

## Neonatal Learning of Faces: Environmental and Genetic Influences

James A. Bednar (jbednar@cs.utexas.edu)

Risto Miikkulainen (risto@cs.utexas.edu)

Department of Computer Sciences, The University of Texas at Austin  
Austin, TX 78712 USA

### Abstract

Newborn face perception is controversial, but the current evidence suggests that (a) newborns follow face-like schematic patterns further than similar patterns, (b) infants can learn individual faces soon after birth, and (c) full face processing abilities develop through months or years of experience with faces. Previous models have not adequately accounted for all three types of results. In prior work, we showed how a biologically based self-organizing system and spontaneous activity patterns can explain newborn face preferences. In this paper we show that this general-purpose learning system can explain both neonatal and later learning. Using computational simulations, we demonstrate that newborn learning need not be based on the external outline, as has been supposed, and that postnatal decreases in response to schematic faces need not represent a decrease in response to real faces. These simulations provide concrete predictions to guide future experiments with infants, while suggesting new techniques for designing complex adaptive systems in general.

### Introduction

Specific regions in the adult visual cortex respond preferentially to human faces. How this face processing capability develops is not yet clear. Many researchers have argued that infants process only general visual properties like size and spatial frequency until after weeks or months of experience (Maurer & Barrera, 1981). Others have found a preference for faces at birth (Goren, Sarty, & Wu, 1975; Johnson, Dziurawiec, Ellis, & Morton, 1991; Simion, Valenza, & Umiltà, 1998) and that infants can learn and discriminate between specific faces even in the first few hours and days after birth (Bushnell, 2001; Pascalis, de Schonen, Morton, Deruelle, & Fabre-Grenet, 1995). Full face processing abilities clearly take several years to develop.

In this paper we show that a single learning system can account for all three types of results, i.e. face preferences at birth, face learning in the first few days after birth, and the gradual development of full face processing abilities. Using the HLISSOM self-organizing model (Hierarchical Laterally Interconnected Synergetically Self-Organizing Map), we have previously shown how prenatal learning of spontaneous neural activity can lead to newborn face preferences (Bednar & Miikkulainen, 2000). In this paper we show that the same self-organizing system can learn from faces in real images, and that the learning process can explain postnatal changes in infant face detection abilities. Together these simulations show how genetic information can be expressed within a highly adaptive system, and provide concrete predictions for future experiments with infants.

### Development of face detection

Face detection abilities change significantly between birth and two months of age. 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 5a-d). Older infants do not show a peripheral schematic face preference (Johnson et al., 1991) but between one and two months they begin to respond to facial features in central vision (Maurer & Barrera, 1981).

Previous models invoke separate visual processing mechanisms for these newborn and later face preferences. For instance, Johnson and Morton (1991) proposed that infants are born with a simple subcortical system they termed CONSPEC. CONSPEC 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 would begin to control behavior after one month, and would gradually develop more sophisticated face processing through learning in central vision.

However, the CONSPEC/CONLERN model does not account for neonatal face learning. For instance, an infant only a few days old will prefer to look at its mother's face, relative to the face of a stranger (Pascalis et al., 1995). This mother preference has been thought to involve the external outline of the face only, in contrast to the internal facial feature learning of CONLERN, because the preference disappears when internal features are masked (Pascalis et al., 1995).

Accordingly, Johnson and Morton (1991) and subsequent authors have proposed extending the CONSPEC/CONLERN model to include face-outline learning in CONSPEC, or a third, separate subsystem for learning face outlines at birth (de Schonen, Mancini, & Leigeois, 1998; Simion et al., 1998). However, recent studies suggest that newborns can also learn internal features (Slater, Bremner, Johnson, Sherwood, Hayes, & Brown, 2000). Such learning could require a fourth subsystem, like CONLERN but for the periphery and operational at birth. (CONLERN itself cannot explain newborn learning of internal features, because were it present at birth, it would no longer explain the shift from peripheral to central face preferences after one month.)

We will show that such increasingly complex models are unnecessary. A single, CONLERN-like system processing the entire visual field is sufficient to explain the experimental data, if CONSPEC is replaced by a system that generates training patterns before birth. We have previously shown that a system trained on such spontaneous activity can account for the measured face preferences of newborns (Bednar, 2002; Bednar & Miikkulainen, 2000). The 3 hypotheses of the present paper are that: (1) networks trained on spontaneous activity learn more robustly after birth, compared to systems exposed only to environmental stimuli, (2) the decline in re-

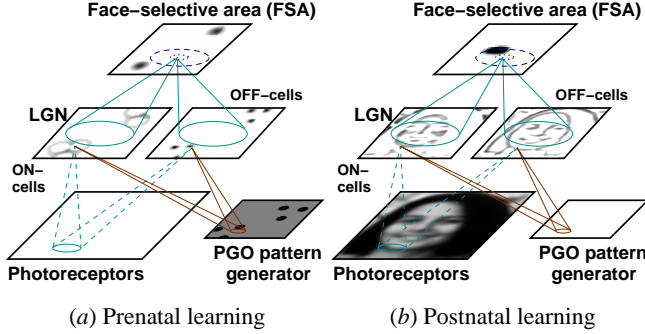


Figure 1: **HLISSOM model of face detection.** The model is a hierarchy of sheets of neural units, modeling the infant visual pathway. The model learns in two phases, each driven by input from a different source. During prenatal learning (a), activity originates internally in the PGO pattern generator. In postnatal learning (b), visual images are drawn on the photoreceptor sheet. For either type of input, the LGN is strongly activated by high-contrast edges and borders. The FSA is activated by the strongly responding units in the LGN, and activity then spreads laterally within the FSA. FSA weights are adapted when the activity settles. The FSA weights are initially uniform and unselective, but through self-organization they become selective for facelike patterns at the corresponding location on the retina. Their response is then information that an organism can use directly to control behaviors like visual fixation.

response to schematic patterns after one month results from learning of real faces and their outlines, not from a shift to a separate system, and (3) learning of both facial features and outlines can explain the development of a mother preference, including why it disappears when the outline is masked. Each of these hypotheses will be tested in a computational experiment with the HLISSOM model. Together the experiments will show that infant face learning and face preferences can be explained by a single, general-purpose learning system, which learns from both internally generated patterns of activity and from the visual environment.

## HLISSOM Model

The architecture for the HLISSOM model is shown in figure 1, and will be briefly reviewed below. (For more details, see Bednar, 2002.) The model consists of a hierarchy of two-dimensional sheets of neural units modeling different areas of the nervous system: two sheets of input units (the retinal photoreceptors and the ponto-geniculo-occipital (PGO) pattern generator, described under *Prenatal learning* below), two sheets of LGN units (ON-center and OFF-center), and a sheet of cortical units (“neurons”) representing a high-level area, the face-selective area (FSA)<sup>1</sup>. Each FSA neuron corresponds to a vertical column of cells through the six anatomical layers of the cortex.

<sup>1</sup>The FSA represents the first region in the ventral processing pathway that has receptive fields spanning approximately 45° of visual arc, i.e. large enough to span a human face at close range. Areas V4v and LO are likely FSA candidates based on adult patterns of connectivity, but the infant connectivity patterns are not known (Rolls, 1990). The generic term “face-selective area” is used rather than V4v or LO to emphasize that the model results do not depend on the region’s precise location or architecture, only on the fact that the region has receptive fields large enough to allow face-selective responses. Cortical areas between the LGN and the FSA have been bypassed for simplicity; see Bednar (2002) for a more complex model including the primary visual cortex (V1).

The input to the model is an activity pattern on a sheet of photoreceptors or the PGO generator (see examples in figure 1). Each cell  $(i, j)$  in the ON- and OFF-center layers of the LGN computes its response  $\eta_{ij}$  as a scalar product of a fixed weight vector and its receptive fields on each input sheet:

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

where  $\sigma$  is a piecewise linear sigmoid activation function,  $\rho$  specifies the input sheet (either photoreceptors or PGO),  $\gamma_{\rho}$  is a constant scaling factor,  $X_{\rho ab}$  is the activation of input unit  $(a, b)$  on sheet  $\rho$ , and  $w_{ij, \rho ab}$  is the corresponding weight value. The lower bound  $\delta$  of the sigmoid acts as an activation threshold; there is no response for activation below  $\delta$ . Each FSA neuron computes its initial response like that of an LGN cell, except that  $\rho$  is either the ON or OFF LGN layer. After the initial response, the FSA activity evolves through short-range excitatory and long-range inhibitory lateral interaction:

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

where  $\rho$  specifies the weight type (either ON channel afferent, OFF channel afferent, lateral excitatory, or lateral inhibitory),  $\gamma_{\rho}$  is a constant scaling factor for each weight type (negative for inhibitory lateral weights), and  $X_{\rho ab}(t-1)$  is the activation of target unit  $(a, b)$  during the previous time step. The FSA activity pattern starts out diffuse, but within a few iterations of equation 2, converges into a small number of stable focused patches of activity, or activity bubbles (as in figure 1). After the activity has settled, the connection weights of each FSA neuron are modified. All FSA weights adapt according to the Hebb rule, normalized so that the sum of the weights of each type  $\rho$  is constant for each neuron  $(i, j)$ :

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

where  $\eta_{ij}$  stands for the activity of neuron  $(i, j)$  in the final activity bubble,  $w_{ij, \rho ab}$  is the connection weight,  $\alpha$  is the learning rate for each type of connection, and  $X_{\rho ab}$  is the presynaptic activity. The larger the product of the pre- and post-synaptic activity  $\eta_{ij} X_{\rho ab}$ , the larger the weight change.

For these experiments, a pair of  $74 \times 74$  ON-center and OFF-center cell layers received input from a  $170 \times 170$  photoreceptor sheet and an  $85 \times 85$  PGO sheet. Each ON/OFF cell had a fixed Difference of Gaussians receptive field (RF) within the photoreceptor array (center  $\sigma = 0.75$ , surround  $\sigma = 1.2$ ). The  $24 \times 24$  FSA neurons each had a circular afferent receptive field of size 25, centered on the location in the central  $24 \times 24$  portion of the ON/OFF cell layer corresponding the neuron’s location in the FSA. This mapping ensures that every neuron has a complete set of circular afferent connections. Initially, the afferent and lateral weights in the FSA had a smooth circular Gaussian profile, and all weights of each type were identical. Other parameters were from Bednar and Miikkulainen (2000), scaled for this cortex size using the equations from Bednar (2002).

## Prenatal learning

The simulations in this paper focus on postnatal learning, but they continue from our earlier results on prenatal learning (Bednar & Miikkulainen, 2000), which we summarize

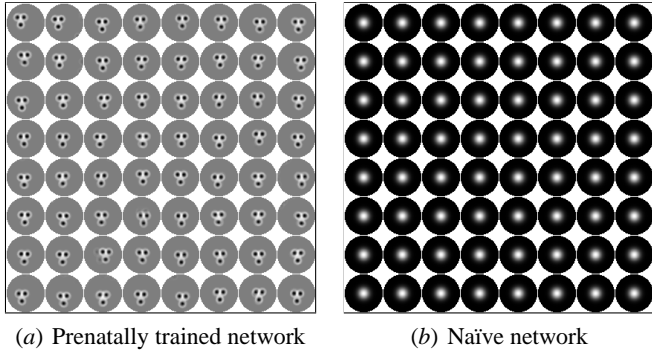


Figure 2: **Starting points for postnatal learning.** These plots show the RFs for every third neuron from the  $24 \times 24$  array of neurons in the FSA. For the prenatally trained network (a), the RFs were visualized by subtracting the OFF weights from the ON. The result is a plot of the retinal stimulus that would most excite that neuron. Like CONSPEC, the prenatally trained network consists of an array of roughly facelike RFs. In contrast, the neurons in the naïve network are initially uniformly Gaussian. The ON and OFF weights were identical, so only the ON weights are shown in (b). Later figures will compare the postnatal learning of each network.

here. We hypothesize that before birth, training patterns arise from ponto-geniculo-occipital (PGO) waves generated during rapid-eye-movement (REM) sleep. Developing embryos spend a large percentage of their time in a precursor of REM sleep, which suggests that this state has a major role in development (Roffwarg, Muzio, & Dement, 1966). During and just before REM sleep, PGO waves originate in the brain stem and travel to the LGN, visual cortex, and many other brain areas (see Callaway, Lydic, Baghdoyan, & 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, & Roffwarg, 1995). PGO waves elicit different distributions of activity in different species, and interrupting them has been shown to increase the influence of the environment on development (Marks et al., 1995).

All of these characteristics suggest that PGO waves may be providing species-specific training patterns for development (see Bednar, 2002 for more details). However, due to limitations in experimental imaging equipment and techniques, the spatial shape of the PGO wave activity patterns has not yet been measured. Based on the CONSPEC model, we chose the three-dot patterns illustrated in the PGO area of figure 1a. Other patterns are also possible, and will provide greater or lesser face selectivity (Bednar, 2002).

As described in previous work (Bednar & Miikkulainen, 2000), FSA neurons exposed to prenatal patterns developed receptive fields (RFs) preferring upright, triangular arrangements of three dots (figure 2a). The resulting map responds to most frontal face images, and not to most objects or backgrounds. At this stage, the trained map can be considered an implementation of CONSPEC, except that it was constructed by learning and will continue to learn after birth.

To determine whether the prenatal training biases subsequent learning (hypothesis 1 above), we also simulated a control condition called the *naïve* network. The naïve network

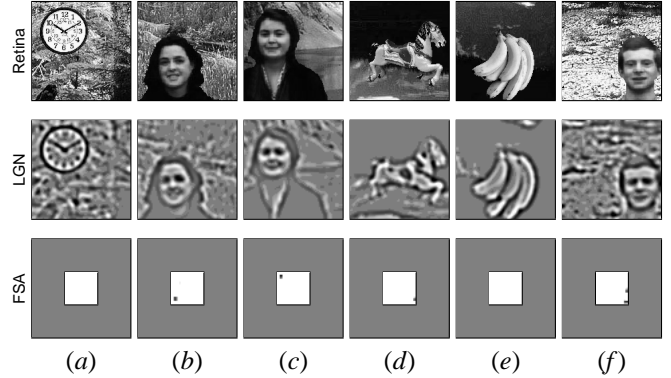
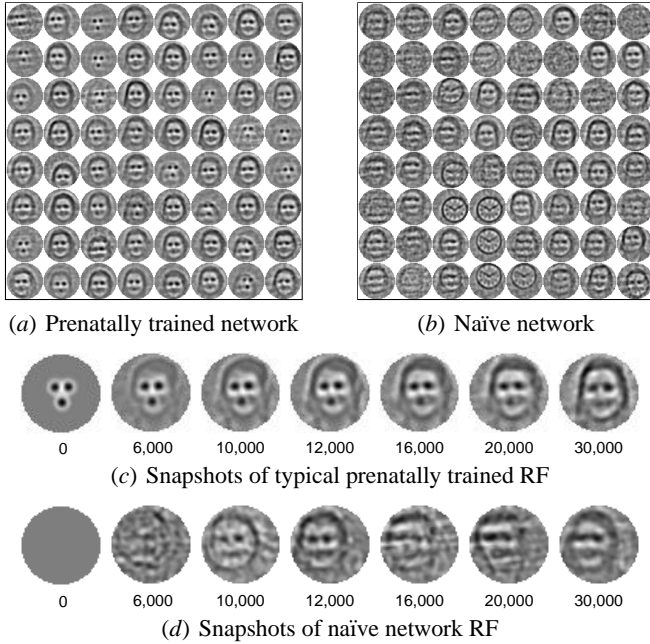


Figure 3: **Sample postnatal learning iterations.** The top row shows six randomly generated images drawn on the retinal photoreceptors at different iterations. Each image contains a foreground item chosen randomly from a set of three men and three women (adapted from Rowley et al., 1998) and six object images (from public domain clip art collections). Each foreground item was overlaid onto a random portion of an image from a database of 58 natural scenes (National Park Service, 1995), at a random location and at a nearly vertical orientation (drawn from a normal distribution around vertical, with  $\sigma = \pi/36$ ). The second row shows the LGN response to each of these sample patterns, visualized by subtracting the OFF cell responses from the ON cell responses. Dark areas indicate high OFF cell response, light indicate high ON cell response, and medium gray indicates no response. The bottom row shows the prenatally trained FSA response to each pattern, at the start of postnatal training. For the FSA, only neurons with complete receptive fields (those in the unshaded inner box) were simulated, because those in the gray area would have RFs cut off by the edge of the retina. The gray area shows the FSA area that corresponds to the same portion of the visual field as in the LGN and retina plots, to facilitate comparison. The FSA responds to groups of dark spots on the retina, such as the eyes and mouths in (b-c,f) and the horse's markings in (d); the location of the FSA activity corresponds to the position of the group of retinal patterns that caused the response.

is so called because it models neurons that have not had experience with coherent activity patterns until after birth. So that the naïve and prenatally organized networks would match on as many parameters as possible, we constructed the naïve network from the prenatally trained network *post hoc* by explicitly resetting afferent receptive fields to their uniform-Gaussian starting point (figure 2b). This procedure removed the prenatally developed face selectivity, but kept the lateral weights and all of the associated parameters the same. The activation threshold  $\delta$  for the naïve FSA network was then adjusted so that for a given training pattern both networks would have similar activation levels; otherwise the parameters were the same for each network. This procedure ensures that the comparison between the two networks will be as fair as possible, because besides the thresholds the networks differ only by whether the neurons have face-selective weights at birth.

## Postnatal testing and learning

The postnatal learning experiments reported in this paper simulate gradual learning of specific individuals and objects seen against a variety of different backgrounds. Figure 3 shows samples of the images we used and describes how they were generated. The prenatally trained and naïve networks were each exposed to the same pseudorandom sequence of 30,000 of these images.



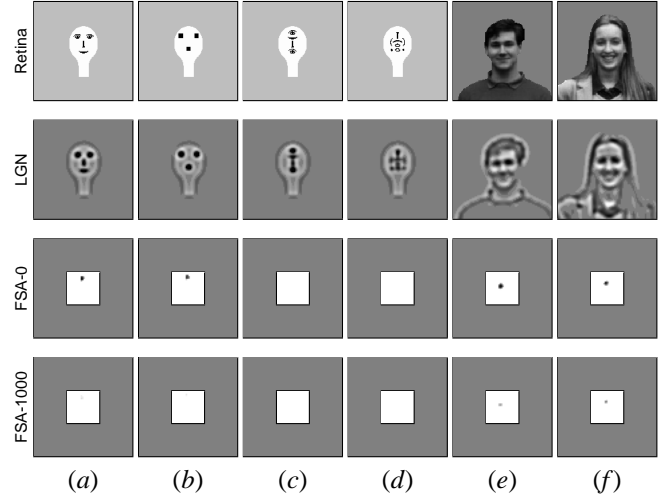
**Figure 4: Prenatal patterns bias postnatal learning in the FSA.** Plots (a) and (b) show the final RFs for every third neuron from the  $24 \times 24$  array of neurons in the FSA, visualized as in figure 2a. As the prenatally trained network learns from real images, the RFs morph smoothly into prototypes, i.e. representations of average facial features and hair outlines (c). By postnatal iteration 30,000, nearly all neurons have learned face-like RFs, with very little effect from the background patterns or non-face objects (a). Postnatal learning is less uniform for the naïve network, as can be seen in the RF snapshots in (d). In the end, many of the naïve neurons do learn face-like RFs, but others become selective for general texture patterns, and some become selective for objects like the clock (b). Overall, the prenatally trained network is biased towards learning faces, while the initially uniform network more faithfully represents the environment. Thus prenatal learning can allow the genome to guide development in a biologically relevant direction.

At the beginning of the postnatal phase, and at intervals throughout, we tested the network using schematic images previously tested with newborns, and with photographs of faces. In order to compare the neural activity in the model to babies’ attentional preferences, we assume that newborns pay more attention to the stimuli that are most effective at activating their visual processing system, focusing on the highest level activated. Patterns activating the FSA will be preferred over those activating only lower areas, and patterns that both activate the FSA will be ranked by their FSA activity. We quantify these comparisons by presenting each stimulus 25 times at different retinal locations, and averaging the sum of the FSA activity. As in the psychological studies we are modeling, differences between patterns will be tested with the two-tailed Student’s t-test, treating  $p$  values below 0.05 as significant.

## Results

### Experiment 1: Bias from prenatal learning

Figure 4 shows that with postnatal exposure to real images, both the naïve and prenatally trained networks develop RFs that are averages (i.e. prototypes) of faces and hair outlines. RFs in the prenatally trained network smoothly increase in



**Figure 5: Decline in response to schematic faces.** Before postnatal training, the prenatally trained FSA (third row from top) responds significantly more to the facelike stimulus (a) than to the three-dot stimulus (b;  $p = 0.05$ ) or the scrambled faces (c-d;  $p = 10^{-8}$ ). Assuming that infants attend most strongly to the stimuli that cause the greatest neural response, these responses replicate the schematic face preferences found by Johnson and Morton (1991) in infants up to one month of age. Some of the Johnson & Morton, 1991 experiments found no significant difference between (a) and (b), which is unsurprising given that they are only barely significantly different here. As the FSA neurons learn from real faces postnatally, they respond less and less to schematic faces. The bottom row shows the FSA response after 1000 postnatal iterations. The FSA now rarely responds to (a) and (b), and the average difference between them is no longer significant ( $p = 0.25$ ). Thus no preference would be expected for the facelike schematic after postnatal learning, which is what Johnson and Morton (1991) found for older infants, i.e. 6 weeks to 5 months old. The response to real faces also decreases slightly through learning, but to a much lesser extent (e-f). The response to real faces declines because the newly learned average face and hair outline RFs are a weaker match to any particular face than were the original three dot RFs. That is, the external features vary more between individuals than do the internal features, and thus their average is not a close match to any particular face. Even so, there is only a comparatively small decrease in response to real faces, because real faces are still more similar to each other than to the schematic faces. Thus HLISSOM predicts that older infants will still show a face preference if tested with more-realistic stimuli, such as photographs.

face selectivity, and eventually nearly all become highly selective for faces (figure 4b). Postnatal self-organization in the naïve network is less regular, and the final result is less face selective. Thus prenatal training biases postnatal learning towards biologically relevant stimuli, i.e. faces (hypothesis 1).

### Experiment 2: Decline in response to schematics

Figure 5 shows that the HLISSOM model replicates the disappearance of peripheral schematic face preferences after one month (hypothesis 2; Johnson et al., 1991). In HLISSOM, the decrease results from the afferent weight normalization (equation 3). 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. Unlike real faces, the facelike schematic patterns match only on these internal features, not the outlines. As a result, the response to schematic facelike patterns decreases as real faces

are learned. Eventually, the response to the schematic patterns approaches and drops below the fixed activation threshold  $\delta$ . At that point, the model response is no longer higher for schematic faces (because there is no FSA response, and V1 responses are similar). 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.

Yet the FSA neurons continue to respond to real faces (as opposed to schematics) throughout postnatal learning (figure 5e-f). Thus the model provides a clear prediction that the decline in peripheral face preferences is limited to schematics, and that if infants are tested with sufficiently realistic face stimuli, no decline in preferences will be found.

### Experiment 3: Mother preferences

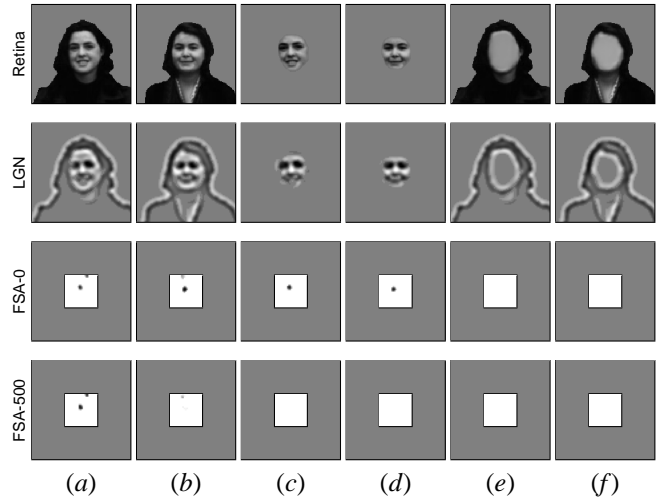
Figure 6a-b shows that when one face (i.e. the mother) appears most often, the FSA response to that face becomes stronger than to a similar stranger. This result replicates the mother preference found in infants a few days old (hypothesis 3; Bushnell, 2001; Pascalis et al., 1995). Interestingly, figure 6c-d shows that the mother preference disappears when the 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. Figure 6(e-f) shows that there is no response to the head and hair outline alone either, and thus that this face learning is clearly *not* outline-only.

In the model, the decreased response with either type of masking results from holistic learning of *all* of the features typically present in 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 also learn faces holistically (e.g. Farah et al., 1998). These results suggest that newborns may learn faces in the same way, and predict that newborns will no prefer their mother when her hair outline is visible but her facial features are masked.

### Discussion and future work

The HLISSOM simulations show that internally generated patterns and a self-organizing system 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 newborn learning appears to depend on the face outline, and why the response to schematic faces decreases over time. These explanations lead to concrete predictions for future infant experiments. 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. The results also show that internally generated patterns allow the genome to steer development towards biologically relevant processing, making learning of more sophisticated abilities quicker and more robust.

The results above do not address one interesting phe-



**Figure 6: 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). Thus face learning in HLISSOM is holistic. Face learning in adults is also thought to be holistic (Farah et al., 1998), and these results show that we do not need to assume that newborns are using a different type of face learning than adults.

nomenon: in central vision, preference for schematic faces is not measurable until 2 months of age (Maurer & Barrera, 1981), and is gone by 5 months (Johnson et al., 1991). This time course is delayed relative to peripheral vision, where preferences are present at birth but disappear by 2 months.

Johnson and Morton (1991) originally proposed that in the periphery the preferences disappear because CONLERN matures and inhibits CONSPEC, while in central vision they disappear because CONLERN learns properties of real faces. HLISSOM provides a unified explanation for both phenomena: a single learning system stops responding to schematic faces because it has learned from real faces.

Why, then, would the time course differ between peripheral and central vision? As Johnson and Morton acknowledged, the retina changes significantly over the first few months. In particular, at birth the fovea is much less mature than the periphery, and may not even be functional yet (Abramov, Gordon, Hendrickson, Hainline, Dobson, & LaBossiere, 1982; Kiorpes & Kiper, 1996). Thus schematic face preferences

in central vision may be delayed relative to those in peripheral vision simply because the fovea matures later. A single cortical learning system like HLISSOM is thus sufficient to account for the time course of both central and peripheral schematic face preferences.

The development of the fovea may also affect mother preferences. Consistent with our results, Bartrip, Morton, and de Schonen (2001) found that infants 19–25 days old do not significantly prefer their mothers when either her internal features or external features are covered. Interestingly, Bartrip et al. found that older infants, 35–40 days old, do prefer their mothers even when the external features are covered. The gradual maturation of the fovea may again explain these later-developing capabilities. Unlike the periphery, the fovea contains many ganglia with small RFs, and which connect to cortical cells with small RFs. These neurons can learn smaller regions of the mother's face, and their responses will allow the infant to recognize the mother even when other regions of the face are covered. Thus simple, documented changes in the retina can explain why mother preferences would differ over time.

In general, the idea that artificially generated training patterns can influence the development of learning systems is powerful, and could be used to construct artificial systems as well. Simple, engineered training patterns can provide an initial or ongoing bias, while learning algorithms incorporate the full complexity of the environment. This approach can allow more complex adaptive systems to be designed and implemented.

## Conclusion

A single learning system can explain the seemingly complex postnatal time course of face processing, if that system is exposed to internally generated patterns. Initial face selectivity develops from these non-visual inputs, and postnatal experience interacts with these genetic factors to develop full face processing abilities. These results provide clear predictions for future infant experiments, and provide new tools for constructing complex artificial systems.

## Acknowledgments

This research was supported in part by the National Science Foundation under grants IRI-9309273 and IIS-9811478.

## References

- Abramov, I., Gordon, J., Hendrickson, A., Hainline, L., Dobson, V., & LaBossier, E. (1982). The retina of the newborn human infant. *Science*, 217 (4556), 265–267.
- Bartrip, J., Morton, J., & de Schonen, S. (2001). Responses to mother's face in 3-week to 5-month-old infants. *British Journal of Developmental Psychology*, 19, 219–232.
- Bednar, J. A. (2002). *Learning to See: Genetic and Environmental Influences on Visual Development*. Doctoral Dissertation, Department of Computer Sciences, The University of Texas at Austin, Austin, TX.
- Bednar, J. A., & Mikkulainen, R. (2000). Self-organization of innate face preferences: Could genetics be expressed through learning?. In *Proceedings of the 17th National Conference on Artificial Intelligence* (pp. 117–122). Cambridge, MA: MIT Press.
- Bushnell, I. W. R. (2001). Mother's face recognition in newborn infants: Learning and memory. *Infant and Child Development*, 10 (1/2), 67–74.
- Callaway, C. W., Lydic, R., Baghdoyan, H. A., & Hobson, J. A. (1987). Pontogeniculooccipital waves: Spontaneous visual system activity during rapid eye movement sleep. *Cellular and Molecular Neurobiology*, 7 (2), 105–49.
- de Schonen, S., Mancini, J., & Leigeois, F. (1998). About functional cortical specialization: The development of face recognition. In (Simion & Butterworth, 1998), pp. 103–120.
- Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). What is "special" about face perception? *Psychological Review*, 105 (3), 482–498.
- Goren, C. C., Sarty, M., & Wu, P. Y. (1975). Visual following and pattern discrimination of face-like stimuli by newborn infants. *Pediatrics*, 56 (4), 544–549.
- Johnson, M. H., Dziurawiec, S., Ellis, H., & Morton, J. (1991). Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition*, 40, 1–19.
- Johnson, M. H., & Morton, J. (1991). *Biology and Cognitive Development: The Case of Face Recognition*. Oxford, UK; New York: Blackwell.
- Kiorpes, L., & Kiper, D. C. (1996). Development of contrast sensitivity across the visual field in macaque monkeys (*Macaca nemestrina*). *Vision Research*, 36 (2), 239–247.
- Marks, G. A., Shaffery, J. P., Oksenberg, A., Speciale, S. G., & Roffwarg, H. P. (1995). A functional role for REM sleep in brain maturation. *Behavioural Brain Research*, 69, 1–11.
- Maurer, D., & Barrera, M. (1981). Infants' perception of natural and distorted arrangements of a schematic face. *Child Development*, 52 (1), 196–202.
- National Park Service (1995). Image database. <http://www.freestockphotos.com/NPS>.
- Pascalis, O., de Schonen, S., Morton, J., Deruelle, C., & Fabre-Grenet, M. (1995). Mother's face recognition by neonates: A replication and an extension. *Infant Behavior and Development*, 18, 79–85.
- Roffwarg, H. P., Muzio, J. N., & Dement, W. C. (1966). Ontogenetic development of the human sleep-dream cycle. *Science*, 152, 604–619.
- Rolls, E. T. (1990). The representation of information in the temporal lobe visual cortical areas of macaques. In Eckmiller, R. (Ed.), *Advanced Neural Computers* (pp. 69–78). New York: Elsevier.
- Rowley, H. A., Baluja, S., & Kanade, T. (1998). Neural network-based face detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20 (1), 23–38.
- Simion, F., & Butterworth, G. (Eds.) (1998). *The Development of Sensory, Motor and Cognitive Capacities in Early Infancy: From Perception to Cognition*. East Sussex, UK: Psychology Press.
- Simion, F., Valenza, E., & Umiltà, C. (1998). Mechanisms underlying face preference at birth. In (Simion & Butterworth, 1998), pp. 87–102.
- Slater, A., Bremner, G., Johnson, S. P., Sherwood, P., Hayes, R., & Brown, E. (2000). Newborn infants' preference for attractive faces: The role of internal and external facial features. *Infancy*, 1 (2), 265–274.