 1998), pp. 121–142.
- Slater A, Morison V, and Somers M (1988). Orientation discrimination and cortical function in the human newborn. *Perception*, 17, 597–602.
- Stellwagen D and Shatz CJ (2002). An instructive role for retinal waves in the development of retinogeniculate connectivity. *Neuron*, 33(3), 357–367.
- Sur M, Angelucci A, and Sharma J (1999). Rewiring cortex: The role of patterned activity in development and plasticity of neocortical circuits. *Journal of Neurobiology*, 41, 33–43.
- Sur M and Leamey CA (2001). Development and plasticity of cortical areas and networks. *Nature Reviews Neuroscience*, 2(4), 251–262.
- Swindale NV (1996). The development of topography in the visual cortex: A review of models. *Network – Computation in Neural Systems*, 7, 161–247.
- Switkes E, Mayer MJ, and Sloan JA (1978). Spatial frequency analysis of the visual environment: Anisotropy and the carpentered environment hypothesis. *Vision Research*, 18(10), 1393–1399.
- Turrigiano GG (1999). Homeostatic plasticity in neuronal networks: The more things change, the more they stay the same. *Trends in Neurosciences*, 22(5), 221–227.
- Ungerleider LG and Mishkin M (1982). Two cortical visual systems. In Ingle DJ, Goodale MA, and Mansfield RJW, editors, *Analysis of Visual Behavior*, pp. 549–586. MIT Press, Cambridge, MA.
- Valenza E, Simion F, Cassia VM, and Umiltà C (1996). Face preference at birth. *Journal of Experimental Psychology: Human Perception and Performance*, 22(4), 892–903.
- von Melchner L, Pallas SL, and Sur M (2000). Visual behaviour mediated by retinal projections directed to the auditory pathway. *Nature*, 404(6780), 871–876.
- Wong ROL (1999). Retinal waves and visual system development. *Annual Review of Neuroscience*, 22, 29–47.