

Creepiness Creeps In: Uncanny Valley Feelings Are Acquired in Childhood

Kimberly A. Brink
University of Michigan

Kurt Gray
University of North Carolina, Chapel Hill

Henry M. Wellman
University of Michigan

The uncanny valley posits that very human-like robots are unsettling, a phenomenon amply demonstrated in adults but unexplored in children. Two hundred forty 3- to 18-year-olds viewed one of two robots (machine-like or very human-like) and rated their feelings toward (e.g., “Does the robot make you feel weird or happy?”) and perceptions of the robot’s capacities (e.g., “Does the robot think for itself?”). Like adults, children older than 9 judged the human-like robot as creepier than the machine-like robot—but younger children did not. Children’s perceptions of robots’ mental capacities predicted uncanny feelings: children judge robots to be creepy depending on whether they have human-like minds. The uncanny valley is therefore acquired over development and relates to changing conceptions about robot minds.

All day, every day, both children and adults try to get inside the minds of others, wondering about their thoughts, feelings, and intentions. Until the past few decades, these minds have been those of flesh and blood—humans and animals—but now we are faced with minds made of metal and silicon, including smart phones and cloud computing. How do we learn to make sense of these artificial minds?

Nowhere is this question more pressing than with robots, who have self-directed mechanical minds dwelling inside human-like bodies. The National Robotics Initiative foresees a future in which “robots are as commonplace as today’s automobiles, computers, and cell phones. Robots will be found in homes, offices, hospitals, factories, farms, and mines; and in the air, on land, under water, and in space” (National Robotics Initiative 2.0: Ubiquitous Collaborative Robots (NRI-2.0; 2017). In fact, robots are already entering homes, not only to help adults with household chores, but also to play with and teach children. Moreover, several robots have been working with children in classrooms, day cares, clinics, and hospitals for years. Robots

are teaching children language skills (Movellan, Eckhardt, Virnes, & Rodriguez, 2009), mathematics (Wei, Hung, Lee, & Chen, 2011), science (Hashimoto, Kobayashi, Polishuk, & Verner, 2013), physical exercises (Mejias et al., 2013), and even social skills (Ricks & Colton, 2010). Dozens of robots have been released in the past year alone designed specifically to interact with children. As robots become increasingly present in our lives and the lives of our children, it becomes more and more important to explore how we reason about the minds of these devices and how this reasoning impacts our interactions and feelings toward them.

Work with adults has identified one phenomenon, in particular, that could shed light on this topic. Specifically, decades of research reveal that while adults prefer robots that are somewhat human like, they find very human-like robots unnerving—the “uncanny valley” phenomenon (MacDorman, Green, Ho, & Koch, 2009; Mori, MacDorman, & Kageki, 2012). According to theories of the uncanny valley, machines become increasingly attractive as they become more human like until they reach a threshold at which they become too human like and are considered uncanny and creepy (see Figure 1). This *dip* in affinity for very human-like robots is the uncanny valley. Closely

The research was funded by a grant from the Varieties of Understanding Project at Fordham University and the John Templeton Foundation to Henry M. Wellman, and a NICHD traineeship to Kimberly A. Brink. The research was conducted in the Living Laboratory[®] at the University of Michigan Museum of Natural History.

Correspondence concerning this article should be addressed to Kimberly A. Brink, Department of Psychology, University of Michigan, 530 Church St, Ann Arbor, MI 48109. Electronic mail may be sent to kabrink@umich.edu.

© 2017 The Authors
Child Development © 2017 Society for Research in Child Development, Inc.
All rights reserved. 0009-3920/2019/9004-0014
DOI: 10.1111/cdev.12999

human-like robots are distinctly creepier than other robots and, in particular, creepier than the more unsettling of machine-like robots. Support for the uncanny valley comes from many studies in which adults report feeling greater unease when presented with extremely human-like robots compared to others (Gray & Wegner, 2012; MacDorman, 2006).

Two theories have been proposed to explain the uncanny valley's origins. One references innate evolutionary concerns (Steckenfinger & Ghazanfar, 2009), including the innate drive to avoid illness. Human-like robots may display "visual defects" that are interpreted as signs of a "communicable disease" thus producing a creepy response. Alternatively, innate face-processing mechanisms may recognize visual defects in very human-like faces compared to real human faces and thus process those faces as unattractive and creepy. Various facial processing mechanisms and standards are in fact apparent even in infants (see Langlois, Roggman, & Rieser-Danner, 1990 for evidence that infants prefer attractive faces). This evolutionary account receives further support from research demonstrating that even monkeys experience an uncanny valley when viewing computer-simulated images of monkey faces (Steckenfinger & Ghazanfar, 2009).

An alternative theory proposes that for humans the uncanny valley is not simply a by-product of evolutionary perceptual responses, but instead

depends on an acquired everyday understanding of what makes humans distinct from machines (MacDorman & Ishiguro, 2006). Feelings of uncanniness may instead emerge when a human-like machine violates our learned expectations of how a machine should look or behave. In the case of robots, for example, when a machine closely resembles a thinking and feeling human being, this would violate our expectations that machines should be inanimate and hence incapable of thought and experience. Specifically, a very human-like appearance in a machine can prompt attributions of a human-like mind (Epley, Waytz, & Cacioppo, 2007), and as human-like minds are seldom ascribed to robots (Gray, Gray, & Wegner, 2007), this mismatch causes feelings of uncanniness (Gray & Wegner, 2012). Indeed, research with adults reveals that the more robots are seen to have human feelings, the more unnerving they seem (Gray & Wegner, 2012). Violations of expectations about the behavior and appearance of machines and humans thus link to the uncanny valley phenomenon in adults.

While the uncanny valley has been studied in adults, its origins have never been studied in children. If evolutionary in nature, the uncanny valley should be evident in even the youngest children. However, if it is related to developing expectations about humans and machines, then it should emerge throughout childhood—perhaps in tandem with exclusive attributions of human-like minds to humans. As the origins and mechanisms of the uncanny valley have yet to be tested in children, we examine them here.

We offer a detailed look at the uncanny valley across development by measuring uncanny responses to videos of robots in children from ages 3 to 18. We used stimuli previously validated with adults (Gray & Wegner, 2012): videos of the same robot that revealed either its machine-like or human-like nature (see Figure 2). We showed these videos to children and then assessed their feelings of creepiness and also their attributions of mind—thinking (agency) and feeling (experience)—toward the robots. By assessing feelings of unease and mind attribution across a large age range, we could detect whether (and when) the uncanny valley develops, and its potential link to children's understandings of robot minds.

We expected one of three possible patterns would likely appear: (a) the uncanny valley would be present at even the youngest ages, offering support for the evolutionary perspective. (b) The uncanny valley would emerge in early childhood in tandem with general perceptions of mind—offering




Figure 1. A schematic depiction of the theoretical uncanny valley (figure closely derived from figure 2 in Mori et al., 2012). The uncanny valley is defined as the precipitous dip in affinity for closely human-like robots. [Color figure can be viewed at wileyonlinelibrary.com]



Figure 2. Still frames from the videos of each robot: Kaspar from the back (left panel), Kaspar from the front (center panel), and Nao (right panel). [Color figure can be viewed at wileyonlinelibrary.com]

support for the developmental perspective. (c) The uncanny valley would emerge in later childhood, when children develop more sophisticated—and specific—understandings of the different kinds of minds possessed by machines, also supporting a developmental perspective.

The evolutionary perspective suggests that even the youngest children should find human-like robots more unnerving than machine-like robots, *irrespective* of attributions of mind. On the other hand, the developmental perspective suggests that the uncanny valley will emerge in childhood, perhaps even early childhood, when children have begun distinguishing humans (and human minds) from other categories. Theory-of-mind research shows that 3- to 5-year-olds become quickly adept at attributing mental states such as beliefs and desires to humans (see meta-analyses by Milligan, Astington, & Dack, 2007; Wellman, Cross, & Watson, 2001). This too might predict that the uncanny valley would be evident in our youngest age group. However, this possibility could, nonetheless, be distinguished from the evolutionary one if the presence of uncanny valley responses in early childhood is related to children's attributions of mind.

Alternatively, the uncanny valley could instead arise around middle childhood—when children develop richer understandings of folk biology and folk psychology, and begin to separate the concepts of minds, brains, bodies, and machines (Wellman, 2014). For example, it is only at about 9–12 years that children truly understand differences between the mind (as more “mental”) and the brain (as more part of the biological body; Johnson & Wellman, 1982; Richert & Harris, 2006). This understanding that the mind stems from the biological brain (i.e., a neurophysiological “machine”) could support the development of the uncanny valley: The uncanny valley may result from the mismatch of perceiving a human mind as stemming from a

machine brain. Indeed, post preschool children, as they age, expect machines to have fewer mental abilities (Kahn et al., 2012, which examined children ages 9–15). It may not be until this later age that children develop an understanding that robots, as machines, should not have minds as humans do, making them uncanny when they seem like they do.

Method

Participants

Two hundred forty children (117 females), 3–18 years old, were recruited from a local natural history museum (218 children) or from a participant database (12 children) between March 2015 and June 2016. Children were questioned in a semi-isolated, quiet space within the museum or (for 12) in an on-campus laboratory space. One child was excluded due to incorrect parental report of their birthdate. Our sample was twice the number of participants ($N = 120$) used in a similar previous task (Gray & Wegner, 2012, study 1). Power analyses indicate that $N = 240$ exceeds .80 statistical power (Cohen, 1988).

Because data were collected in a public space, we did not collect information regarding children's race, ethnicity, or socioeconomic status. Written parental consent and verbal child assent were obtained first; children received a small toy for participating.

Videos

Children were randomly assigned to watch short videos of either a closely human-like robot or a more machine-like robot (Figure 2), the two used in Gray and Wegner (2012). For the human-like robot, 119 children watched 16s of Kaspar moving its

head, filmed from the front with its human-like face clearly visible. In the machine-like robot condition, 120 children watched 16s of the robot Kaspar moving its head while filmed from behind, where *only* its wiring and electrical components could be seen, no human-like features were visible. Respondents could not infer that these views were of the same robot because (paralleling Gray & Wegner, 2012) these videos were presented between subjects. This focal comparison controls for many irrelevant differences between the human-like and machine-like robots because the videos contain the same robot making the same movement but from different views (front vs. back).

Although our focal contrast compares the human-like and machine-like versions of Kaspar, we also included a video of Nao—a commercially available abstractly humanoid robot—as a baseline condition (see Figure 2). We filmed Nao to mimic the human-like robot—only head, face, and torso visible, moving its head from side to side, with no changes in expression. After viewing either the machine-like robot or human-like robot, 234 children watched the 16s video of Nao. We implemented this baseline condition using Nao as one indicator that children used our rating scale and terms appropriately (see following).

Nao has been used in previous studies with children (4–9 years) to effectively comfort them during stressful events (e.g., receiving a vaccination; Beran, Ramirez-Serrano, Vanderkooi, & Kuhn, 2013) and so is presumably not creepy or uncanny. Nao is also unlikely to be considered creepy because it resembles the friendly, animated robot protagonists portrayed in children’s films like *WALL-E* (Stanton, 2008) and Baymax in *Big Hero 6* (Hall & Williams, 2014). Thus, if children appropriately use our scale, they should provide low ratings of uncanniness for Nao. Although we expect that Nao should be rated low on feelings of uncanniness, we did not use Nao as an indicator of the presence or absence of the uncanny valley. First, Nao has never been empirically placed on the hypothesized uncanny valley gradient. Moreover, a contrast between Nao and the very human-like robot would be insufficient evidence to prove the existence of the uncanny valley. The uncanny valley is more specifically defined as the dip in affinity when very human-like robots are perceived as creepier than even somewhat creepy machine-like robots. Finding that a very human-like robot (Kaspar from the front) is creepier than a not very creepy, and even comforting, humanoid robot (Nao) would not sufficiently demonstrate the presence or absence of the uncanny valley.

Task and Design

After viewing each robot (machine-like or very human-like, then Nao), children answered multiple questions presented in a two-part format. In the first part, children chose one of two options. For example, when asked “Does the robot think for itself?,” children answered either yes or no verbally or by pointing to a “thumbs up” (yes) or a “thumbs down” (no) card (see Figure 3). If children answered yes, they then answered a second Likert-type scale question. For example, “How much does the robot think for itself?”: “a little bit,” “a medium amount,” or “a lot.” Children could answer verbally or by indicating on a scale with increasingly tall bars (see Figure 3).

Robot Beliefs Interview

Children were assessed on their feelings of uncanniness via two questions gauging the extent to which children felt the robot was creepy or unsettling: (a) “Do you feel the robot is nice or creepy?” If children reported the robot was creepy, we asked, “How creepy do you feel it is?” (b) “Does the robot make you feel weird or happy?” and “How weird does it make you feel?” This question format resulted in a 4-point scale for each question coded as 0 (*nice/happy*), 1 (*creepy/weird—a little bit*), 2 (*creepy/weird—a medium amount*), and 3 (*creepy/weird—a lot*).

Then children answered 11 additional questions, 10 of them addressing the robots’ mental capacities (adapted from Gray et al., 2007; Severson & Lemm, 2016; see Appendix for complete interview). Previous interviews with adults have identified components of mental capacity labeled “agency” and “experience.” Questions were designed to encompass similar factors in our sample: psychological agency (does the robot “do things on purpose?,” “choose to move?,” “think for itself?,” and “know the difference between good and bad?”) and perceptual experience (would the robot “feel pain?,” “feel scared?,” and “feel hungry?”). The same two-part question format resulted in a 4-point scale for each of these “mind” questions coded as 0 (*no*), 1 (*yes—a little bit*), 2 (*yes—a medium amount*), and 3 (*yes—a lot*).

Procedure

Children were instructed that they would view videos of robots and answer questions about them. Children then answered three warm-up questions




Figure 3. Images that were shown to children to aid them in answering the two-part survey questions: thumbs up (yes), thumbs down (no), and a scale with bars increasing in height (“a little bit,” “a medium amount,” or “a lot”). These exact depictions were taken from Severson and Lemm (2016).

and were randomly assigned to watch a video of either the closely human-like or machine-like robot on an iPad. After watching the video, children completed the *Robot Beliefs Interview*. The video was paused so that a still frame of the robot was visible during the interview. Upon completion of the interview, children watched a video of Nao and performed the *Robot Beliefs Interview* once more for Nao, while a still frame of Nao remained visible on the iPad.

Data Analysis

Exploratory factor analysis

It was unclear whether children’s perceptions and feelings toward robots would reveal the same factor structure as past work with adults, so we performed an exploratory factor analysis separately for each of the two conditions (very human like and machine like). Using an oblique rotation, Kaiser’s criterion (eigenvalues ≥ 1), a scree plot, and model fit indices (Preacher & MacCallum, 2003), three factors were identified: uncanniness (machine like: $\alpha = .62$, human like: $\alpha = .75$), agency ($\alpha = .72$, $.64$), and experience ($\alpha = .73$, $.85$). Five additional items were pruned (see Appendix) that had low factor loadings, cross-loaded on to multiple factors, or were not easily interpretable; these items were not included in the final factor analysis.

As shown in Table 1, the three factors (covering eight items)—uncanniness, agency, and experience—had identical factor structures across the two

conditions (very human like and machine like) and provided high overall fit. Table 1 reveals that all items had a loading of at least .40 on their respective factors. Model fit indices also support the three-factor solution within each condition. For the machine-like condition, chi-square goodness of fit, $\chi^2(7) = 3.64$, $p = .82$, root mean square error of approximation (RMSEA) = 0.00, 90% CI [0.00, 0.07], and Tucker Lewis index (TLI) = 1.07, were all within their established cutoff ranges (Hu & Bentler, 1999). For the human-like condition, model fit indices approximated or were within established cutoff ranges: chi-square goodness of fit, $\chi^2(7) = 10.53$, $p = .16$, RMSEA = 0.07, 90% CI [0.00, 0.14], and TLI = 0.953. Cronbach’s α ranged from .62 to .85 across all three factors and both conditions.

Attributions of Mind Score

Agency was measured by averaging the items “does the robot choose to move?,” “think for itself?,” and “know the difference between good and bad?” Thought, decision making, and morality have been linked to psychological agency for adults (Gray et al., 2007). Experience was measured by averaging the items: “would the robot feel pain?,” “feel scared?,” and “feel hungry?.” These items also resembled those for perceptual experience in adult research (Gray et al., 2007). The aggregates for agency and experience were highly correlated, $r(236) = .49$. Thus, for conceptual reasons and to avoid issues of multicollinearity, agency and experience were averaged to create a composite measure

Table 1
Exploratory Factor Analysis of Interview Items

Items	Factors		
	Uncanniness	Agency	Experience
Do you feel the robot is creepy?	.81, .99	-.05, .02	.01, -.07
Does the robot make you feel weird?	.57, .55	.13, -.18	-.01, .12
When the robot moves, does it choose to move?	.08, .10	.59, .69	.06, .07
Does the robot think for itself?	.01, -.04	.77, .53	-.04, .34
Does the robot know the difference between good and bad?	-.14, -.10	.66, .47	.02, -.08
Would the robot feel pain?	-.03, .02	-.02, -.04	.93, .86
Would the robot feel scared?	.01, .03	.27, .03	.40, .82
Would the robot feel hungry?	.03, -.09	.27, .06	.51, .76
α	.62, .75	.72, .64	.73, .85

Note. Using an oblique rotation, Kaiser's criterion (eigenvalues ≥ 1), a scree plot, and model fit indices, three factors were identified. The first nine rows represent the factor loadings for each item and in the bottom row, the α values for each factor. Bolded numbers identify the items that were used to calculate the aggregates for their respective factors. The first number in each cell represents values for the machine-like robot and the second for the human-like robot, respectively (as indicated by the two numbers in each column).

of attributions of mind. This approach has also been used in the adult literature (see Gray, Knickman, & Wegner, 2011).

Uncanniness Difference Scores

To account for individual differences in children's use of the scale, the dependent variable of *uncanniness* was converted to a difference score. This difference score was calculated by subtracting the uncanniness score for the baseline condition, Nao, from the uncanniness score for the focal robot that a child viewed. For example, for children that viewed the human-like robot, their uncanniness score for Nao was subtracted from their uncanniness score for the human-like robot. This difference score allows the comparison between our focal two robots to be individualized to the extent that we use Nao as a baseline, anchoring the score of the focal robots on an empirically verified comforting

robot, Nao (Beran et al., 2013). This score thus also provides a control for children's unfamiliarity with robots in general which can vary from child to child and age to age (see Dalecki & Willits, 1991 for an explanation and justification of the statistical advantages of such comparisons). We did not use Nao as a regressor in our subsequent regression analyses because preliminary analyses showed that feelings of uncanniness for Nao did not significantly predict feelings of uncanniness for the focal robots. Uncanniness difference scores will be referred to as uncanniness scores.

Results

The Development of the Uncanny Valley

Attributions of mind, robot type (very human like vs. machine like), and age as well as interactions between mind and age, and robot and age were entered into a regression analysis predicting uncanniness scores. As shown in Table 2, there were significant associations between uncanniness scores and attributions of mind and robot type, qualified by the interaction between attributions of mind and age and the interaction between robot type and age. There was no main effect of age.

The interaction between robot type and age indicates that the uncanny valley develops. The positive interaction indicates that as children age, the human-like robot is perceived as increasingly creepier than the machine-like robot. A plot of the interaction can be seen in Figure 4. Tests of simple slopes indicated that the human-like robot does not become creepier than the machine-like robot until

Table 2
Regression Analyses Predicting Uncanniness Difference Scores

	Estimate	SE	β	t	p
Human-like robot	.32	.15	.14	2.10	.037*
Mind	-.25	.11	-.21	-2.31	.022*
Age	-.23	.14	-.20	-1.68	.094
Human-Like Robot \times Age	.32	.15	.13	2.08	.039*
Mind \times Age	.20	.10	.17	2.04	.043*
Intercept	.89	.12		7.41	< .0001****

Note. Robot type, attributions of mind, the child's age, and two interaction terms were entered into a regression analysis to predict uncanniness difference scores. Attributions of mind and age were centered. Robot type, attributions of mind, the interaction between robot type and age, and the interaction between attributions of mind and age all predicted reports of uncanniness, $R^2 = .12$.

* $p < .05$. **** $p < .0001$.

approximately 9 years of age. Before 9 years of age, robot type did not predict feelings of uncanniness (i.e., both the human-like and machine-like robots were equally creepy): at age 4, $\beta = -.08$, $p = .57$ and at age 8, $\beta = .13$, $p = .11$. At 9 years, however, the uncanny valley effect emerges with robot type significantly predicting uncanniness, $\beta = .18$, $p = .02$, where the human-like robot is significantly more uncanny than the machine-like robot. The uncanniness of human-like robots relative to machine-like robots continues to increase up to 16 years, $\beta = .55$, $p = .002$.

Uncanniness and Mind

The development of the uncanny valley is also linked to attributions of mind, shown by the significant interaction between attributions of mind and age on uncanniness scores (as shown in Table 2). This interaction indicates that, as children get older, the association between attributions of mind and feelings of uncanniness changes. Specifically, tests of simple slopes show that, in young children, increased attributions of mind tend to predict *decreased* feelings of uncanniness in young children: For children ages 4–9, attributions of mind negatively correlated with the uncanny response—at

4 years, $\beta = -.55$, $p < .001$, and at 9 years, $\beta = -.23$, $p = .04$. For older children (ages 10–18), this negative correlation started to disappear and began to trend positive, although not significantly so—at 10 years, $\beta = -.16$, $p = .21$, and at 16 years, $\beta = .23$, $p = .42$. The upper age limit of our sample, however, had fewer participants and therefore less power to test for a statistically positive association. This descriptively positive trend at the oldest range of our child sample fits qualitatively with findings in the adult literature in which feelings of uncanniness become positively associated with attributions of mind (Gray & Wegner, 2012). In total, this significant interaction suggests that the emergence of the uncanny valley is associated with children's perceptions of mind, particularly for younger children.

Uncanniness Responses for Nao

Figure 5 shows the raw scores for feelings of uncanniness for all three robots. Inspection of this figure shows that Nao was consistently rated low on uncanniness across all ages in our sample. A regression analysis predicting raw uncanniness scores while controlling for age and comparing Nao with the least creepy of the two focal robots, the machine-like robot, showed that Nao is less creepy than the machine-like robot, $\beta = -.82$, $t(110) = -9.34$, $p < .0001$. Nao is also less creepy than the closely human-like robot, $\beta = -1.05$, $t(111) = -11.58$, $p < .0001$. Statistically, appropriate baseline conditions should have low, stable scores on the variable of interest, thus dispelling the possibility of a yes bias. Our expectation for Nao to be minimally uncanny (as explained earlier) was thus confirmed empirically, a result that supports its use as a baseline condition for creating uncanniness difference scores. Furthermore, raw uncanniness scores for Nao did not differ between children who first saw the human-like ($M = 1.31$, $SD = 0.67$) or the machine-like robot ($M = 1.36$, $SD = 0.69$), $t(224.25) = -0.53$, $p = .59$, 95% CI $[-0.23, 0.13]$. Because the robot that children saw first did not impact their responses to Nao, this provides additional empirical justification for its use as a baseline condition.




Figure 4. The interaction between robot type and age, $\beta = .13$, $t(223) = 2.08$, $p = .04$, shows that the uncanny valley develops. The positive interaction indicates that as children get older the human-like robot becomes increasingly creepier than the machine-like robot, as shown here. The development of the uncanny valley effect is demonstrated by the increasing distance between the two lines (machine like vs. human like) with age.

Appropriate Understanding of “Creepy” and “Weird”

Methodologically, the comparison between Nao and the other robots shows that even the youngest children respond to our uncanny valley questions and test format with varied answers across the conditions, indicating that they offered meaningfully




Figure 5. Measure of uncanniness, the aggregate of the two questions that measured whether children perceived the robot to be creepy or weird, for each of the three robots: the machine-like robot (Kaspar from the back), the human-like robot (Kaspar from the front), and Nao. Error bars represent standard errors.

differential responses; children did not merely demonstrate a yes bias to our scale. Although their uncanny ratings, based on use of the terms “creepy” and “weird,” do not distinguish the machine-like and the human-like robots, they do distinguish between Nao and these two.

For further confirmation of their understanding of these two key terms, we tested a separate sample of 20 young children on their understanding of the terms “weird” and “creepy” via a brief interview. These children came from the same local population as those in our main study and were equivalent in age to children at the lower end of the age range in our sample ($M = 3.40$ years, $SD = 0.34$, range = 3.01–3.99). These children were presented with two paired images of a typical toy (i.e., a stuffed giraffe or tricycle) and a clearly strange toy (toys modeled after the creepy/weird toys in toy story; Lasseter, 1995) shown in Figure 6. For one pair, children were asked to select which toy was creepy and, for the other pair, which toy made them feel weird (see questions in Figure 6). On these tasks, 95% of these young children appropriately chose the strange toy as creepy and 85% chose the other strange toy as making them feel weird. Overall children were 90% correct on these items that used terms and phrasing closely similar to those in the two items that constituted our uncanny

index. Thus, even children at our youngest age are capable of appropriately using the two words necessary for meaningfully employing our uncanny rating scale.

Discussion

We provide three novel findings. First, the uncanny valley develops: Younger children found the closely human-like and machine-like robots equally not very creepy, whereas older children found the closely human-like robot much creepier than the machine-like robot—similar to adults. Second, we identified the approximate age at which the uncanny valley emerges. Differences in feelings about the two focal robots emerged progressively over age, but it was not until middle childhood that children had a greater uncanny response to a closely human-like robot than a contrasting machine-like robot. Third, children’s perceptions of mind were correlated with this change in uncanny responses. For younger children, increasing perceptions of mind predicted *decreased* uncanniness. For older children, this association trended in the reverse direction, though not significantly so. Thus, from younger children to older children, the correlation between mind and feelings of uncanniness



Figure 6. Stimuli used for follow-up interview to address whether children appropriately understood the words “creepy” and “weird.” [Color figure can be viewed at wileyonlinelibrary.com]

increases from negative to trending positive, and the broader literature indicates that this correlation continues to increase and becomes positive for adults (Gray & Wegner, 2012). Of course, caution is needed in such cross-study child–adult comparisons (even though our videos for the Kaspars are exactly those used by Gray & Wegner, 2012). Regardless, we clearly demonstrate that feelings of uncanniness emerge and change over childhood, and are associated with differing perceptions of robots’ minds.

This research addresses important questions in psychology and robotics by providing initial evidence on the origins of the uncanny valley. One theory suggests that the uncanny valley is grounded in an innate mechanism (Steckenfinger & Ghazanfar, 2009), which means it should be present at an early age. Indeed, a priori, it is easy to imagine that young children could have responded to human-like robots as adults do—perceiving them as creepy. However, our results suggest that the uncanny valley emerges through development, and tracks changing understandings of mind. Our results clearly show that only children 9 years and older—those who have clear expectations about human and robot minds—feel unease toward very human-like robots, consistent with the second

hypothesis outlined in the introduction that a robot is considered creepy when it violates our learned expectations of how a machine should look or behave.

The absence of the uncanny valley in younger children may reflect that they *expect* robots to have a myriad of mental abilities. In fact, other research on children’s understanding of robots supports this speculation. Young children report that robots have perceptual abilities, like sight and touch: 3-year-olds claimed that a robot dog could see and be tickled (Jipson & Gelman, 2007). In conversations with parents, 3- to 5-year-olds also attributed biological, psychological, and sensory abilities, in addition to features of artifacts, to a robot dog (Jipson, Gülgöz, & Gelman, 2016). As young children seemingly expect some robots to have mental abilities, the perception of a human-like mind may be a welcome familiarity for them. Indeed, our results are consistent with this explanation: Young children found robots to be more pleasing (less uncanny) when they perceived the robots to have more mental abilities.

Older children, on the other hand, seemingly have different expectations about robot minds and expect robots to have reduced mental capacities.

Five-year-olds are less likely to claim that a robot dog could think or feel happy compared to 3-year-olds (Jipson & Gelman, 2007) and are less likely to report that a robot has emotions, desires, or is capable of autonomous action (Mikropoulos, Misailidi, & Bonoti, 2003). Fifteen-year-olds were also less likely to believe that the robot Robovie could have interests, experience emotions, or be a friend compared to 9- and 12-year-olds (Kahn et al., 2012). These changes in expectations with age have been linked to children's increasing experiences with and growing knowledge of technological devices (Bernstein & Crowley, 2008). Although the ages and robots varied across these studies, the general trend is clear: with age, children begin to deny psychological, emotional, social, and perceptual abilities to robots. For our older children, judgments of mind were no longer negatively associated with their ratings of uncanniness: attributions of mind stopped predicting a decrease in feelings of uncanniness. By hypothesis, older children in our sample may have evidenced an uncanny response to the human-like robot due to emerging changes in their expectations about the mental abilities of robots.

Our research was exploratory in the sense that, in advance of collecting the needed, relevant data, we had no firm prediction for which of various developmental patterns might emerge. Although exploratory, our results clearly demonstrate that the uncanny valley effect develops with age, and that it occurs in tandem with judgments of mind.

Questions about the uncanny valley and children will become only more important over time as more robots are being made to interact and play with children. iPal, Jibo, and Zenbo, three Pixar-like robots, are designed to play games, answer questions, read stories, and watch children unsupervised (Glaser, 2016; Low, 2016; Wong, 2016). Nao, Ursus (a robotic bear), and Kaspar are all robots used to teach typically developing children, and those with motor disorders and autism spectrum disorders, a variety of skills including language (Movellan et al., 2009), physical exercises (Mejías et al., 2013), and social skills (Ricks & Colton, 2010). We should likely ensure that children, both typically developing and those with special needs, actually like these robots before extensively using them as companions or teachers.

While our research speaks to important avenues for future research, there are noteworthy potential limitations to our study. First, we acknowledge that our data are only correlational—unlike some work with adults (Gray & Wegner, 2012), there is no causal evidence for the link between understandings of

mind and the uncanny valley in childhood. Future studies should therefore more explicitly test developments in thinking about minds (and machines) more generally with the development of the uncanny valley. This could involve both experiments and longitudinal studies that track children's developing concepts of minds and machines (Gelman, 2003).

Second, for older children, judgments of uncanniness became dissociated with attributions of mind; for them, that link was no longer statistically significant. Although at first glance this result may seem problematic, a clear developmental picture emerges when these childhood data are coupled with data from adults in prior studies: The correlation between mind and uncanniness increases over age from young (negative, i.e., increasing attributions of mind predict decreased feelings of uncanniness) to middle (zero) to older (trend positive) children and to adults (positive). On the other hand, it is also still possible that significant positive links between uncanniness and mind (increasing attributions of mind predict increased feelings of uncanniness) appear in middle to late childhood, in advance of adulthood. In order to be child friendly, our scale for uncanniness had a restricted range (e.g., allowing three degrees of creepiness and feeling weird) in comparison with scales for adults, and it may be that a more nuanced scale would reveal similar effects to those with adults. A direction for future research would be to look further at participants aged 10 to adulthood but with a revised rating scale to assess the development of a positive association between mind and uncanniness in later childhood.

Third, our results do not speak to degrees of the uncanny valley, as we used a binary comparison for machine-like and very human-like robots (consistent with Gray & Wegner, 2012). The classic uncanny valley proposal is that liking of robots follows a nonlinear curve (Mori et al., 2012) as in Figure 1, and future research with children should explore the full range of its trajectory. Still, our study provides three initial data points on this trajectory: a machine-like robot, an anthropomorphized robot, and a human-like robot. Although our study design (being a mix of within and between subjects) complicates the analysis of these three robot types somewhat, we can descriptively say that the anthropomorphic robot (Nao) was the least uncanny, followed by the machine-like robot, and finally the human-like robot as most uncanny.

It is moreover extremely likely that "humanness" is more than just a single dimension, thus plausibly

robots could be human like in several different multidimensional ways. Robots could be closely human like in face, limbs, behavior, language, and more. Future research should investigate which of these features of human likeness is considered creepy and at what ages. And given our data, an important question would be which features are tied to mind? Future research examining a larger variety of robots across childhood is clearly needed.

Fourth, one might argue that our results do not speak to children's changing perceptions of robots, but instead their developing understanding of words such as "weird," "feel," or "think." Our results with Nao, however, coupled with our additional data on young children's appropriate understanding of "weird" and "creepy" speak against this. Ample research also demonstrates that even preschoolers have appropriate understanding of the terms used in our questions eliciting children's attributions of mind, such as "think," "know," "feel," and "on purpose" (e.g., Bartsch & Wellman, 1995). Such findings make it difficult to argue that children have only shifted their understanding of the key words—creepy, weird, think, etc.—and are much more consistent with children appropriately using these terms to convey their developing conceptions of robots and robotic uncanniness and mind.

Understanding the development of the uncanny valley as an outgrowth of children's basic assumptions about robots coupled with increasing insights into minds provides a new perspective on this important phenomenon—it also suggests that one day the uncanny valley may disappear. As human-like robots become more commonplace and expand their abilities, children may come to expect that robots, although machines, *can* look surprisingly human, and *do* have minds, encompassing at least some human-like experiences. At which point, even highly human-like robots may be comfortably familiar to children—even as they continue to unnerve today's adults.

References

- Bartsch, K., & Wellman, H. M. (1995). *Children talk about the mind*. New York, NY: Oxford University Press.
- Beran, T. N., Ramirez-Serrano, A., Vanderkooi, O. G., & Kuhn, S. (2013). Reducing children's pain and distress towards flu vaccinations: A novel and effective application of humanoid robotics. *Vaccine, 31*, 2772–2777. <https://doi.org/10.1016/j.vaccine.2013.03.056>
- Bernstein, D., & Crowley, K. (2008). Searching for signs of intelligent life: An investigation of young children's beliefs about robot intelligence. *The Journal of the Learning Sciences, 17*, 225–247. <https://doi.org/10.1080/10508400801986116>
- Cohen, J. (1988). *Statistical power analysis for the social sciences*, 2nd ed. Hillsdale, NJ: Lawrence Earlbaum Associates.
- Dalecki, M., & Willits, F. K. (1991). Examining change using regression analysis: Three approaches compared. *Sociological Spectrum, 11*, 127–145. <https://doi.org/10.1080/02732173.1991.9981960>
- Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review, 114*, 864. <https://doi.org/10.1037/0033-295X.114.4.864>
- Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. New York, NY: Oxford University Press.
- Glaser, A. (2016). *Jibo is like Alexa and a puppy inside one adorable robot*. Retrieved from <http://www.wired.com/2016/06/jibo-like-alexa-puppy-inside-one-adorable-robot/>
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science, 315*, 619. <https://doi.org/10.1126/science.1134475>
- Gray, K., Knickman, T. A., & Wegner, D. M. (2011). More dead than dead: Perceptions of persons in the persistent vegetative state. *Cognition, 121*, 275–280. <https://doi.org/10.1016/j.cognition.2011.06.014>
- Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition, 125*, 125–130. <https://doi.org/10.1016/j.cognition.2012.06.007>
- Hall, D., & Williams, C. (Writers). (2014). *Big Hero 6* [Motion Picture]. United States: Walt Disney Studios Motion Pictures.
- Hashimoto, T., Kobayashi, H., Polishuk, A., & Verner, I. (2013). *Elementary science lesson delivered by robot*. Paper presented at the Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction, Tokyo, Japan.
- Hu, L., & Bentler, P. M. (1999). Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling: A Multidisciplinary Journal, 6*, 1–55. <https://doi.org/10.1080/10705519909540118>
- Jipson, J. L., & Gelman, S. A. (2007). Robots and rodents: Children's inferences about living and nonliving kinds. *Child Development, 78*, 1675–1688. <https://doi.org/10.1111/j.1467-8624.2007.01095.x>
- Jipson, J. L., Gülgöz, S., & Gelman, S. A. (2016). Parent-child conversations regarding the ontological status of a robotic dog. *Cognitive Development, 39*, 21–35. <https://doi.org/10.1016/j.cogdev.2016.03.001>
- Johnson, C. N., & Wellman, H. M. (1982). Children's developing conceptions of the mind and brain. *Child Development, 53*, 222–234. <https://doi.org/10.2307/1129656>
- Kahn, P. H., Kanda, T., Ishiguro, H., Freier, N. G., Severson, R. L., Gill, B. T., & Shen, S. (2012). "Robovie, you'll

- have to go into the closet now”: Children’s social and moral relationships with a humanoid robot. *Developmental Psychology*, 48, 303–314. <https://doi.org/10.1037/a0027033>
- Langlois, J. H., Roggman, L. A., & Rieser-Danner, L. A. (1990). Infants’ differential social responses to attractive and unattractive faces. *Developmental Psychology*, 26, 153. <https://doi.org/10.1037/0012-1649.26.1.153>
- Lasseter, J. (Writer). (1995). *Toy Story* [Motion Picture]. United States: Buena Vista Pictures.
- Low, A. (2016). *Asus’ adorable Zenbo home robot will be released in 9 to 12 months*. Retrieved from <http://www.cnet.com/news/asus-zenbo-will-only-be-released-in-9-12-months/>
- MacDorman, K. F. (2006). *Subjective ratings of robot video clips for human likeness, familiarity, and eeriness: An exploration of the uncanny valley*. Paper presented at the ICCS/CogSci-2006 Long Symposium: Toward social mechanisms of android science, Vancouver, BC.
- MacDorman, K. F., Green, R. D., Ho, C.-C., & Koch, C. T. (2009). Too real for comfort? Uncanny responses to computer generated faces. *Computers in Human Behavior*, 25, 695–710. <https://doi.org/10.1016/j.chb.2008.12.026>
- MacDorman, K. F., & Ishiguro, H. (2006). The uncanny advantage of using androids in cognitive and social science research. *Interaction Studies*, 7, 297–337. <https://doi.org/10.1075/is.7.3.03mac>
- Mejías, C. S., Echevarría, C., Nuñez, P., Manso, L., Bustos, P., Leal, S., & Parra, C. (2013). Ursus: A robotic assistant for training of children with motor impairments. In J. L. Pons, D. Torricelli, & M. Pajaro (Eds.), *Converging Clinical and Engineering Research on Neurorehabilitation* (pp. 249–253). Berlin, Heidelberg: Springer.
- Mikropoulos, T., Misailidi, P., & Bonoti, F. (2003). Attributing human properties to computer artifacts: Developmental changes in children’s understanding of the animate–inanimate distinction. *Psychology: The Journal of the Hellenic Psychological Society*, 10, 53–64.
- Milligan, K., Astington, J. W., & Dack, L. A. (2007). Language and theory of mind: Meta-analysis of the relation between language ability and false-belief understanding. *Child Development*, 78, 622–646. <https://doi.org/10.1111/j.1467-8624.2007.01018.x>
- Mori, M., MacDorman, K. F., & Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine*, 19, 98–100. <https://doi.org/10.1109/MRA.2012.2192811>
- Movellan, J., Eckhardt, M., Virnes, M., & Rodriguez, A. (2009). *Sociable robot improves toddler vocabulary skills*. Paper presented at the Proceedings of the 4th ACM/IEEE international conference on Human robot interaction, La Jolla, CA.
- National Robotics Initiative 2.0: Ubiquitous Collaborative Robots (NRI-2.0). (2017). Retrieved from <https://www.nsf.gov/pubs/2017/nsf17518/nsf17518.htm>
- Preacher, K. J., & MacCallum, R. C. (2003). Repairing Tom Swift’s electric factor analysis machine. *Understanding Statistics: Statistical Issues in Psychology, Education, and the Social Sciences*, 2, 13–43. https://doi.org/10.1207/S15328031US0201_02
- Richert, R. A., & Harris, P. L. (2006). The ghost in my body: Children’s developing concept of the soul. *Journal of Cognition and Culture*, 6, 409–427. <https://doi.org/10.1163/156853706778554913>
- Ricks, D. J., & Colton, M. B. (2010). *Trends and considerations in robot-assisted autism therapy*. Paper presented at the 2010 IEEE International Conference on Robotics and Automation (ICRA), Anchorage, AK.
- Severson, R. L., & Lemm, K. M. (2016). Kids see human too: Adapting an individual differences measure of anthropomorphism for a child sample. *Journal of Cognition and Development*, 17, 122–141. <https://doi.org/10.1080/15248372.2014.989445>
- Stanton, A. (Writer). (2008). *WALL-E* [Motion Picture]. United States: Walt Disney Studios Motion Pictures.
- Steckenfinger, S. A., & Ghazanfar, A. A. (2009). Monkey visual behavior falls into the uncanny valley. *Proceedings of the National Academy of Sciences*, 106, 18362–18366. <https://doi.org/10.1073/pnas.0910063106>
- Wei, C.-W., Hung, I.-C., Lee, L., & Chen, N.-S. (2011). A joyful classroom learning system with robot learning companion for children to learn mathematics multiplication. *TOJET: The Turkish Online Journal of Educational Technology*, 10(2), 11–23.
- Wellman, H. M. (2014). *Making minds: How theory of mind develops*. New York: Oxford University Press.
- Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72, 655–684. <https://doi.org/10.1111/1467-8624.00304>
- Wong, J. C. (2016). “This is awful”: robot can keep children occupied for hours without supervision. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2016/sep/29/ipal-robot-childcare-robot-business-san-jose>

Appendix

Interview

Warm-up questions:

1. Do you like candy?
 - 1.1. “Do you like candy?”
 - 1.2. “How much do you like candy? A little bit, a medium amount, or a lot?” or “How much do you not like candy? A little bit, a medium amount, or a lot?”
2. Do you like broccoli?
 - 2.1. “Do you like broccoli?”
 - 2.2. “How much do you like broccoli? A little bit, a medium amount, or a lot?” or “How much do you not like broccoli? A little bit, a medium amount, or a lot?”

- 3. Do you like carrots?
- 3.1. "Do you like carrots?"
- 3.2. "How much do you like carrots? A little bit, a medium amount, or a lot?" or "How much do you not like carrots? A little bit, a medium amount, or a lot?"

Interview questions:

- 4. Do you feel the robot is nice or creepy? +
- 4.1. "Do you feel the robot is nice (thumbs up) or creepy (thumbs down)?"
- 4.2. "How creepy do you feel it is? A little bit, a medium amount, or a lot?"
- 5. Does the robot make you feel weird or happy? +
- 5.1. "Does the robot make you feel weird (thumbs down) or happy (thumbs up)?"
- 5.1. "How weird does it make you feel? A little bit, a medium amount, or a lot?"
- 6. Would you want to play with the robot?
- 6.1. "Would you want to play with the robot?"
- 6.2. "How much would you want to play with it? A little bit, a medium amount, or a lot?"
- 7. Can the robot do things on purpose? * ++
- 7.1. "Can the robot do things on purpose?"
- 7.2. "How much can the robot act on purpose? A little bit, a medium amount, or a lot?"
- 8. When the robot moves, does it choose to move? +
- 8.1. "When the robot moves, does it choose to move?"
- 8.2. "How many things can the robot choose to do? A few things, a medium amount of things, or a lot of things?"
- 9. Does the robot think for itself? ++
- 9.1. "Does the robot think for itself?"
- 9.2. "How much does it think for itself? A little bit, a medium amount, or a lot?"
- 10. Some actions are bad, like hitting. And some actions are good, like helping. Does this robot know the difference between good and bad? +

- 10.1. "Does this robot know the difference between good and bad?"
- 10.2. "How much does it know the difference between good and bad? A little bit, a medium amount, or a lot?"
- 11. If I pinched the robot, would it feel pain? +
- 11.1. "If I pinched the robot, would it feel pain?"
- 11.2. "How much can this robot feel pain? A little bit, a medium amount, or a lot?"
- 12. Does the robot have feelings, like happy and sad? * ++
- 12.1. "Does the robot have feelings, like happy and sad?"
- 12.2. "How much does the robot have feelings? A little bit, a medium amount, or a lot?"
- 13. If the robot saw a snake, would it feel scared? +
- 13.1. "If the robot saw a snake, would it feel scared?"
- 13.2. "How much can the robot feel scared? A little bit, a medium amount, or a lot?"
- 14. If the robot did not eat breakfast, would it feel hungry? +
- 14.1. "If the robot did not eat breakfast, would it feel hungry?"
- 14.2. "How much can the robot feel hungry? A little bit, a medium amount, or a lot?"
- 15. Is this robot like a human? *
- 15.1. "Is this robot like a human?"
- 15.2. "How much is the robot like a human? A little bit, a medium amount, or a lot?"
- 16. Does the robot know it's a robot? * ++
- 16.1. "Does the robot know it's a robot?"
- 16.2. "How much does it know it's a robot? A little bit, a medium amount, or a lot?"

* Item not included in final factor analysis due to cross loading or low factor loadings.

+ Item derived from Gray and Wegner (2012) and/or Gray et al. (2007).

++ Item derived from Severson and Lemm (2016).