Chapter to appear in *Face Processing in Infancy and Early Childhood*,
O. Pascalis and A. Slater (eds.), New York: Nova Science Publishers, 2003.

The Role of Internally Generated Neural Activity in Newborn and Infant Face Preferences

James A. Bednar
jbednar@cs.utexas.edu
Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712 USA

September 30, 2002

Abstract

Infant face perception is controversial, but the current evidence suggests that (a) newborns orient to and follow face-like schematic patterns more 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. Together, these capabilities have proved difficult to explain in terms of either environment-driven learning or genetically hardwired abilities. Accordingly, researchers have proposed that multiple visual processing areas may be involved, some hardwired and some plastic. New discoveries of widespread spontaneous neural activity during development suggest an alternative explanation: a single plastic visual processing system may learn from both spontaneous and visually evoked activity. In simulations with a biologically based computational model, we show that such internally generated patterns and a learning system can account for a wide range of these seemingly contradictory experimental results. The results suggest that learning of both internally generated and environmentally evoked activity may be a general feature of brain development.

INTRODUCTION

Specific regions in the adult visual cortex respond more strongly to human faces than to other similar stimuli (Kanwisher, McDermott, & Chun, 1997). 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. Thus, face perception appears to be substantially organized at birth, yet only develops fully through postnatal experience.

This paper summarizes work demonstrating 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. Further details are presented in Bednar (2002). Using the HLISSOM self-organizing model (Hierarchical Laterally Interconnected Synergetically Self-Organizing Map), we have shown how prenatal learning of internally generated activity can lead to newborn face preferences (Bednar & Miikkulainen, 2000, 2002a). We have also shown 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 (Bednar & Miikkulainen, 2002b). 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

Newborns appear to have some face detection abilities, but these abilities change significantly over early development. 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 2*b-e*). 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 & Morton (1991) proposed that infants are born with a basic subcortical system they termed CONSPEC (see also Johnson, this volume). 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 example, an infant only a few days old will prefer 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). 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 external features are masked (Pascalis et al., 1995).

Accordingly, Johnson & 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, i.e. a learning system similar to CONLERN, but for the periphery, and operational at birth.

Alternatively, CONLERN itself could be present 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 are unnecessary. A single cortical learning system present at birth is sufficient to explain the experimental data, if it is exposed to internally generated patterns broadly similar to faces. That is, CONLERN alone is sufficient, if instead of CONSPEC there is a mechanism that generates training patterns before birth.

INTERNALLY GENERATED PATTERNS

Spontaneous neural activity has been documented in many cortical and subcortical areas as they develop, including the visual cortex, the retina, the auditory system, and the spinal cord (reviewed in O'Donovan, 1999; Wong, 1999). This activity has been shown to be responsible for the segregation of the LGN (the lateral geniculate nucleus of the thalamus) into eye-specific layers before birth, indicating that internally generated activity patterns are crucial for normal visual development (Shatz, 1996; Stellwagen & Shatz, 2002).

In this work, we focus on one common type of spontaneous activity: ponto-geniculo-occipital (PGO) waves. These activity patterns are 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 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, & 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 also 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 (Jouvet, 1998). However, due to limitations in experimental imaging equipment and techniques, it has not yet been possible to measure the two-dimensional shape of the activity resulting from the PGO waves; only the temporal behavior is known (Rector, Poe, Redgrave, & Harper, 1997). How closely neonatal REM sleep and PGO waves are related to those of adults is also controversial (Dugovic & Turek, 2001; Jouvet, 1999). Even so, it is clear that there is substantial brainstem-generated activity in the visual areas of newborns during sleep.

We hypothesized that if these activity patterns have a simple spatial configuration of three active areas (or other similar patterns), they could account for the measured face-detection performance of human newborns and for subsequent postnatal learning. We tested this hypothesis in a series of computational experiments with the HLISSOM model. The experiments show that much of what is known about 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.

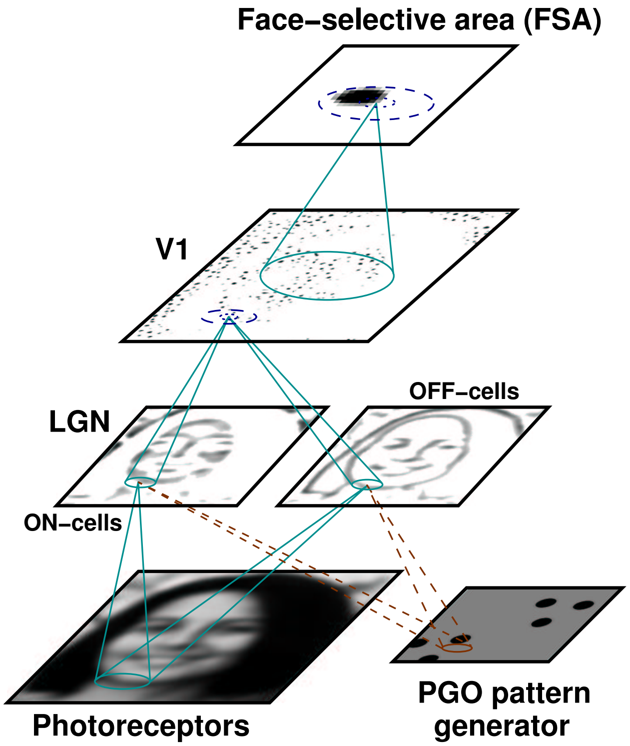

Figure 1: **HLISSOM model.** The model is a hierarchy of sheets of neural units, modeling the infant visual pathway. Each unit computes its response from the activity in its receptive fields (RFs); the RFs of one unit in each map are shown. Units have receptive fields on previous levels (afferent connections) and on the same level (lateral connections). Higher levels process more complicated visual properties, with the top-level area (the FSA) containing units that become selective for the location of a human face. In addition to the visual processing areas, HLISSOM includes input from internal sources of activity, represented here by the PGO pattern generator.

HLISSOM MODEL

The architecture for the HLISSOM model is shown in figure 1, and will be briefly reviewed below. (For more details, including parameter settings, 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 PGO pattern generator, two sheets of LGN units (ON-center and OFF-center), a sheet of primary visual cortex (V1) units, and another sheet of cortical units ("neurons") called the face-selective area (FSA). 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. Cortical areas between V1 and the FSA have been bypassed for simplicity, and postnatal simulations also omit V1 to focus on learning in the FSA. Each cortical neuron in V1 or the FSA corresponds to a vertical column of cells through the six anatomical layers of the cortex.

The input to the model is an activity pattern on a sheet of photoreceptors or the PGO generator (see examples in figure 1). The cells in the ON- and OFF-center layers of the LGN compute their responses as a scalar product of their retinal receptive fields and a weight vector. Like the ON and OFF cells, each FSA neuron also computes its initial response as a weighted sum of activity in its receptive field, subject to an activation threshold. After initial activation, the neural responses repeatedly propagate within the FSA through the excitatory and inhibitory lateral connections, and evolve into coherent activity "bubbles" (see example FSA activity in figure 1). After the activity stabilizes, weights of the active neurons are adapted by a normalized Hebbian rule. That is, weights are strengthened between units that are both active, while other connections to those neurons are weakened.

To determine the influence of internally generated activity, we simulated a variety of prenatal training conditions. Most simulations used the three-dot patterns shown in figure 1, presented at randomly chosen positions and random near-vertical orientations. For control purposes we also trained networks using other internally generated patterns, as well as testing an untrained (naïve) network. 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).

After prenatal training, and at intervals throughout postnatal training, we tested the network using schematic images previously tested with newborns, and with photographs of faces. In order to compare the resulting model neural activity to babies' attentional preferences, we assume 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 FSA activity.

RESULTS

Each of the following subsections describes the result of a computational experiment with the HLISSOM model.

Newborn schematic face preferences

When trained on three-dot patterns at various locations and orientations, FSA neurons developed receptive fields (RFs) preferring upright, triangular arrangements of three dots (figure 2*a*, row RF 10000). The network was then tested on 18 schematic patterns used in previous experimental studies. The total response of these neurons to each pattern ranked them in the same preference order for all 22 of the statistically significant preferences found in newborns (Bednar & Miikkulainen, 2002a). Figure 2 shows the response to four such patterns; the rest are shown in Bednar (2002). These results demonstrate that a network exposed to three-dot patterns is sufficient to explain the experimental results with schematic patterns. At this stage, the trained map can be considered a testable implementation of CONSPEC, Johnson & Morton's subcortical system, except that it was constructed by learning and will continue to learn after birth.

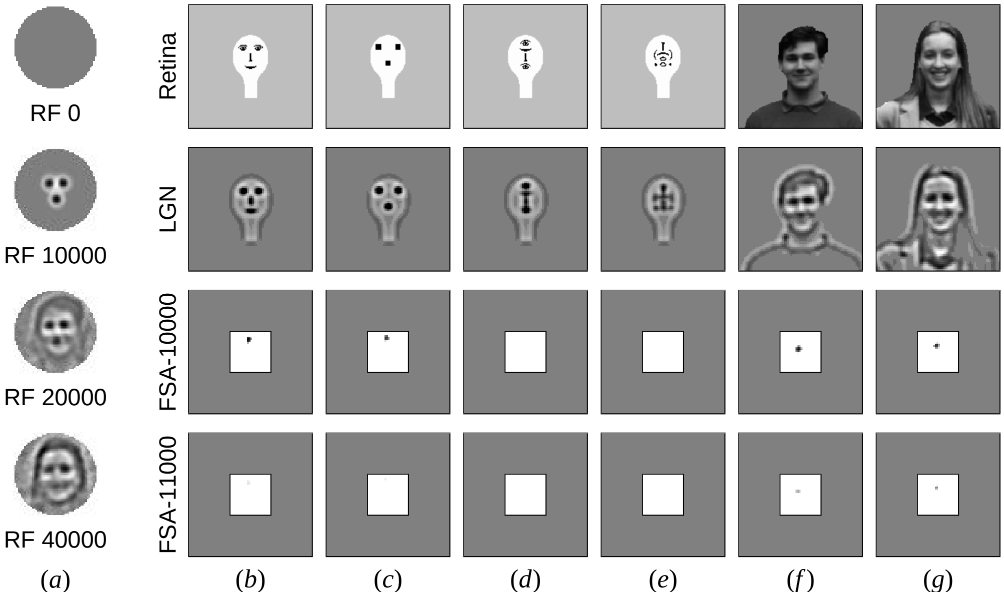

Figure 2: **Model RFs and responses.** The plots in (*a*) show the RF for a typical neuron in the FSA at different stages of training; the pattern shown is the LGN activity that would most excite this neuron. At first the neuron is unselective (RF 0), but after prenatal exposure to many three-dot patterns 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 & 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.

Newborn preferences for real face images

Although newborn studies usually use schematic patterns for practical reasons, it is 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 & Miikkulainen, 2002a). 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 does not respond to most other stimuli, except when they contain accidental three-dot patterns. The model predicts that human newborns will have a similar pattern of responses in the face-selective cortical regions.

Effect of prenatal training pattern shape

How crucial is the specific training pattern for the prenatal preferences? In addition to the three-dot pattern proposed by Johnson & Morton (1991), we tested a variety of other training pattern shapes (Bednar, 2002). For all training patterns, test faces matching the overall training pattern size gave higher responses than did other natural images. Thus simply matching the size of typical faces gives some degree of face selectivity in real images. However, different training patterns led to large differences in responses to the schematic patterns. Of all training patterns tested, the three-dot pattern provided the most overall face selectivity, and provided 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 shows no face preference. These results show that the pattern shape is important, but that the shape does not need to be strictly controlled to provide face preferences.

Postnatal bias from prenatal learning

What effect does the prenatal training have on postnatal learning? When given postnatal training of 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 & Miikkulainen, 2002b). However, the postnatal time course and final result differed. RFs in the prenatally trained network smoothly increase in face selectivity, and eventually nearly all become highly selective for faces. Postnatal self-organization in the naïve network is less regular, and the final RFs are usually less face selective. Overall, prenatal training appears to bias postnatal learning towards biologically relevant stimuli, i.e. faces. This result parallels one proposed effect of CONSPEC, but within a single learning system.

Postnatal decline in response to schematics

The bottom row of figure 2 shows that the HLISSOM model replicates the disappearance of peripheral schematic face preferences after one month (Johnson et al., 1991). In HLISSOM, the decrease results from the 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 on the internal features. Eventually, the response to the schematic patterns drops below the fixed activation threshold. At that point, the model response is no longer higher for schematic faces than for 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 2*f*-*g*). Thus the model provides a clear prediction that the decline in peripheral face preferences is limited to schematics, and that no decline will be found if infants are tested with sufficiently realistic face stimuli.

Mother preferences

In Bednar & Miikkulainen (2002b) we showed 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 (Bushnell, 2001; Pascalis et al., 1995). Interestingly, the model's 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. 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 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 learn faces holistically (e.g. Farah, Wilson, Drain, & Tanaka, 1998). The simulation results suggest that newborns may learn faces in the same way, and predict that newborns will not 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 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.

In future work, it will be important to measure the actual activity patterns present 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 model can be tested behaviorally with infant experiments. That is, 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. Finally, additional behavioral tests can help distinguish between HLISSOM and models predicting less face specificity, like Simion, Macchi Cassia, Turati, and Valenza (this volume).

CONCLUSION

A single learning system can explain the seemingly complex postnatal time course of face preferences, if that system is exposed to internally generated patterns. Initial face selectivity develops from these non-visual inputs, and postnatal experience interacts with the initial learning to develop full face processing abilities. These results provide clear predictions for future infant experiments, and provide 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. The simulations and papers referenced were done in collaboration with Risto Miikkulainen, and I thank Mark H. Johnson, Les Cohen, and others for comments on earlier paper drafts. Software, demonstrations, and related publications are available at `http://www.cs.utexas.edu/users/nn/`.

References

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., & Miikkulainen, 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.

Bednar, J. A., & Miikkulainen, R. (2002a). Learning innate face preferences. Submitted. Available as Technical Report AI-01-291, Department of Computer Sciences, University of Texas at Austin.

Bednar, J. A., & Miikkulainen, R. (2002b). Neonatal learning of faces: Interactions between genetic and environmental inputs. In *Proceedings of the 24th Annual Conference of the Cognitive Science Society* (pp. 107–112).

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, F., & Butterworth, G. (Eds.), *The Development of Sensory, Motor and Cognitive Capacities in Early Infancy: From Perception to Cognition* (pp. 103–120). East Sussex, UK: Psychology Press.

Dugovic, C., & Turek, F. W. (2001). Similar genetic mechanisms may underlie sleep-wake states in neonatal and adult rats. *Neuroreport*, *12* (14), 3085–3089.

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.

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*. Cambridge, MA: MIT Press.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, *17* (11), 4302–4311.

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.

O’Donovan, M. J. (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., & Fabre-Grenet, M. (1995). Mother’s face recognition by neonates: A replication and an extension. *Infant Behavior and Development*, *18*, 79–85.

Rector, D. M., Poe, G. R., Redgrave, P., & Harper, R. M. (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, H. P., Muzio, J. N., & Dement, W. C. (1966). Ontogenetic development of the human sleep-dream cycle. *Science*, *152*, 604–619.

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.

Shatz, C. J. (1996). Emergence of order in visual system development. *Proceedings of the National Academy of Sciences, USA*, *93*, 602–608.

Simion, F., Valenza, E., & Umiltà, C. (1998). Mechanisms underlying face preference at birth. In Simion, F., & Butterworth, G. (Eds.), *The Development of Sensory, Motor and Cognitive Capacities in Early Infancy: From Perception to Cognition* (pp. 87–102). East Sussex, UK: Psychology Press.

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.

Stellwagen, D., & Shatz, C. J. (2002). An instructive role for retinal waves in the development of retinogeniculate connectivity. *Neuron*, *33* (3), 357–367.

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