



Published in final edited form as:

Dev Psychobiol. 2012 March ; 54(2): 124–132. doi:10.1002/dev.20583.

The Development of the Uncanny Valley in Infants

David J. Lewkowicz¹ and Asif A. Ghazanfar²

¹ Department of Psychology & Center for Complex Systems & Brain Science, Florida Atlantic University, Boca Raton, FL 33431

² Neuroscience Institute, Departments of Psychology and Ecology & Evolutionary Biology, Princeton University, Princeton, NJ 08540

Abstract

When adults view very realistic humanoid robots or computer avatars they often exhibit an aversion to them. This phenomenon, known as the “uncanny valley”, is assumed to be evolutionary in origin, perhaps tapping into modules for disgust or attractiveness that detect violations of our normal expectations regarding social signals. Here, we test an alternative hypothesis that the uncanny valley is developmental in origin and, thus, that specific early experience with real human faces leads to its eventual emergence. To test this idea, we measured visual preferences in response to all possible pairs of a human face, realistic avatar face, and an unrealistic avatar face in groups of 6-, 8-, 10-, and 12-month-old infants. Consistent with the developmental hypothesis, we found that the uncanny valley effect emerges at 12 months of age suggesting that perceptual experience with real human faces is critical to its emergence.

Highly realistic human-looking robots or computer avatars tend to elicit negative feelings in humans (MacDorman, Green, Ho, & Koch, 2009; Mori, 1970; Seyama & Nagayama, 2007). Mori (1970) named this perceptual effect the “uncanny valley” (Figure 1). Although many human-looking robots and computer avatars bear a striking resemblance to the real human face, they are nearly always slightly imperfect and detection of these imperfections produces the uncanny feeling in observers. The effects of two kinds of imperfection – large eyes and skin texture - have been investigated in adults. These two types of imperfections induce an uncanny feeling and the two features interact with one another. For example, human-like avatar faces that have atypically large eyes tend to elicit feelings of unpleasantness (Seyama & Nagayama, 2007). Likewise, computer-generated faces with unrealistic skin textures tend to be seen as “more eerie” than computer-generated faces with photorealistic skin textures. When, however, faces with unrealistic skin textures are combined with disproportionately sized eyes, subjects find such faces to be less disturbing than faces with photorealistic skin textures and large eyes (MacDorman, et al., 2009).

It is assumed that the negative feelings that avatars elicit are caused by a conflict. On the one hand, avatars are assumed to elicit the concept of ‘human’. On the other, avatars are assumed to fail to live up to the concept of human because one or more character traits (e.g., eye size or skin texture) fall outside the spectrum of everyday social experience. Most hypotheses regarding the mechanisms and origins of the uncanny valley suggest that it is based on an evolved “module” for disgust or attractiveness (MacDorman, et al., 2009). Indeed, the evolutionary hypothesis is supported by comparative studies showing that adult macaque monkeys also exhibit the uncanny valley in that they look less at realistic monkey avatars than at real monkey faces or unrealistic monkey avatars (Steckenfinger & Ghazanfar, 2009). The fact that monkeys exhibit the uncanny valley effect suggests that whatever mechanism induces this effect, it is probably shared by the common ancestor of Old World monkeys (or at least, macaques) and humans.

Although the evolutionary hypothesis is reasonable, the mechanism need not be a specific module for this or that facial feature. It is more likely that developmental and evolutionary mechanisms, working hand-in-hand, have led to the emergence of this phenomenon. More specifically, the evolved mechanism could be a developmental manifold shared by both monkeys and humans (Steckenfinger & Ghazanfar, 2009). If that is the case, then it is possible that the proximal cause of the uncanny valley effect is the early and highly selective perceptual experience that species have with faces of their own species as opposed to the faces of other species. In humans, this experience consists of massive exposure to human faces from birth onwards, whereas in monkeys this experience consists of massive exposure to monkey faces. This highly selective exposure to conspecific faces presumably biases infant learning toward the conspecific face prototype. Once infants learn the prototype, they presumably acquire sufficient perceptual expertise to detect the slight anomalies inherent in realistic but imperfect avatar faces and begin to exhibit the uncanny valley effect. If our developmental hypothesis is correct then the study of human infant responses to human-like avatar faces provides a unique way to gain better insight into the effects of early experience on perceptual development as well as into the relationship between development and evolution.

The possibility that the uncanny valley of perception has developmental roots is reasonable on both theoretical and empirical grounds. On theoretical grounds, it is now well-accepted that early perceptual development depends greatly on experience, that this involves learning and differentiation, and that the end result of this process is the emergence of perceptual expertise (Gibson, 1969; Gottlieb, 1991; Nelson, 2001; Thelen & Smith, 1994; Werner, 1973). Thus, it would not be surprising if the uncanny valley of perception were a product of early developmental processes. On empirical grounds, a large body of evidence has now clearly documented that experience plays a central role in the development of perceptual skills. This body of research has shown that infants perceive faces, voices, and their combination from an early age, that these perceptual abilities are relatively crude at the start of life, and that they improve rapidly during the first year of life (Lewkowicz & Ghazanfar, 2009; Pascalis & Kelly, 2009; Scott, Pascalis, & Nelson, 2007; Simion, Leo, Turati, Valenza, & Dalla Barba, 2007; Werker & Tees, 2005). Moreover, this body of evidence has demonstrated that as perceptual abilities improve with development, infants become increasingly better at extracting more complex features from their perceptual array and, at the same time, that experience fine-tunes and narrows their responsiveness to just those stimulus features that are most frequent in their normal ecology. For example, starting at birth, infants exhibit preferences for face-like patterns and their specific configuration. These preferences are characterized by longer looking at visual patterns whose internal elements are arranged in a top-heavy configuration (i.e., two blobs at the top and a single blob below) than at patterns whose internal elements are inverted (Cassia, Turati, & Simion, 2004; Turati, Simion, Milani, & Umiltà, 2002; Valenza, Simion, Cassia, & Umiltà, 1996). Importantly, however, newborns prefer these kinds of top-heavy face-like patterns regardless of whether such patterns represent real human faces or static geometric versions of real faces. This indicates that newborns are not selective with respect to the particular features that contribute to face-like patterns and, thus, that they are broadly rather than specifically tuned to perceptual input.

The sort of broad tuning that newborn infants exhibit in response to face-like patterns has been found in many other studies, suggesting that it actually reflects a domain-general aspect of perceptual development. These other studies have shown that young infants are not only broadly tuned to faces (Pascalis, Haan, & Nelson, 2002) but to speech (Werker & Tees, 1984) and music (Hannon & Trehub, 2005) and that this broad tuning also characterizes young infants' response to face-voice relations (Lewkowicz & Ghazanfar, 2006; Pons, Lewkowicz, Soto-Faraco, & Sebastián-Gallés, 2009). In addition, these studies have shown

that the initial broad tuning narrows during the first year of life in a seemingly paradoxical way in that as infants acquire perceptual experience, and as they learn to detect increasingly finer perceptual structure, they stop responding to some of the perceptual features that they respond to earlier in life. For example, studies of infant response to native and non-native faces have found that whereas 6-month-old infants can discriminate different human and different monkey faces, 9-month-old infants can only discriminate human faces (Pascalis, et al., 2002). Similarly, studies of the other-race effect - where adults have difficulty discriminating the faces of other races (Chiroro & Valentine, 1995) - have found that young infants can discriminate the faces of different races but that older infants only discriminate the faces of their own race (Kelly, et al., 2005; Kelly, et al., 2007; Kelly, et al., 2009; Sangrigoli & de Schonen, 2004).

There is direct support for the conclusion that the various perceptual narrowing effects found in infancy are due to the specific early experience that infants acquire as they grow. The other-race effect, for example, can be reversed by exposing children to other-race faces at a time when they normally are not exposed to such faces (Sangrigoli, Pallier, Argenti, Ventureyra, & de Schonen, 2005). Similarly, the narrowing of responsiveness to the faces of other species can be prevented in infancy if infants are given extra exposure to non-native faces during the period when the narrowing process is known to begin (Pascalis, et al., 2005; Scott & Monesson, 2009). Finally, it has been found that narrowing of visual speech perception – watching the mouth utter silent speech - occurs in monolingual infants but not in bilingual infants, suggesting that continued exposure to both target languages maintains responsiveness to both of them (Weikum, et al., 2007).

Although perceptual narrowing reduces the range of stimulus features that infants respond to, it also leads to the emergence of perceptual expertise for the stimulus features that are part of the infant's native ecology. As a result, older infants become capable of detecting more complex and finer stimulus structure. This general developmental trend is illustrated by a variety of findings. For instance, 5-month-old infants do not discriminate audiovisual affect (i.e., sad from happy) but 7-month-olds do (Caron, Caron, & MacLean, 1988), 5-month-old infants do not exhibit matching of visual and audible affect but 7-month-old infants do (Walker-Andrews, 1986), and 4-month-old infants do not perceive audiovisual gender information but 8-month-olds do (Patterson & Werker, 2002).

It is now clear that the emergence of perceptual expertise is the result of the complimentary processes of perceptual narrowing, on the one hand, and improving perceptual detection skills on the other. If that is the case, and if the uncanny valley effect depends on the ability to detect relatively subtle imperfections, then it is highly likely that this effect is a developmental phenomenon. Moreover, given that most of the perceptual narrowing effects and the concurrent emergence of perceptual expertise becomes most evident in the latter half of the first year of life, it is also likely that the uncanny valley effect emerges during this period. If this hypothesis is correct, then infants should begin to detect the subtle differences that differentiate human faces from realistic but imperfect avatar faces and, as a result, should exhibit preferences consistent with the uncanny valley by the end of the first year of life.

To test our prediction, we investigated 6-, 8-, 10- and 12-month-old infants' visual preferences for human, realistic avatar, and uncanny avatar (realistic but imperfect avatar) faces. We did so by conducting three experiments across which we presented all possible pairs of these three types of faces (Figure 2). In Experiment 1, we paired the uncanny avatar face with the human face and, based on our developmental hypothesis, we expected that that the youngest infants would, at a minimum, not look less at the uncanny avatar face but that the oldest infants might look longer at the human face. Given that the most obvious

imperfection that distinguished the uncanny avatar from the human face was the disproportionate size of the eyes in the avatar face, in Experiment 2 we paired the uncanny avatar face with its large eyes with a realistic avatar face that had normally sized eyes. If infants were sensitive to eye size then we expected them to look for different amounts of time at these two avatars. Finally, if infants' response in Experiment 1 was based on eye size and not on how closely the faces resembled the human prototype, infants would not be expected to distinguish between a human face and an avatar face with normally sized eyes (Experiment 3).

General Method

Participants—We tested a total of 96 infants for the three experiments. This consisted of separate groups of 24 mostly Caucasian infants, each, at six months (mean age = 26 wks, SD = 0.8 wks; 12 boys), eight months (mean age = 34.3 wks, SD = 0.6 wks; 12 boys), 10 months (mean age = 43.2 wks, SD = 0.6 wks; 12 boys), and 12 months (mean age = 52.3 wks, SD = 1.3 wks; 14 boys). Each group of 24 infants at the four ages, respectively, consisted of separate cohorts of eight infants and each of these cohorts was tested in one of the three experiments (the exception to this was that eight of the 96 infants were tested in two different experiments across different ages). The data from an additional 10 infants were not included because of fussing (6), technical difficulties (1), or prior ear infection or low birth-weight (3). Infants were full-term and healthy and were recruited from local birth records.

Materials and procedure—The majority of the infants tested were seated in an infant seat at a distance of approximately 50 cm from the computer monitors on which they could see the stimuli. The few infants who were fussy were tested on the parent's lap and the parents were not aware of the hypothesis under test. Infants were given two 30 s paired-preference trials during which they viewed pairs of video clips on two 17 inch side-by-side computer monitors. The video clips showed a human face, realistic avatar face, or an uncanny avatar face uttering the syllable /ba/ silently every 2–3 s. All the stimuli were dynamically identical, meaning that neither the eyes nor the eyebrows moved, nor was there blinking or smiling. We chose dynamic faces because Mori (1970) predicted that movement accentuates the uncanny valley effect and because infants respond more to moving than static faces (Wilcox & Clayton, 1968).

Each cohort of eight infants/age was divided into two groups. Figure 2 depicts single frames of the specific face-pairs that we presented to each of the two groups across the three different experiments, respectively. As can be seen, two different human faces as well as two different avatar sets were presented. The avatar stimuli were produced using Poser Pro (Smith Micro Software, Inc., Aliso Viejo, CA) and exported as videos (30 frames/s). To produce the uncanny version of each of the avatar faces, we adopted the same approach as in the Seyama and Nagayama (2007) and the MacDorman et. al. (2009) studies and scaled up the eyes in size to 150%. Both Seyama and Nagayama (2007) and MacDorman et. al. (2009) found that this manipulation induced the uncanny valley effect in adult humans when combined with an otherwise realistic avatar face. The side on which the specific member of the pair was presented was counterbalanced across the two test trials. A closed-circuit camera, positioned between the two monitors, recorded visual fixations that were later coded by trained observers who were blind to the testing conditions. Inter-observer reliability, computed on a sample of randomly chosen subjects by computing the degree of agreement on the total duration of looking to each side was 99%.

EXPERIMENT 1

Adults find an uncanny looking face (i.e., with abnormally large eyes) more unpleasant to look at than a realistic face without abnormal features (MacDorman, et al., 2009; Seyama & Nagayama, 2007). If developmental experience plays no role in the uncanny valley effect and if it is present early in life then infants might exhibit similar preferences regardless of their age. Alternatively, if the uncanny valley effect emerges as a function of early experience then age-related differences would be expected. In particular, the age-related differences would be expected during the second half of the year when the various perceptual narrowing effects begin to emerge (Lewkowicz & Ghazanfar, 2009). To test these contrasting predictions in this experiment we investigated 6–12 month-old infants' looking at a human face paired with an uncanny avatar face.

Results & Discussion

To control for age-related differences in the absolute amount of looking, we converted the raw looking times to proportion-of-total-looking-time (PTLT) scores. To do so, we first computed the total amount of time each infant spent looking at each face type over the two test trials and then divided this amount by the total amount of time the infant looked at both face types. To determine whether responses were affected by the specific set of faces that were presented to the two groups of infants at each age, we conducted a preliminary analysis of the data. We entered the PTLT scores into a mixed analysis of variance (ANOVA), with Face Type as a within-subjects factor and Age and Group as between-subjects factors. Results of this preliminary analysis indicated that the Group factor did not have any effects and, as a result, we then collapsed the data across the Group factor and ran a new repeated-measures ANOVA on the PTLT scores, with Face Type as a within-subjects factor and Age as a between-subjects factor. Results of this analysis indicated that there was a significant Face Type x Age interaction, $F(3, 28) = 6.55, p = .0017; p_{\text{rep}} = .95, \eta_p^2 = .41$. This interaction is depicted in Fig. 3A where it can be seen that the proportion of looking at the two face types switched rather dramatically over the six month age-span. To determine whether the complimentary nature of the PTLT scores might have influenced the results of the statistical analysis, we also analyzed the absolute looking time scores from each infant for the two face types summed over the two test trials. This ANOVA, with Face Type as a within-subjects factor and Age as a between-subjects factor, yielded a nearly identical result for the Face Type x Age interaction, $F(3, 28) = 6.47, p = .0018$.

Based on results from adult studies, we predicted that infants would look longer at the human face and, thus, we used one-tailed Bonferroni-adjusted t-tests (a p value of less than .0125 was required for statistical significance) to determine whether looking at the two face types differed at each age. These tests (based on the PTLT scores) revealed that the 6-month-old infants exhibited marginally greater looking at the uncanny avatar face than the human face, $t(7) = 2.64, p = .0167$ and that the 12-month-old infants exhibited significantly greater looking at the human face than the uncanny avatar face, $t(7) = 3.38, p = .0059$.

The results from this experiment provided evidence of a rather dramatic shift in infant response to uncanny avatar and human faces. This evidence is particularly impressive given the small number of infants tested per age in this experiment. In addition, it is rather striking that the preference that we found at six months of age is opposite to what is typically observed in adult studies (MacDorman, et al., 2009; Seyama & Nagayama, 2007). That is, the 6-month-old infants preferred the uncanny avatar face over the human face whereas the 12-month-old infants preferred the human face over the uncanny avatar face (please note that we cannot conclude based on looking times alone that the older infants found the uncanny avatar “unpleasant”). These findings suggest that the uncanny valley effect is

absent in the first half of the first year of life and that it begins to emerge gradually during the second half of the first year.

EXPERIMENT 2

The age-related response differences to the human-uncanny avatar face pair found in Experiment 1 raise questions about the specific feature(s) that might have mediated responsiveness. The most obvious feature that renders the uncanny avatar abnormal is the uncharacteristic size of its eyes. As indicated earlier, unusually large eyes induce an uncanny feeling in human adults (MacDorman, et al., 2009; Seyama & Nagayama, 2007). As a result, to determine whether this was the critical feature that caused infants to differentiate between the two faces in Experiment 1, here we paired the uncanny avatar face with a realistic avatar face. As Fig. 2 shows, these two faces were identical except for the size of their eyes. If eye size was the feature that was responsible for the infants' differential response in Experiment 1 then the infants in this experiment should differentiate between the realistic avatar face and the uncanny avatar face.

Results & Discussion

As in Experiment 1, we first conducted a preliminary repeated-measures ANOVA on the PTLT scores to determine whether group membership affected outcome. The results of this analysis indicated that it did not and, as a result, we collapsed the data across the Group factor and reanalyzed the PTLT scores by way of a repeated-measures ANOVA, with Face Type as a within-subjects factor and Age as a between-subjects factor. Consistent with our prediction, and as can be seen in Fig. 3B, infants looked longer at the realistic avatar face than at the uncanny avatar face across the 6–12 month age range. This difference was evident in a significant main effect of Face Type, $F(1, 28) = 13.19, p = .0011; p_{\text{rep}} = .939, \eta_p^2 = .32$. The ANOVA of the absolute looking time scores also yielded a significant main effect of Face Type, $F(1, 28) = 10.89, p = .0026$.

The finding that infants preferred the realistic avatar face shows that they were highly sensitive to eye size differences and that they preferred the face with eyes that are normal in size relative to the human face prototype. In addition, the findings from this experiment indicate unambiguously that the decline in looking at the uncanny avatar face found in Experiment 1 was due to the disproportionately large eyes in the uncanny avatar rather than to the synthetic nature of the computer-generated avatar face. Finally, the current findings are consistent with the results from adult studies showing that adults are highly sensitive to abnormal facial features and that they find them to be unpleasant (MacDorman, et al., 2009; Seyama & Nagayama, 2007).

EXPERIMENT 3

The conclusion that it is not the synthetic nature of the face *per se* and its associated slight imperfections that caused the 12-month-old infants to look less at the uncanny avatar face in Experiment 1 is a reasonable one. It should be noted, however, that this conclusion was not tested directly in Experiment 1. As a result, the purpose of Experiment 3 was to test this possibility directly by pairing a human face with a realistic avatar face. We expected that, despite the fact that these two faces differed on a number of dimensions, infants would look equally at them because the eyes were proportionately correct in each and because this feature dominated responsiveness.

Results & Discussion

The preliminary ANOVA indicated that the grouping factor did not affect outcome and, as a result, we collapsed over it and conducted a repeated-measures ANOVA, with Face Type as

a within-subjects factor and Age as a between-subjects factor. The results of this analysis indicated that there was no significant difference in looking at the two faces, $F(1, 28) = 1.49$, n.s. (see Fig. 3C). The ANOVA of the absolute looking time scores also indicated that there was no significant main effect of Face Type, $F(1, 28) = 2.53$, n.s. This result confirms the conclusion that infants were attending primarily to the eyes and that they did not detect the synthetic nature of the realistic avatar. This, in turn, further supports the main finding that infants looked less at the uncanny avatar because of its unusually large eyes.

General Discussion

The results of this study are consistent with the hypothesis that the proximal cause of the uncanny valley is developmental experience. The developmental hypothesis is supported by the pattern of visual preferences obtained in Experiment 1. For the youngest infants, the pattern of visual preferences is opposite to the findings reported from studies of adult humans' and adult monkeys' response to uncanny faces. In studies of adult humans, it has been found that they dislike looking at realistic but uncanny avatars (MacDorman, et al., 2009; Seyama & Nagayama, 2007). Similarly, in studies of monkeys, it has been found that they look less at realistic but uncanny synthetic monkey faces than at unrealistic synthetic monkey faces and real monkey faces (Steckenfinger & Ghazanfar, 2009). In contrast, here we found that 6-month-old infants exhibited a marginal preference for the uncanny avatar over the human face and that it was not until 12 months of age that infants began to exhibit a significant preference for the human face over the uncanny avatar face. This unexpected developmental shift in preferences, and the relatively late emergence of a pattern of visual preferences that is consistent with the uncanny valley effect, is a testament to the powerful effects of early perceptual experience and its central role in shaping perceptual expertise and the concomitant emergence of a preference for the human face prototype.

Our findings raise interesting questions about the processes underlying the developmental changes observed here and the ultimate emergence of the uncanny valley. Two complimentary processes most likely underlie the developmental changes. As noted earlier, one is perceptual learning and differentiation of increasingly finer stimulus features (Gibson, 1969; Gottlieb, 1991; Nelson, 2001; Thelen & Smith, 1994; Werner, 1973) and the other is perceptual narrowing (Lewkowicz & Ghazanfar, 2009; Scott, et al., 2007; Werker & Tees, 2005). Together, these two processes contribute to the development of the perceptual expertise that is required for the perception of subtle facial imperfections. The fact that infants exhibited a dramatic increase in their looking at the human face relative to their looking at the uncanny avatar face between six and 12 months of age in Experiment 1 (the linear trend in looking at the human face over age was highly significant, $F(1, 28) = 16.41$, $p < .001$) suggests that everyday experience with human as opposed to other-kind faces, and the association of human faces with generally positive consequences, gradually confers special status on them. As this occurs, and as infants continue to learn about human faces, they gradually become experts at perceiving increasingly finer facial features. This expertise is further enhanced by perceptual narrowing which enables infants to deal with a more restricted range of stimulus attributes that can now be explored in a more detailed manner. Together, the processes of perceptual learning/differentiation and narrowing provide the foundation for the emergence of the uncanny valley that has been found in studies of adult humans and monkeys (MacDorman, et al., 2009; Seyama & Nagayama, 2007; Steckenfinger & Ghazanfar, 2009).

What makes our findings particularly intriguing is that they clearly demonstrate that the level of perceptual expertise achieved between six and 12 months of life is sufficient to distinguish between different types of faces, including differences between synthetic ones, but that this does not include the ability to detect the relatively subtle features that distinguish between human faces and synthetic agents. This limitation, particularly at the

end of the first year of life, is interesting because infants of this age have already become sufficiently specialized for human faces that they no longer discriminate the faces of other-species (Pascalis, et al., 2002) and of other races (Bar-Haim, Ziv, Lamy, & Hodes, 2006; Sangrigoli, et al., 2005). Despite this, however, the expertise achieved by 12 months is sufficient to allow infants to detect uncanny facial features.

In sum, the pattern of findings from the three experiments in the current study is consistent with the conclusion that the uncanny valley of perception is the result of early and specific developmental experience. The product of such experience is a useful behavioral adaptation because it enables observers to quickly detect anomalies (e.g., disease) and/or the esthetic value (i.e., beauty) of a face. As a result, it is likely that the uncanny valley was conserved in evolution via developmental manifolds (Gottlieb, 2002) that recreate the developmental conditions that render the kinds of preferences reflected in the uncanny valley adaptive across successive generations. If so, the current findings, together with previous findings from monkeys (Steckenfinger & Ghazanfar, 2009), demonstrate how behavioral characteristics that turn out to be adaptive during ontogeny may be reflected in a species' phylogenetic history (Gottlieb, 1992; Lewkowicz, 2011; Lickliter & Harshaw, 2010).

Acknowledgments

This work was supported, in part, by NSF grant BCS-0751888 and NIH grant HD057116 to DJL and by an NSF CAREER award BCS-0547760 to AAG. We thank Shawn Steckenfinger for creating the avatar stimuli and Kelly McMullen and Amy Hansen for their help in testing the infants.

References

- Bar-Haim Y, Ziv T, Lamy D, Hodes RM. Nature and Nurture in Own-Race Face Processing. *Psychological Science*. 2006; 17(2):159–163. [PubMed: 16466424]
- Caron AJ, Caron RF, MacLean DJ. Infant discrimination of naturalistic emotional expressions: The role of face and voice. *Child Development*. 1988; 59(3):604–616. [PubMed: 3383670]
- Cassia VM, Turati C, Simion F. Can a nonspecific bias toward top-heavy patterns explain newborns' face preference? *Psychological Science*. 2004; 15(6):379–383. [PubMed: 15147490]
- Chiroro P, Valentine T. An investigation of the contact hypothesis of the own-race bias in face recognition. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*. 1995; 48A(4):879–894.
- Gibson, EJ. *Principles of perceptual learning and development*. New York: Appleton; 1969.
- Gottlieb G. Experiential canalization of behavioral development: Theory. *Developmental Psychology*. 1991; 27(1):4–13.
- Gottlieb, G. *Individual development & evolution: the genesis of novel behavior*. New York: Oxford University Press; 1992.
- Gottlieb, G. Emergence of the developmental manifold concept from an epigenetic analysis of instinctive behavior. In: Lewkowicz, DJ.; Lickliter, R., editors. *Conceptions of development: Lessons from the laboratory*. New York: Psychology Press; 2002. p. 31-56.
- Hannon EE, Trehub SE. Metrical categories in infancy and adulthood. *Psychological Science*. 2005; 16(1):48–55. [PubMed: 15660851]
- Kelly DJ, Quinn PC, Slater AM, Lee K, Gibson A, Smith M, Ge L, Pascalis O. Three-month-olds, but not newborns, prefer own-race faces. *Developmental Science*. 2005; 8(6):F31–F36. [PubMed: 16246233]
- Kelly DJ, Quinn PC, Slater A, Lee K, Ge L, Pascalis O. The other-race effect develops during infancy: Evidence of perceptual narrowing. *Psychological Science*. 2007; 18(12):1084–1089. [PubMed: 18031416]
- Kelly DJ, Liu S, Lee K, Quinn PC, Pascalis O, Slater AM, Ge L. Development of the other-race effect during infancy: evidence toward universality? *Journal of Experimental Child Psychology*. 2009; 104:105–114. [PubMed: 19269649]

- Lewkowicz DJ, Ghazanfar AA. The decline of cross-species intersensory perception in human infants. *Proceedings of the National Academy Sciences USA*. 2006; 103(17):6771–6774.
- Lewkowicz DJ, Ghazanfar AA. The emergence of multisensory systems through perceptual narrowing. *Trends in Cognitive Sciences*. 2009; 13(11):470–478. [PubMed: 19748305]
- Lewkowicz DJ. The biological implausibility of the nature-nurture dichotomy and what it means for the study of infancy. *Infancy*. 2011; 16(4):331–367. [PubMed: 21709807]
- Lickliter, R.; Harshaw, C. Canalization and Malleability Reconsidered: The Developmental Basis of Phenotypic Stability and Variability. In: Hood, K.; Halpern, C.; Greenberg, G.; Lerner, R., editors. *Handbook of Developmental Science, Behavior, and Genetics*. New York: Wiley-Blackwell; 2010.
- MacDorman K, Green R, Ho C, Koch C. Too real for comfort? Uncanny responses to computer generated faces. *Computers in Human Behavior*. 2009; 25(3):695–710.
- Mori M. The uncanny valley. *Energy*. 1970; 7(4):33–35.
- Nelson CA. The development and neural bases of face recognition. *Infant & Child Development*. 2001; 10(1–2):3–18.
- Pascalis O, Haan Md, Nelson CA. Is face processing species-specific during the first year of life? *Science*. 2002; 296(5571):1321–1323. [PubMed: 12016317]
- Pascalis O, Scott LS, Kelly DJ, Shannon RW, Nicholson E, Coleman M, Nelson CA. Plasticity of face processing in infancy. *Proceedings of the National Academy of Sciences USA*. 2005; 102(14):5297–5300.
- Pascalis O, Kelly DJ. The origins of face processing in humans: Phylogeny and ontogeny. *Perspectives on Psychological Science*. 2009; 4(2):200–209.
- Patterson ML, Werker JF. Infants' ability to match dynamic phonetic and gender information in the face and voice. *Journal of Experimental Child Psychology*. 2002; 81(1):93–115. [PubMed: 11741376]
- Pons F, Lewkowicz DJ, Soto-Faraco S, Sebastián-Gallés N. Narrowing of intersensory speech perception in infancy. *Proceedings of the National Academy of Sciences USA*. 2009; 106(26):10598–10602.
- Sangrigoli S, de Schonen S. Recognition of own-race and other-race faces by three-month-old infants. *Journal of Child Psychology and Psychiatry*. 2004; 45(7):1219–1227. [PubMed: 15335342]
- Sangrigoli S, Pallier C, Argenti AM, Ventureyra VAG, de Schonen S. Reversibility of the other-race effect in face recognition during childhood. *Psychological Science*. 2005; 16(6):440–444. [PubMed: 15943669]
- Scott LS, Pascalis O, Nelson CA. A domain general theory of the development of perceptual discrimination. *Current Directions in Psychological Science*. 2007; 16(4):197–201. [PubMed: 21132090]
- Scott LS, Monesson A. The origin of biases in face perception. *Psychological Science*. 2009; 20(6):676–680. [PubMed: 19422630]
- Seyama J, Nagayama RS. The uncanny valley: Effect of realism on the impression of artificial human faces. *Presence: Teleoperators and Virtual Environments*. 2007; 16(4):337–351.
- Simion, F.; Leo, I.; Turati, C.; Valenza, E.; Dalla Barba, B. How face specialization emerges in the first months of life. In: von Hofsten, C.; Rosander, K., editors. *Progress in Brain Research*. Vol. 164. Elsevier; 2007.
- Steckenfinger S, Ghazanfar A. Monkey visual behavior falls into the uncanny valley. *Proceedings of the National Academy of Sciences*. 2009; 106(43):18362.
- Thelen, E.; Smith, LB. *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT Press; 1994.
- Turati C, Simion F, Milani I, Umiltà C. Newborns' preference for faces: What is crucial? *Developmental Psychology*. 2002; 38(6):875–882. [PubMed: 12428700]
- Valenza E, Simion F, Cassia VM, Umiltà C. Face preference at birth. *J Exp Psychol Hum Percept Perform*. 1996; 22(4):892–903. [PubMed: 8756957]
- Walker-Andrews AS. Intermodal perception of expressive behaviors: Relation of eye and voice? *Developmental Psychology*. 1986; 22:373–377.

- Weikum WM, Vouloumanos A, Navarra J, Soto-Faraco S, Sebastián-Gallés N, Werker JF. Visual language discrimination in infancy. *Science*. 2007; 316(5828):1159. [PubMed: 17525331]
- Werker JF, Tees RC. Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior & Development*. 1984; 7(1):49–63.
- Werker JF, Tees RC. Speech perception as a window for understanding plasticity and commitment in language systems of the brain. *Developmental Psychobiology Special Issue: Critical Periods Re-examined: Evidence from Human Sensory Development*. 2005; 46(3):233–234.
- Werner, H. *Comparative psychology of mental development*. New York: International Universities Press; 1973.
- Wilcox BM, Clayton FL. Infant Visual Fixation on Motion Pictures of the Human Face. *Journal of Experimental Child Psychology*. 1968; 6(1):22–32. [PubMed: 5642131]




Figure 1. The hypothetical plot of the relationship between degree of realism and response in adults as proposed by Mori (1970) in his hypothetical uncanny valley of perception.

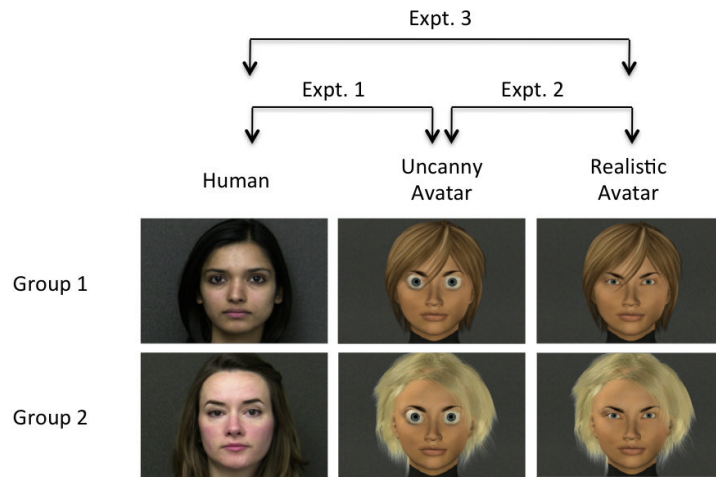


Figure 2. Still images of the three types of faces presented and the specific face-pairings presented in each experiment, respectively, to each group of infants.




Figure 3. Proportion of total looking time (PTLT) accorded to each of the two faces presented in each experiment as a function of age. Error bars represent the standard errors of the mean. (A) Experiment 1, (B) Experiment 2, (C) Experiment 3.