Article

PERCEPTION

Human Perception of Animacy in Light of the Uncanny Valley Phenomenon

Perception
2017, Vol. 46(12) 1386–1411

DOI: 10.1177/0301006617722742
journals.sagepub.com/home/pec
SAGE

Shensheng Wang and Philippe Rochat
Emory University, Atlanta, GA, USA

Abstract

The uncanny valley hypothesis by Japanese roboticist Masahiro Mori posits a nonlinear relation between human replicas' human likeness and the emotional responses they elicit. In three studies, we corroborated the uncanny valley hypothesis, using the uncanny phenomenon as a vehicle to shed a new light on human animacy perception. In Study 1, 62 participants rated emotional responses and human likeness of 89 artificial and human faces. In Study 2, another 62 participants conducted a visual looming task with the same 89 faces allowing for the measurement of perceived threat. Results support the uncanny valley hypothesis, suggesting that the uncanny feeling may serve a function to wary humans of the potential danger of entities crossing the animate–inanimate boundary. In Study 3, 36 participants sorted faces as either real or unreal as quickly as possible in a reaction time sorting task allowing for the measurement of categorical uncertainty associated with animacy perception. Faces associated with longer sorting reaction times were also those associated with the highest ratings of negative emotions, suggesting that categorical uncertainty in animacy detection is related to the uncanny feeling. Results are discussed in light of human animacy perception and new directions for future research are suggested.

Keywords

animacy, face, uncanny valley, visual looming, categorical uncertainty

A challenge for roboticists and filmmakers draws scientists' attention to the uncanny valley hypothesis. The uncanny valley hypothesis was first coined by Japanese roboticist Masahiro Mori (1970/2005), who predicted that as robots become more and more realistically human, they become increasingly likable to people. However, just prior to becoming 100% humanlike in appearance, there exists a sudden drop in likability accompanied by an intense uncanny feeling (Mori, 1970/2005).

By analogy, Mori (1970/2005) described that this uncanny feeling is one of the sort we experience when shaking someone's hand until we are caught by surprise noticing that it is cold and part of a prosthetic arm. This experience is typically associated with an eerie sensation (Mori, 1970/2005; Poliakoff, Beach, Best, Howard, & Gowen, 2013).

Corresponding author:
Shensheng Wang, Emory University, 36 Eagle Row, Atlanta, GA 30322, USA.
Email: swan232@emory.edu

Mori argued that human likeness of a human replica increases its likability, yet becomes repulsive when too realistic. Anecdotal evidence corroborates the negative effects of high human likeness in robots and computer-generated characters. Robots like KASPAR Two (Gray & Wegner, 2012) and Repliee Q2 (Saygin, Chaminade, Ishiguro, Driver, & Frith, 2012) were reported to elicit negative reactions of eeriness, repulsion, and unease in people. Similar effects were reported for wax figures such as those in the film *House of Wax* (Collet-Serra, 2005), realistic wax sculptures, mannequins, and dolls (Brenton, Gillies, Ballin, & Chatting, 2005), as well as computer-generated characters in films like *The Polar Express* (Zemeckis, 2004; see also horror movies like *The Hills Have Eyes*, Aja, 2006).

Discussion of the ‘‘uncanny’’ feeling associated with human replicas is not new and can be dated back to the early 20th century when Ernst Jentsch (1906/1997) and Sigmund Freud (1919) both mentioned the ‘‘uncanny’’ experience in the context of Hoffmann’s story *Sandman* (1817/1994) featuring a lifelike doll Olympia. Nevertheless, the uncanny experience is not limited to dolls and androids. Freud associated the uncanny feeling with the fear of being castrated, buried alive, or the eerie experience of coincidences. Aside from mere introspective speculation and casual clinical observations, the psychoanalytic interpretation of ‘‘uncanniness’’ is difficult to test experimentally.

Mori and more recently other researchers have argued that the uncanny feeling toward human replicas is elicited by the imperfections in robots’ appearance and behavior (MacDorman & Ishiguro, 2006; Mori, 1970/2005). Research has identified a variety of imperfections, including physical appearance of the face (Hanson, 2005; MacDorman, Green, Ho, & Koch, 2009; Seyama & Nagayama, 2007), emotional expressions (Tinwell, Grimshaw, Nabi, & Willams, 2011), body movements (Groom et al., 2009; Saygin et al., 2012), and voice (Mitchell et al., 2011), alone or in combination (Tinwell, Grimshaw, & Williams, 2010), that may elicit the uncanny feeling.

Although various imperfections of a robot supposedly elicit the uncanny feeling, researchers have yet to agree on the ways in which these imperfections are perceived by humans to provoke the uncanny feeling. Mori (1970/2005) proposed that this uncanny feeling would derive from human replicas’ high levels of human likeness. Nevertheless, Mori’s proposal has been criticized (for a review, see Wang, Lilienfeld, & Rochat, 2015). Evidence in support of the nonlinear relation between human likeness and emotional responses has been mixed and at times contradictory (Bartneck, Kanda, Ishiguro, & Hagita, 2007; Hanson, 2005; MacDorman, 2006; MacDorman & Ishiguro, 2006; Mathur & Reichling, 2016; Poliakoff et al., 2013; Seyama & Nagayama, 2007). In sum, Mori’s account of the uncanny feeling as deriving from human replicas’ high levels of human likeness is either unconvincing or at least incomplete.

Given the inconsistent findings and interpretations in the literature, in Study 1 and Study 2, we sought to first validate the uncanny valley hypothesis as to ascertain the extent to which the uncanny feeling might correspond to the level of human likeness of human replicas. Even more importantly, we aimed at validating the uncanny valley hypothesis, with the ultimate goal to understand human animacy perception (e.g., Looser, Guntupalli, and Wheatley, 2013; Looser & Wheatley, 2010; Wheatley, Weinberg, Looser, Moran, & Hajcak, 2011). If humans were not particularly tuned to perceiving animacy (e.g., the discrimination between real and unreal faces), we would expect that with gradually increasing levels of human likeness would come with a monotonic increase in likability of the human replicas, a result that would run counter to what the uncanny valley hypothesis predicts.

Researchers have proposed different hypotheses explaining the uncanny feeling, including pathogen avoidance, evolutionary aesthetics, mortality salience, violation of expectation, categorical uncertainty, and mind perception. From an evolutionary perspective, the pathogen avoidance hypothesis conceives the uncanny feeling as deriving from evolved mechanisms of pathogen avoidance (Curtis, Aunger, & Rabie, 2004; Ho, MacDorman, & Pramono, 2008; MacDorman & Ishiguro, 2006; Rozin & Fallon, 1987). In contrast, the evolutionary aesthetics hypothesis links the uncanny feeling to evolved mechanisms of mate selection (MacDorman et al., 2009; MacDorman & Ishiguro, 2006; Rhodes et al., 2001). Based on the terror management theory (Greenberg, Pyszczynski, Solomon, Simon, & Breus, 1994), MacDorman and Ishiguro (2006) proposed the mortality salience hypothesis, which suggests that the robots eliciting the uncanny feeling might invite implicit thoughts of death. According to these hypotheses, the uncanny feeling is construed as relating to either disgust, attractiveness, or a fear of death, respectively.

In Study 1, we drew on these different interpretations of the uncanny feeling and conducted a survey on 89 static facial images to validate the uncanny valley hypothesis. Participants rated the human likeness of each face as well as eight different emotional responses to it, including "eeriness," "disgustingness," "unsettlingness," "attractiveness," "threateningness," "likableness," "real" and "aliveness." We then plotted ratings of these eight different emotional responses against those of human likeness, trying to validate for each the hypothetical uncanny valley depicted by Mori (Figure 1; Mori, 1970/2005, p. 33). As would be predicted by Mori (1970/2005), we expected to find a dip and a rebound in the

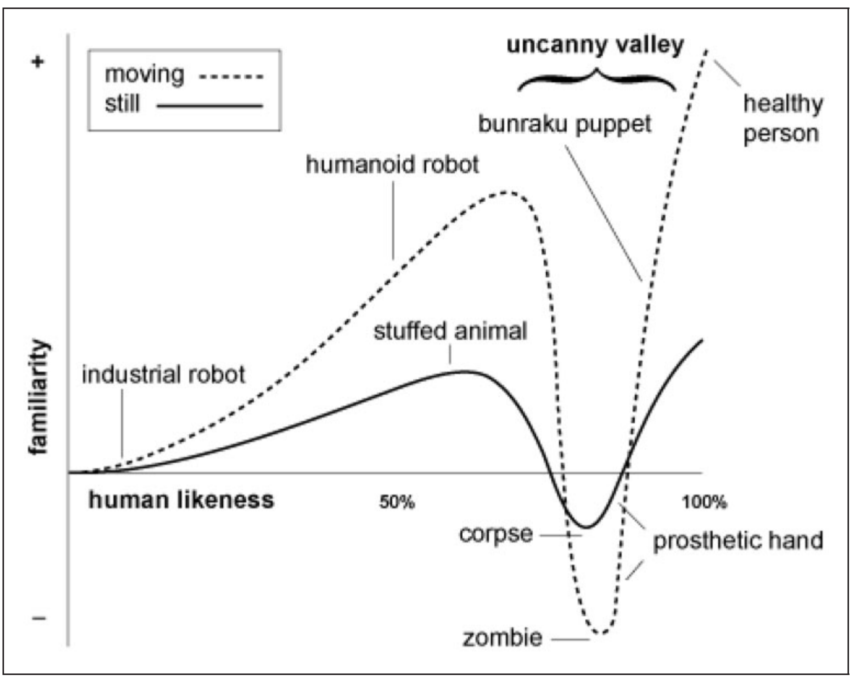

Figure 1. Reprinted from Mori (1970/2005). The uncanny valley depicts familiarity as a function of human likeness of the human replicas. As robots increasingly resemble a person, their familiarity increases. However, prior to eventually reaching the level of a person, the robots' familiarity suddenly drops to a negative value accompanied by an intense uncanny feeling. Mori (1970/2005) hypothesized that movement deepens the uncanny valley.

Note. The Japanese term "shinwakan," which appeared on the *y*-axis, was first translated to "familiarity" (Mori, 1970/2005), but this translation was then replaced in favor of terms such as "affinity" or "likability" (Mori, MacDorman, & Kageki, 2012) .

curve as faces approached human likeness. By searching the existing literature, we found little consensus on the criteria researchers adopted to determine the uncanny valley (Bartneck et al., 2007; Hanson, 2005; MacDorman, 2006; Mathur & Reichling, 2016; Mori, 1970/2005; Seyama & Nagayama, 2007), except for the agreement that a null hypothesis of the uncanny valley is characterized by a linear relation between emotional responses and human likeness. Therefore, we did not commit to a precise shape of the uncanny valley, simply searching for the evidence that allows us to reject this linear relation.

In Study 2, we conducted a computer-based visual looming task to validate further the uncanny valley hypothesis by adding an implicit measure (estimated relative to actual time-to-contact) of emotional response (i.e., perceived threat) to our analysis. To the extent that perceived threat might capture the uncanny feeling, we aimed at probing the link between animacy perception and threat detection. In the visual looming task, we presented the same 89 faces looming in succession towards the participant on a computer screen. After the looming face disappeared, we asked participants to estimate when the looming face would make contact with their own faces. We systematically recorded variations in their reaction times in relation to each looming face. As previous research has shown, more threatening objects were judged to move faster than less threatening objects (Vagnoni, Lourenco, & Longo, 2012). Therefore, estimated relative to actual time-to-contact was used as an objective measure of participants' perceived threat. Assuming that the uncanny feeling is driven by a fear of death according to the mortality salience hypothesis, we predicted that the faces eliciting the uncanny feeling would be perceived by participants as being more threatening and looming faster toward them, thereby eliciting shorter estimated time-to-contact. We expected to validate further the uncanny valley hypothesis by showing that a dip and a rebound in estimated relative to actual time-to-contact when plotting it against the ratings of human likeness of the 89 faces established in Study 1.

In addition to the different evolutionary hypotheses of the uncanny feeling, Jentsch (1906/1997) proposed that the uncanny feeling is linked to a sense of uncertainty. In his article "On the psychology of the uncanny," Jentsch associated uncanniness with "a lack of orientation" (p. 8). In particular, he attributed the uncanny effect of Hoffmann's (1817/1994) "The Sandman" to the uncertainty about whether the automaton Olympia—a protagonist in the story—was a real person. More modern and experimental approaches have echoed what Jentsch intuited over a century ago. For example, Yamada, Kawabe, and Ihaya (2013) tested whether the uncanny feeling derives from the difficulty in categorizing morphed images. The researchers created morphed images of a real human face, a stuffed human face, and a cartoon human face and asked participants to categorize them as belonging to either the human or the doll category as quickly as possible. Participants also rated the likability of these morphed images in an evaluation task without time constraint. The results showed that faces associated with the highest category ambiguity were the same faces that elicited the longest response latency and the lowest ratings of likability. The researchers replicated this finding using morphed images of the faces of a real and a doll dog but did not find the same effect using morphed images of a male and a female's faces. Overall, these findings suggest that categorical uncertainty associated with animate–inanimate discrimination may contribute to the uncanny feeling.

In light of this categorical uncertainty hypothesis, in Study 3, we examined whether animacy detection might be related to the uncanny feeling using a reaction time-based sorting task. In the sorting task, we presented each of the 89 faces on a computer screen one at a time and asked participants to sort them as either "real" or "unreal," doing so as quickly as possible. We analyzed the choices as well as the reaction times to test if the faces that were judged with less certainty would correspond to those that elicited stronger negative emotional responses in Study 1. Congruent with the idea that categorical uncertainty is a good candidate for explaining the uncanny feeling, we expected a tendency toward significantly higher reaction times in the real–unreal dichotomous sorting task for faces that in Study 1 were ranked high in overall negative emotions.

Study 1

In this first study, adult participants were asked to rate their emotions and judge the relative human likeness of a large collection of faces varying in human likeness, from mechanical, lifeless, and artificial looking to real human faces ($N = 89$). The primary goal was to further validate the uncanny valley hypothesis by plotting emotional ratings in relation to the rating of human likeness of all the faces.

Method

Participants. Sixty-two undergraduate psychology students (39 females; $M_{age} = 24.18$ years, $SD = 9.01$) from Georgia Highlands College participated in return for research participation credits. All participants provided oral informed consent. The experimenter provided debriefing to all the participants at the end of the study. Based on a posttest interview of the participants, none of them acknowledged knowing the uncanny valley phenomenon before participating in this study.

Stimuli and procedure. All 89 faces, including 40 "uncanny" faces and 49 "normal" looking human faces by searching "uncanny valley" and "human faces" in Google Images, respectively. The uncanny images consisted of human replicas (e.g., androids, dolls, wax figures, mannequins, computer-generated characters, and Zombies) and other pathological, computer, or plastic surgery altered faces of real people (e.g., human Barbie dolls, Bell's palsy and symmetrical faces created by mirroring one side of the face). Gender, relative age, and facial features like facial hair were equated among uncanny and normal human faces. All images were resized (300-pixel width and 380-pixel height) and were cropped to an oval shape in Adobe Photoshop software to control for background. Figure 2 shows the collection of all 89 faces used in the experiment for emotional and human likeness ratings.

A paper-based survey testing took place in a large classroom with 62 participants. Before the survey, a female experimenter informed the participants that they were going to see some faces and instructed them to rate each face on eight statements on a 7-point Likert scale (1 = *Not at all* and 7 = *Very much*, see Table A1). During the survey, the experimenter presented in succession (the order of each face was randomly assigned and remained constant across participants) via PowerPoint the slides of the 89 faces on a large classroom screen (170 cm width 127 cm height) located 4 m from the front row. Slides were presented one at a time for 40 seconds each giving all participants sufficient time to rate a face on the eight statements described in Table A1. Completion of the whole survey took approximately 1 hour.

The eight statements of the survey corresponded to two possible forms, either "This figure makes me feel_____," or "This figure looks _____," the blank filled with one of eight corresponding adjectives: "eerie," "disgusted," "unsettling," "attractive," "threatening," "likable," "real," and "alive." Among these terms, "real" was included to measure human likeness. "Disgusted," "threatening," "alive," and "attractive" were included to test whether we could detect the uncanny valley in reference to existing hypotheses regarding what might account for the uncanny feeling, in particular pathogen avoidance

Figure 2. A reproduction of 89 faces used in the three studies, including 49 "normal" looking human faces and 40 "uncanny" faces by searching "uncanny valley" and "human faces" in Google Images, respectively. As shown in Figure 2, the "normal" looking human faces included the faces in the three uppermost rows and the fourth row up to the ninth column. All the other faces in Figure 2 were "uncanny" faces. Selection criteria included (a) faces are in frontal view, (b) there are no obvious facial expressions of emotions, and (c) resolution of the original images allows cropping and resizing to 300 pixel × 380 pixel.

(Curtis et al., 2004; Rozin & Fallon, 1987), mortality salience (MacDorman & Ishiguro, 2006), and evolutionary aesthetics (Hanson, 2005; MacDorman et al., 2009). Ratings of "eerie" and "unsettling" emotions associated with each face were viewed as synonymous of "uncanny." "Likable" was used as the *y*-axis variable—*shinwakan*—of the uncanny valley (see also Bartneck et al., 2007). The general rationale was that the uncanny valley hypothesis would be corroborated if emotional evaluation varies as a function of human likeness ("real" evaluation on the questionnaire), which would map onto the hypothetical curve shown in Figure 1.

For this first study, we hypothesized that as a function of human likeness ("real" evaluation), there would be an increase followed by a valley in the plotted curve for both ratings of attractiveness, likableness, and aliveness, as well as the opposite values of negative ratings of eeriness, disgustingness, unsettlingness, and threateningness (Seyama & Nagayama, 2007).

Due to potential individual differences in the sensitivity to the uncanny phenomenon (MacDorman & Entezari, 2015), participants may adopt different criteria to evaluate their emotional responses and rate images on different scales. Therefore, for data analysis, we first calculated *z* scores of the ratings of the 89 faces on each statement by each individual participant. We then averaged the *z* scores of each face to establish an overall evaluation by all participants for a particular face. We then plotted the average *z* scores for each face

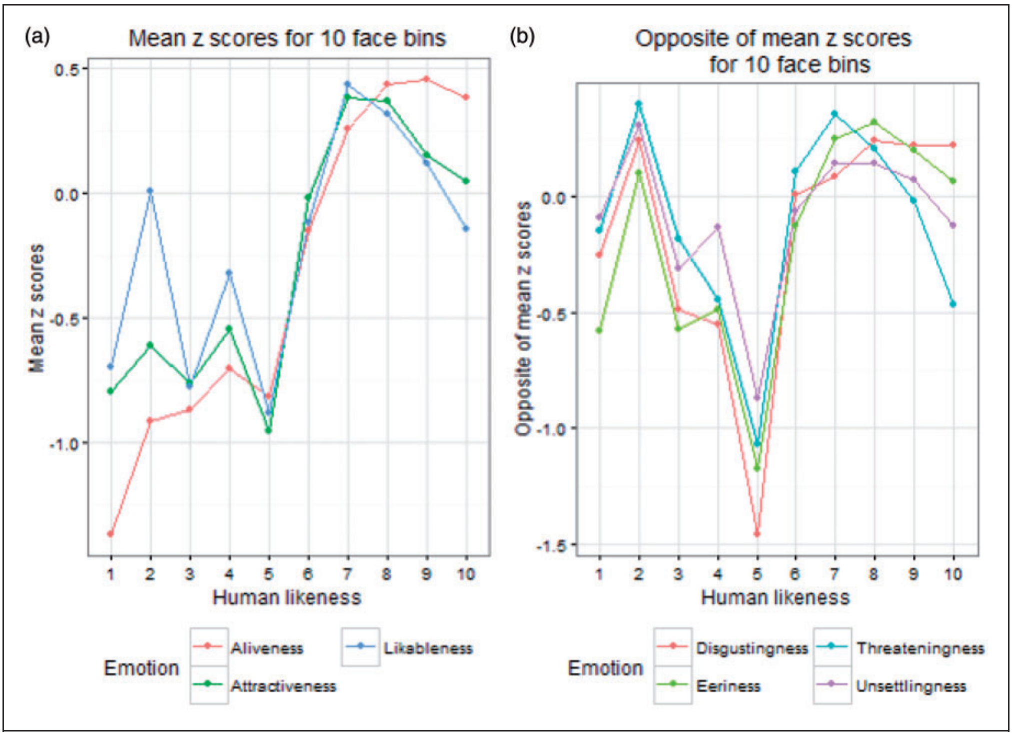

Figure 3. (a) Mean group *z* scores on the three positive emotions (e.g., aliveness, attractiveness, and likableness) and (b) the opposite of the four negative emotions (e.g., disgustingness, eeriness, threateningness, and unsettlingness) as a function of human likeness (average "real" ratings by 10 bins).

obtained for each of the seven evaluations (*eerie*, *disgusted*, *unsettling*, *attractive*, *threatening*, *likable*, and *alive*) against the average *z* scores on the human likeness ("real") evaluation of each face. We thus obtained seven line graphs potentially capturing the uncanny valley.

In order to filter psychometric noises, we first divided the range of the *z* scores of human likeness by 10. We then divided the 89 faces, according to their *z* scores of human likeness, into the resulting 10 progressive bins (from low to high human likeness) in the human likeness ("real") dimension on the *x*-axis, and the corresponding seven other emotional ratings on the *y*-axis. We thus plotted group means per bin rather than individual face average. To map onto Mori's hypothetical curve, we calculated the opposite values of negative ratings (i.e., eeriness, disgustingness, unsettlingness, and threateningness evaluations).

Results

Figure 3 presents the seven-line graphs for each of the emotional ratings as a function of the average ratings of human likeness of the faces divided by 10 bins, from lowest to highest ratings of human likeness.

Figure 3 shows that except for the evaluation of aliveness, there was clear evidence of an uncanny "valley" in all the regression curves. A Pearson's correlation was then conducted to confirm what the line graph revealed as a linear relation between human likeness and aliveness. Results showed a strong positive correlation between ratings of aliveness and ratings of human likeness, $N = 89$, $r(87) = .88$, $p < .001$, 95% CI [0.82 0.92], suggesting that the two variables are highly interdependent. Results showed significant and negative correlations between ratings of human likeness and ratings of eeriness ($r = -.52$) and disgustingness ($r = -.34$). The correlation between ratings of human likeness and ratings of threateningness and unsettlingness did not reach statistical significance. Finally, ratings of human likeness were positively and significantly correlated with ratings of attractiveness ($r = .58$) and likableness ($r = .47$).

For ratings in relation to the opposite of negative emotions, the results showed that the fifth face bin corresponded to the trough of valley and the second and the eighth face bin corresponded to the first and second peak, respectively (except for threateningness, the second peak corresponded to the seventh bin). We compared the ratings at the trough of the valley to those at the two peaks. We used Bonferroni adjustments with a significance level set at $\alpha = .016$ for all pairwise comparisons across the studies. A one-way repeated measures analysis of variance (ANOVA) yielded a significant difference in the opposite of eeriness ratings among the three groups of faces, $F(2, 122) = 109.70$, $p < .01$, partial $\eta^2 = .63$. Faces belonging to the fifth bin were rated significantly eerier compared with the faces belonging to the second and eighth bins. Pairwise comparisons showed that eeriness ratings at the fifth bin ($M = 1.17$; $SD = 0.52$) were significantly higher than those at the second bin ($M = -0.10$; $SD = 0.86$) and the eighth bin ($M = -0.32$; $SD = 0.12$), ps $< .01$. We found similar results for disgustingness, $F(2, 122) = 345.9$, $p < .01$, partial $\eta^2 = .85$; unsettlingness, $F(2, 122) = 40.10$, $p < .01$, partial $\eta^2 = .40$; and threateningness, $F(2, 122) = 151.50$, $p < .0001$, partial $\eta^2 = .71$ (ps $< .01$).

Congruent with the preceding data, the ratings of positive emotions were lowest for faces belonging to the fifth bin. Regarding attractiveness ratings, a one-way repeated measures ANOVA yielded significantly lower ratings at the fifth bin compared with the two peaks framing the valley, $F(2, 122) = 222.00$, $p < .01$, partial $\eta^2 = .78$. We found similar results for likableness, $F(2, 122) = 71.60$, $p < .01$, partial $\eta^2 = .54$. The two peaks corresponded to bins 2 and 7 for likableness; bins 4 and 7 for attractiveness (see Figure 3; ps $< .01$).

Discussion

Except for the ratings of aliveness, our results show that ratings of all the other emotional responses—eeriness, disgustingness, unsettlingness, attractiveness, threateningness, and likableness followed a nonlinear relation with the ratings of human likeness. These data corroborated the existence of the uncanny valley, consistent with the recent findings of Mathur and Reichling (2016) based on a similar methodological approach. In particular, we found that the valley corresponded to the fifth bin and the two peaks corresponded to the second and the eighth bins. One might wonder why the peaks tended to occur at bins 2 and 8, rather than at the extremes of 1 and 10. We believe that the first peak did not appear at the first bin, because the human likeness of images in the first bin may not be sufficiently humanlike to elicit more positive emotions in the participants. However, the interpretation of the location of the second peak is not straightforward. First, the location of the second peak slightly varied across the emotional ratings. Second, following the second peak, there was a decline in the positive ratings or the opposite of negative ratings of the images. We interpreted this decline as evidence that does not run counter to the uncanny valley but indicates that observers might adopt different criteria in emotional ratings of faces for humans and robots (for a discussion, see Bartneck et al., 2007). This is supported by the fact that with increasing levels of human likeness, there was an increasing proportion of human faces in each bin.

However, our findings are inconsistent with those of Bartneck et al. (2007), MacDorman (2006), and Poliakoff et al. (2013), who did not find evidence of the uncanny valley. A variety of methodological differences may be related to these inconsistent findings. First, except for MacDorman (2006) who presented video clips of 13 robots and 1 person, all of the other studies used images of real-world human replicas. As MacDorman (2006) suggested, the different contexts in which the videos were produced and whether the depicted robots were capable of humanlike movements or not added considerable psychometric noises to the data. Second, the particular part of a human replica under examination varied from faces (Bartneck et al., 2007; Mathur & Reichling, 2016) to hands (Poliakoff et al., 2013) and a mix of faces and the whole body (MacDorman, 2006). Third, the stimuli varied substantially across these studies, both in size and in variability, from MacDorman (2006) who obtained a sample of 1 individual person and 13 robots of varied degrees of human likeness, to Bartneck et al. (2007) who obtained a sample of 18 entities, 3 for each of the 6 categories (e.g., toy robots, humanoids, androids, computer-generated characters, manipulated humans, and real humans), to Poliakoff et al. (2013) who obtained a sample of 22 human, prosthetic, and mechanical hands, and to Mathur and Reichling (2016) who obtained a sample of 80 real-world robots of various levels of human likeness. Finally, these studies differed in their measures of the independent and the dependent variables. For example, researchers measured human likeness in a variety of ways, ranging from a 9-point scale (1 = *very mechanical* and 9 = *very humanlike*, MacDorman, 2006) to 7-point semantic differential scales (e.g., fake/natural, machinelike/humanlike, unconscious/conscious, and artificial/lifelike, Bartneck et al., 2007), a 9-point scale on human likeness, with the term "human likeness" being defined as "having human form or attributes" (Poliakoff et al., 2013, p. 999), and to a continuous visual analog scale (by indicating "how mechanical/human does this robot face look?" on a 0 to 100 scale, Mathur & Reichling, 2016, p. 24). As it stands, it is thus difficult to compare and evaluate evidence from these studies that would decisively either support or reject the uncanny valley hypothesis.

Noteworthy is also the fact that we found no evidence of a valley when considering the relation between the ratings of human likeness ("real" ratings) and the ratings of aliveness. The positive correlation between these two variables runs counter to our prediction. The fact that ratings of aliveness positively correlate with ratings of human likeness suggests that individuals' evaluations of human likeness and aliveness are interdependent.

Considering that each participant rated both human likeness and aliveness of all 89 faces, the ratings of these two dimensions might have influenced each other via priming or positive interference yielding to their high correlation. Independent ratings of these two dimensions in a between-subject design could control for such interference. Furthermore, the combined use of explicit and implicit measures (e.g., implicit association tests) of these dimensions would help in further assessing the relative psychological correspondence between human likeness and aliveness.

In general, the present findings demonstrate that except for the ratings of aliveness, all the others (eeriness, disgustingness, unsettlingness, attractiveness, threateningness, and likableness) follow a nonlinear relation with the ratings of human likeness, therefore corroborating Mori's prediction. Nevertheless, based on these findings, it was not yet possible to specify which emotional response might best characterize the uncanny feeling, thus providing more or less support for existing hypotheses regarding the uncanny feeling (e.g., pathogen avoidance, mortality salience, and evolutionary aesthetics). Converging evidence suggests that the uncanny feeling, instead of being a unitary emotion, may involve a variety of emotions, including fear, disgust, and nervousness (Ho et al., 2008) or fear, disgust, and attractiveness (Burleigh, Schoenherr, & Lacroix, 2013). Mangan (2015) conceived the uncanny feeling as being an imminent fringe experience, and involves "a rare and complex mix of familiarity, wrongness, and a feeling of threat" (p. 198) and in some cases disgust. Our findings that all of the emotional responses of eeriness, disgustingness, unsettlingness, attractiveness, threateningness, and likableness contribute to the uncanny feeling are broadly consistent with this view.

One limitation in the present study concerns the assessment of the dependent variable—Shinwakan (Mori, 1970/2005). Drawing on previous interpretations of Shinwakan as well as on the various existing hypotheses put forth in the literature to account for the uncanny feeling, we assessed Shinwakan and the uncanny valley hypothesis in reference to emotional terms like eeriness, disgustingness, unsettlingness, attractiveness, threateningness, likableness, and aliveness. Existing literature, however, provides no direct evidence or mixed evidence that either terms or hypotheses attached to them might fully explain the uncanny feeling (MacDorman & Ishiguro, 2006). In a recent review of the uncanny valley research, we argued that it is important to distinguish between the empirical testing of the uncanny valley hypothesis (e.g., whether there exists a nonlinear relation between human likeness and likability as Mori's hypothetical curve depicted) from explaining the mechanisms that might account for the uncanny feeling as proposed by the aforementioned hypotheses of pathogen avoidance and mortality salience, among other hypotheses (Wang et al., 2015). In light of this distinction, future research should adopt a standard interpretation of the term "shinwakan," namely, likability or affinity (Mori et al., 2012), rather than trying to incorporate the purported interpretations of the uncanny feeling into the validation of the uncanny valley hypothesis.

Finally, and motivating our second study, previous research has shown that the use of self-report is sometimes associated with lower measurement accuracy. For example, Ho and MacDorman (2010) point to the fact that correlating self-report measures of human likeness and emotional reactions in the Godspeed indices (Bartneck, Kulić, Croft, & Zoghbi, 2009) significantly lowers the accuracy of the plotted graphs. This effect is ostensibly attributable to the fact that participants' subjective ratings of human likeness invite not only their evaluations of the physical similarity between human replicas and real humans but also their perception of animacy in and the attribution of human characteristics to the human replicas. These processes of animacy perception (Looser & Wheatley, 2010) and anthropomorphism in turn are interweaved with participants' evaluations of likability (Gray & Wegner, 2012). In the second experiment, we therefore tested the 89 faces in an implicit paradigm (looming task), adding an implicit measure (estimated relative to actual time-to-contact) to the self-report based analysis we reported so far. We probed whether we could reproduce and corroborate the uncanny valley hypothesis based on an implicit measure such as this.

Study 2

In a second study, we thus adopted a behavioral test—visual looming task—to evaluate emotional responses, particularly fear, in relation to the faces and in relation to their human likeness ("real" ratings). Assuming that the emotion of fear captures the uncanny feeling, as suggested by the mortality salience hypotheses in particular, we tested whether plotting the estimated relative to actual time-to-contact of each individual face looming toward the participant on a computer screen would reproduce the uncanny valley.

Objects that systematically expand in size in relation to the self create a visual equivalent of impending contact in the field of view of the participants, eliciting avoidant responses. The visual looming effect is documented in human adults (Coker-Appiah et al., 2013), in human infants from 2 months of age (Ball & Tronick, 1971), as well as in adult rhesus monkeys (Schiff, Caviness, & Gibson, 1962), thus is deeply rooted both in development and in evolution. Studies indicate that the looming effect is modulated by perceived threat of the looming stimuli with threatening animals like spiders and snakes perceived as moving faster than nonthreatening animals (e.g., butterflies and rabbits, Brendel, DeLucia, Hecht, Stacy, & Larsen, 2012; Vagnoni et al., 2012). Specifically, there is evidence of greater avoidant behavior for the former as indexed by systematic underestimation of imagined time-to-contact. In adults, the magnitude of underestimation correlates with the degree of self-reported fear (Vagnoni et al., 2012).

If uncanny faces primarily elicit fear as would be predicted by the mortality salience hypothesis, the uncanny valley hypothesis should find some support in the estimation of time-to-contact of faces looming toward the participant. We therefore tested the uncanny valley hypothesis by plotting participants' estimated relative to actual time-to-contact against the ratings of human likeness (real rating) of that particular face, probing whether at this implicit level we could detect the nonlinear feature of the uncanny valley.

Method

Participants. Sixty-two undergraduate psychology students (43 females, $M_{\text{age}} = 19.60$ years, $SD = 1.89$) participated in return for course credit as part of a large introduction to psychology class. All participants were blind to both the rationale and hypothesis guiding the study prior to the poststudy debrief.

Stimuli and procedure. Participants sat at a table about 40 cm in front of a 19-in. monitor (75-Hz refresh rate) with their head resting comfortably on a chin rest affixed to the table, instructed to remain in it all through testing. The same 89 faces from Study 1 were presented in a program run by MATLAB (Mathworks, Natick, MA). On each trial, a face gradually enlarged for a second at a given rate, corresponding to 2.5, 3.5, and 4.5 s of time-to-contact. On the first frame, the width of the face was either 400 or 500 pixels (15.1° or 18.9° visual angle). The size of image differed so that estimation of time-to-contact involved both the speed and the initial size of an image.

Ninety images were presented twice and were randomly assigned to each of six (3 time-to-contact × 2 initial size) conditions. The resulting 180 trials were divided into 30 blocks (each block consisted of six trials corresponding to six conditions). Since it was impossible to distribute 89 faces evenly to six presentation conditions, one face (Face 89) was duplicated to match the demand of the experimental design. The order of trials was randomized. As a result, the order of the 90 faces varied across the participants. After each participant responded to one trial, the next trial began after a random interval of 300 to 800 ms.

Participants were all instructed that they would see faces expanding in size as if they were approaching and that the faces would disappear after a second. From the moment the face disappeared, the participant was instructed to imagine it continuing to approach at the same speed and to press the space bar of the computer keyboard when they judged that the face would have come into contact with the participant's own face. Reaction times of estimated time-to-contact were calculated based on the time elapsed between the moment each looming face disappeared from the screen and the space bar being pressed by the participant, measuring the estimated time-to-contact. Prior to actual testing with the faces, each participant ran a block of eight practice trials with looming geometric shapes. Before launching the testing session, the experimenter repeated the practice session until participants indicated that they followed the instruction. The testing session lasted on average 10 to 15 minutes.

Results

Raw data showed extreme reaction times probably attributable to fatigue or attention failure by some of the participants. (*Note*. The computer will not register a button pressing as a valid response until the face disappeared. Failing to press the button after the face disappears results in extremely long reaction times.). In order to eliminate outliers, and because reaction times were obtained from six different presentation conditions varying along actual time-to-contact and initial size (see Method section), we standardized reaction times by the same participant for each face in the six presentation conditions separately. Cases with a z score above 3 or below -3 were operationally considered as outliers and excluded from the final analysis (total of 109 trials or 0.99% of all trials).

A repeated measures ANOVA on the two presentation factors (i.e., actual 2.5, 3.5, and 4.5 s time-to-contact and 400 vs. 500 pixels wide initial size) yielded significant main effects for both, and no significant interaction. Descriptive statistics of reaction times on the six presentation conditions were reported in Table A2. Overall, participants perceived faces that started closer to them as making contact with their own faces within a shorter period of time, $F(1, 61) = 41.37$, $p < .05$, partial $\eta^2 = .404$. The same was true for those with shorter actual time-to-contact, $F(2, 122) = 48.02$, $p < .05$, partial $\eta^2 = .44$. Pairwise comparisons showed that the faces looming toward the participants with an actual time-to_contact of 2.5 s ($M = 2907$; $SD = 946$) elicited significantly shorter reaction times than those with an actual time-to-contact of 3.5 s ($M = 3458$; $SD = 1387$; $t_{61} = -7.34$, $p < .01$) and 4.5 s ($M = 3811$; $SD = 1624$; $t_{61} = -7.56$, $p < .0001$), and that the reaction times associated with an actual time-to-contact of 3.5 s were significantly shorter than those associated with an actual time-to-contact of 4.5 s ($t_{61} = -4.50$, $p < .01$). In general, results indicated that participants tended systematically to underestimate the time-to-contact of looming faces. Six participants showed the reversed pattern (i.e., they did not demonstrate effect of underestimation of time-to-contact in the visual looming paradigm) and were excluded from further analysis. Since there was no significant interaction between actual time-to-contact and initial size, we collapsed trials across initial size levels.

In Study 2, participants viewed each image twice under any of the three levels of time-to-contact condition, which rendered a direct comparison of the images under the same viewing condition impossible. In order to obtain an index that would reflect the emotional modulating effect of each face, we averaged the reaction times for each face under each of the three actual time to contact levels across participants. To do so, we first converted reaction times into z scores for each level of actual time to contact for each participant. We then averaged the standardized reaction times of each face in the three conditions across participants. We then plotted the three means of standardized reaction times against the three corresponding actual time to contact in a Cartesian coordinate plane, resulting in three points in the plane. A line through the two points at both ends was drawn. Finally, the slope for each face was calculated based on the averaged standardized reaction times of both estimated and actual time to contact. The slopes of the lines were used to indicate the overall threat-modulated underestimation for each face. Smaller slopes indicate that participants perceived the corresponding faces as more threatening, whereas larger slopes indicate that participants perceived the corresponding faces as less threatening.

To further filter the data, as in the first study, we used the established 10 face subgroups (bins) to capture the relation between the mean human likeness of the face bin and its corresponding slope computed using standardized reaction times in the looming task. We plotted the slopes against the mean ratings of human likeness of the 10 bins, resulting in a line graph (Figure 4) that partially resembled the uncanny valley.

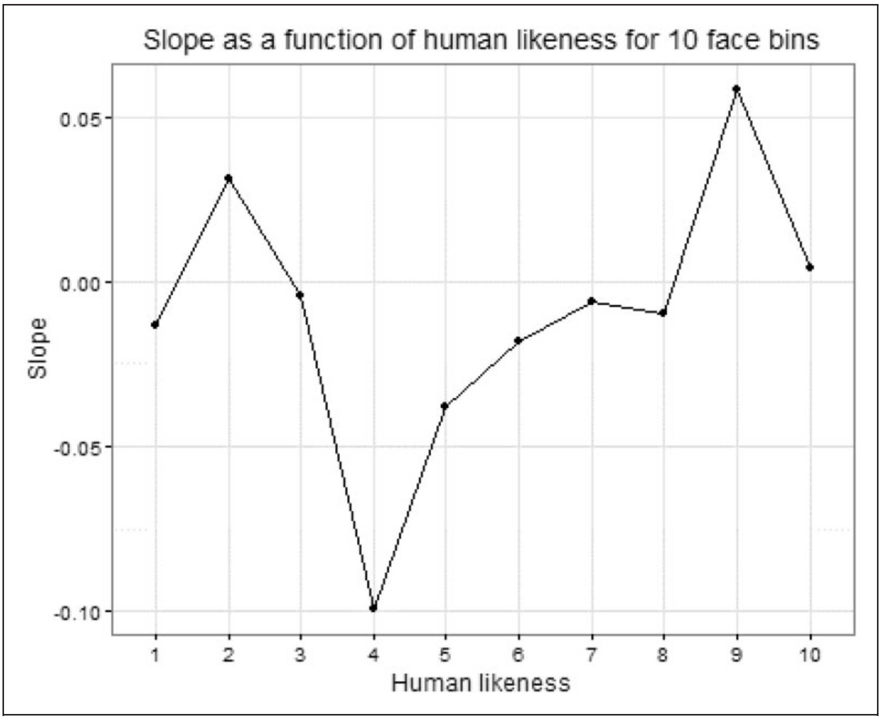

Figure 4. Slope of standardized (z score) estimated relative to actual time-to-contact as a function of ratings of human likeness divided by 10 face bins (see above).

Discussion

To date, an overwhelming majority of studies testing the uncanny valley hypothesis have used self-report questionnaires to measure participants' emotional responses in terms of the uncanny feeling, eeriness, or creepiness (Burleigh & Schoenherr, 2015; Ferrey, Burleigh, & Fenske, 2015; MacDorman & Ishiguro, 2006; Seyama & Nagayama, 2007). Although self-reported measures typically allow easier interpretation compared with behavioral or physiological ones, they lack external validity, which is partially attributable to the heterogeneity of self-report items (Kätsyri, Förger, Mäkäräinen, & Takala, 2015) and the fact that researchers rarely defined these items (e.g., eeriness) during the studies.

In Study 2, we introduced a visual looming task to implicitly measure participants' emotional responses to faces in relation to their human likeness (real rating of Study 1). A "slope" characteristic of the emotional modulating effect for each face (or group of faces) was obtained by plotting the estimated time-to-contact against the actual time-to-contact. One may argue that the participants might engage in spending more or less time ruminating about a stimulus rather than judging its time-to-contact. Nevertheless, we doubt that the rumination of images would be devoid of their emotional content in the visual looming task. Future research should further examine the cognitive underpinnings of the emotional modulating effect in the visual looming task to dissociate these alternative explanations.

By plotting the slopes against the ratings of human likeness for the 89 faces, divided in 10 bins, we detected a curve that resembled the uncanny valley. Nevertheless, the face bin (i.e., the fourth bin) that fell into the "valley" did not match the one (i.e., the fifth bin) detected in the first study using subjective measures. The inconsistent finding is attributable in part to the fact that explicit and implicit measures might capture different aspects of participants' subjective experiences (Gawronski, 2002). One explanation for the inconsistent findings may be that for self-reported measures, participants had longer time to evaluate the human likeness of the faces compared with the behavioral measures, in which participants had to make a quick response. We suggest that longer exposure time to the faces allowed participants to detect subtle cues that suggested nonhuman characteristics of the categorically ambiguous faces, thereby shifting the categorical boundary toward the human end of the dimension. Furthermore, because the current paradigm (i.e., the visual looming task) was driven by the mortality salience hypothesis, characterizing the uncanny feeling in terms of fear or perceived threat, the validity of this measurement tool relies on the feasibility of this hypothesis. As noted previously, the uncanny feeling may be a multifaceted affect that is not reducible to a single emotion (Ho et al., 2008). Likewise, the reaction times in the visual looming task may not exclusively or fully capture the "uncanny" feeling. Finally, it is unclear the extent to which emotional content other than perceived threateningness may moderate the estimated time-to-contact in the visual looming task. Therefore, the specificity of the visual looming task as an implicit measure of the uncanny feeling needs further scrutiny in future investigation.

To our knowledge, two studies used physiological or behavioral measures to examine the uncanny feeling (Strait & Scheutz, 2014; Złotowski et al., 2015) but generated divergent results. Strait and Scheutz (2014) presented 80 images of humans and robots to participants and asked them to rate human likeness and eeriness on 2 three-item surveys. They then re-presented the 80 images and used functional near infrared spectroscopy to assess participants' hemodynamic activity in the anterior medial prefrontal cortices, a region associated with emotion regulation. The subjective measures of affective responses indicated an inverted U-shape relation between eeriness and human likeness, corroborating the uncanny valley hypothesis. The implicit measure, indexed by hemodynamic activity in the prefrontal cortices, also captured the subjective feeling of eeriness that matched the uncanny valley. Nevertheless, using a brief implicit association test (bIAT) paradigm to implicitly measure eeriness, Złotowski et al. (2015) failed to replicate the findings obtained by using explicit measures. During the experiment, participants interacted with either an android (Geminoid HI-2) or a mechanical-looking robot ("*Robo*") for three times as part of an interview. After each interaction, they rated the eeriness of the interview robot and performed a bIAT. During the bIAT, participants classified a series of words as belonging to either one of the "Robo" and "Eeriness" categories (by pressing "K" key) versus neither category (by pressing "D" key). The bIAT paradigm predicts that participants should react faster if the two categories are strongly associated compared with a pair of unrelated categories. The results showed that for both interview robots, repeated interactions significantly lowered the explicit ratings of eeriness but not its implicit measure as indexed by the bIAT reaction times. These inconsistent findings suggest that the implicit association test paradigm may not be a valid measurement tool for capturing the uncanny feeling. With our second study that tries to come closer to an objective measure of the uncanny feeling, we are able to show that with our novel implicit (visual looming) task, the uncanny valley hypothesis is supported. The implicit measures of the uncanny feeling, however, await further validation and inquiry into what might be specifically measured in the bIAT paradigm and in the visual looming task.

Study 3

From Study 1 and Study 2, we examined both the subjective ratings and behavioral measures of people's emotional responses to artificial and human faces and demonstrated that there is a nonlinear relation between ratings of human likeness and emotional responses that matches the uncanny valley. In the third study, we attempted to examine further what might be the cognitive mechanisms contributing to this relation. We tested the hypothesis that the uncanny feeling may be attributable to categorical uncertainty, in particular the uncertainty of whether an entity is either animate or inanimate (Jentsch, 1906/1997; Ramey, 2006). Previous studies have generated mixed results regarding whether categorical uncertainty accounts for the uncanny feeling associated with human and artificial faces (e.g., avatars). Using morph images of human and avatar faces, some studies show that the most categorically ambiguous faces were also those rated as the least likable or most eerie (Burleigh et al., 2013; de Borst & de Gelder, 2015; Ferrey et al., 2015; Yamada et al., 2013), whereas other studies suggest that ambiguous faces do not elicit stronger negative emotions (Cheetham, Wu, Pauli, & Jäncke, 2015). A recent review on the categorization difficulty hypothesis (Kätsyri et al., 2015) questioned the image morph procedure, arguing that morphing artifacts (e.g., "ghosting" and "color interpolation," see Kätsyri et al., 2015, p. 8 for a discussion) might lower the validity of these studies. Factoring this criticism, Ferrey et al. (2015) used bistable images of animal images (e.g., a duck and an elephant) and human–avatar morphs to reproduce the uncanny valley. The results showed that the uncanny valley occurred at all morphed spectra and the negative affect was stronger at the midpoint compared with the endpoints of the spectrum. The researchers concluded that the uncanny valley is not limited to human likeness; instead, it occurs also for nonhuman stimuli. They also argued that that the valley is located not always close to the human endpoint, but around the midpoint that affords the highest degrees of competing interpretations of the stimuli during recognition. On the one hand, the researchers defined participants' self-reported negative affects as the uncanny feeling based on detecting the uncanny valley; on the other hand, the shape of the uncanny valley was arbitrarily determined by the researchers. Therefore, the findings that stimulus-inhibition accounts for the uncanny phenomenon are vulnerable to confirmation bias.

More importantly, we argue that the negative affects elicited by stimuli of nonhuman categories in this study bear questionable relevance to the uncanny feeling elicited by human replicas that are ambiguously human. This is particularly valid given that there is no operational definition regarding what types of negative affect capture the uncanny experience. In our view, the uncanny phenomenon associated with human replicas is defined and embedded in animacy perception and is not about any category of stimuli that would generate an "uncanny valley"-shape curve. After all, we agree with Ferrey et al. (2015) that the precise shape of the uncanny valley graph depicted by Mori (1970/2005) should not be taken literally as the criteria for detecting the uncanny phenomenon.

In the present study, we tested the categorical uncertainty hypothesis (see introduction) with real-world human and artificial faces using a novel reaction time-based sorting task. In particular, we probed the various degrees of categorical uncertainty associated with the 89 faces used in the preceding two studies (see Figure 2). This paradigm is similar to the two-alternative forced-choice classification task (Cheetham, Suter, & Jäncke, 2011, 2014; Cheetham et al., 2015) for establishing categorical ambiguity among morphed images. In the present study, participants sorted one by one and as quickly as possible, each of the 89 faces (see Figure 2) as either real or unreal. For each face, we systematically recorded their reaction times in sorting as a measurement of categorical uncertainty. Following the categorical uncertainty hypothesis, we predicted that the faces sorted as neither real nor unreal (e.g., the faces in-between, indexed by the percentage of agreement among participants on sorting a face as "real") would elicit the longest average reaction time. Furthermore, and more to the point here, if categorical uncertainty has anything to do with the uncanny feeling, longer sorting reaction times should be associated with faces that had the strongest negative emotional ratings in Study 1 (i.e., associated uncanny feeling).

Method

Participants. Thirty-six undergraduate students (20 females, $M_{age} = 20.36$ years, $SD = 2.29$) from Emory University participated in the sorting task, among whom 11 participated in the visual looming task first.

Stimuli and procedure. In the sorting task, participants sat in front of a 19-in. monitor (75-Hz refresh rate), where the 89 faces were presented one at a time in a random order at the center of the screen. Each trial began with a fixation cross at the center of the screen for 800 ms and then followed by a face. The task was to decide whether the face is real or unreal as quickly as possible by pressing one of two buttons: A green button on the left indicated "real" and a red button on the right indicated "unreal." At the moment the participant pressed the button, the face disappeared and was followed by a blank screen for 500 ms before the next trial began. The reaction time between the onset of the face and the response of the participant was recorded as a measure of uncertainty.

Results

Eleven of 36 participants participated in Study 2 before participating in Study 3. To test whether Study 2 interfered with Study 3, we compared the data from participants who only participated in Study 3 ($N = 25$) to those who were also tested in Study 2 ($N = 11$). We found a high correlation between the two groups of participants, both in mean reaction time, $N = 89$, $r(87) = .61$, $p < .01$, two tailed, 95% CI [0.46, 0.73] and in the percentage of participants that sorted a face as real, $N = 89$, $r(87) = .95$, $p < .01$, two tailed, 95% CI [0.93, 0.97]. Because the two samples generated similar results, we combined data from both samples in further analyses.

Of the 89 faces, 18 faces were consensually sorted as real and 7 faces as unreal. In other words, all of the 36 participants sorted the 18 faces as real and the 7 faces as unreal. Although the sorting of the other faces was inconsistent, a binomial test indicated that if equal or more than 23 of 36 participants sorted a face as real (or unreal), the agreement reached a significant level ($p < .05$). We called the faces whose agreement of sorting neither reached a significant level for "real" nor for "unreal" judgments "in-between" faces. We therefore categorized the 89 faces into three groups: 21 faces were sorted as significantly real, 59 faces were sorted as significantly unreal, and 9 faces were sorted as in between. A one-way repeated measures ANOVA yielded a significant difference in reaction time in sorting the three groups of faces, $F(2, 70) = 9.34$, $p < .05$, partial $\eta^2 = .21$. Pairwise comparisons showed that in-between faces ($M = 1401$; $SD = 111$) elicited significantly longer reaction times than the real ($M = 1157$; $SD = 76$; $t_{35} = 3.06$, $p < .01$) and the unreal faces ($M = 1048$; $SD = 67$, $t_{35} = 3.57$, $p < .01$), whereas the reaction times did not differ between the real and the unreal faces ($t_{35} = 1.57$, $p > .37$).

Using the z scores of ratings of emotional responses from Study 1 as a dependent measure, we then conducted one-way repeated measures ANOVAs to test whether, based on the sorting task, the in-between faces were also those that were rated higher on negative and lower on positive emotional ratings in Study 1. Results showed that, as predicted, the in-between faces of the sorting task were associated with significantly higher ratings of eeriness, disgustingness, unsettlingness, and threateningness, as well as significantly lower ratings of attractiveness and likableness compared with those categorized as significantly real or unreal.

Figure 5(a) to (c) provides graphic representations of the mean z scores of the emotional ratings in Study 1 across all three categories of faces—the real, the in-between, and the unreal. Mean z scores and confidence intervals as well as significant findings based on the statistical analysis are also reported in Table A3.

Eeriness. A one-way repeated measures ANOVA yielded a significant difference in eeriness ratings among the three groups of faces (real, in-between, and unreal)—$F(2, 122) = 53.20$, $p < .01$, partial $\eta^2 = .466$. Pairwise comparisons showed that in-between faces ($M = 0.30$; $SD = 0.26$) elicited significantly higher ratings of eeriness than the real ($M = -0.10$; $SD = 0.08$; $t_{61} = 10.11$, $p < .01$) and the unreal faces ($M = 0.15$; $SD = 0.20$; $t_{61} = 3.68$, $p < .01$). The ratings for the unreal faces were significantly higher than the real faces ($t_{61} = 6.92$, $p < .01$).

Disgustingness. A one-way repeated measures ANOVA revealed similar results for ratings of disgustingness—$F(2, 122) = 93.50$, $p < .01$, partial $\eta^2 = .61$. Pairwise comparisons showed that in-between faces ($M = 0.40$; $SD = 0.25$) elicited significantly higher ratings of disgustingness than the real ($M = -0.09$; $SD = 0.08$; $t_{61} = 12.31$, $p < .01$) and the unreal faces ($M = 0.09$; $SD = 0.18$; $t_{61} = 8.63$, $p < .01$). The ratings for the unreal faces were significantly higher than the real faces ($t_{61} = 5.57$, $p < .01$).

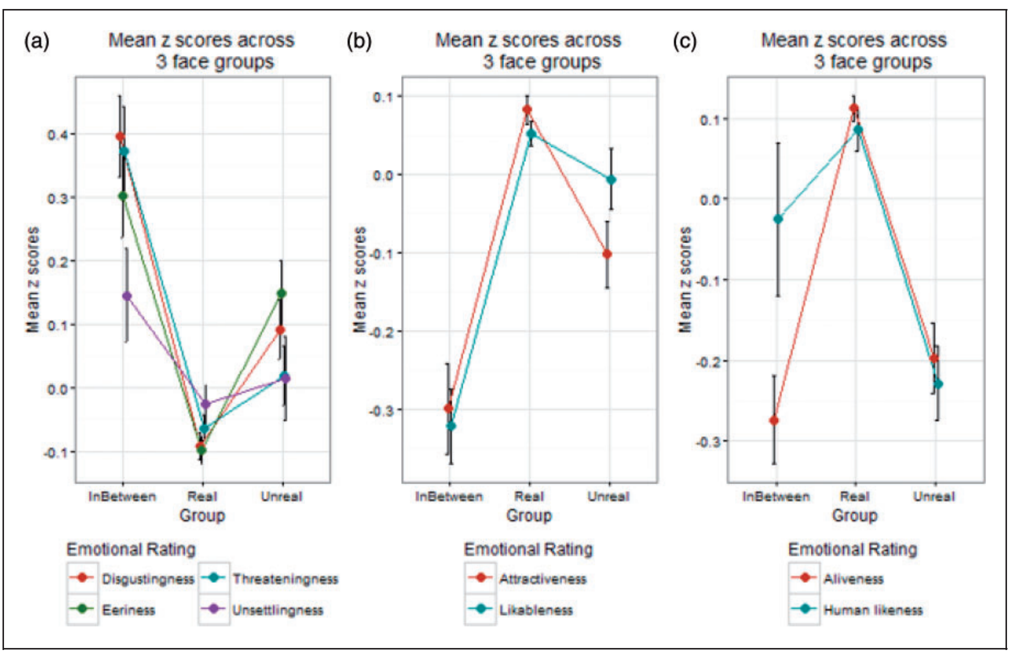

Figure 5. Emotional ratings (Study 1) as a function of face groups (real, in between, and unreal) established based on the sorting task (Study 3). (a) The ratings of negative emotions, including disgustingness, eeriness, threateningness, and unsettlingness; (b) the ratings of positive emotions, including attractiveness and likableness; and (c) the ratings of aliveness and human likeness, across all three face categories. Error bars represent standard errors.

Unsettlingness. A one-way repeated measures ANOVA revealed similar results for ratings of unsettlingness—$F(2, 122) = 8.20$, $p < .01$, partial $\eta^2 = .119$. Pairwise comparisons showed that in-between faces ($M = 0.15$; $SD = 0.29$) elicited significantly higher ratings of unsettlingness than the real ($M = -0.03$; $SD = 0.12$; $t_{61} = 3.50$, $p < .01$) and the unreal faces ($M = 0.01$; $SD = 0.26$; $t_{61} = 3.72$, $p < .01$)]. The ratings of unsettlingness for the unreal faces and the real faces did not differ.

Threateningness. A one-way repeated measures ANOVA revealed similar results for ratings of threateningness—$F(2, 122) = 70.80$, $p < .01$, partial $\eta^2 = .54$. Pairwise comparisons showed that in-between faces in-between faces ($M = 0.37$; $SD = 0.28$) elicited significantly higher ratings of threateningness than the real ($M = -0.06$; $SD = 0.08$; $t_{61} = 10.49$, $p < .01$) and the unreal faces ($M = 0.02$; $SD = 0.19$; $t_{61} = 8.49$, $p < .01$)]. The ratings of threateningness for the unreal faces and the real faces did not differ.

For ratings in relation to positive emotions, the results showed that the in-between faces elicited significantly lower ratings of attractiveness and likableness compared with the real and the unreal faces.

Attractiveness. A one-way repeated measures ANOVA yielded a significant difference in attractiveness ratings among the three groups of faces—$F(2, 122) = 66.27$, $p < .01$, partial $\eta^2 = .52$. Pairwise comparisons showed that in-between faces ($M = -0.30$; $SD = 0.23$) elicited significantly lower ratings of attractiveness than the real ($M = 0.08$; $SD = 0.07$; $t_{61} = -10.95$, $p < .01$) and the unreal faces ($M = -0.10$; $SD = 0.17$; $t_{61} = -5.69$, $p < .01$), and the unreal faces were significantly less attractive than the real faces ($t_{61} = -6.21$, $p < .01$).

Likableness. A one-way repeated measures ANOVA revealed similar results for likableness—$F(2, 122) = 98.00$, $p < .001$, partial $\eta^2 = .62$. Pairwise comparisons showed that in-between faces ($M = -0.32$; $SD = 0.19$) elicited significantly lower ratings of likableness than the real ($M = 0.05$; $SD = 0.06$; $t_{61} = -13.18$, $p < .01$) and the unreal ($M = -.01$; $SD = .15$; $t_{61} = -10.14$, $p < .01$) faces, whereas the unreal and the real faces did not differ in likableness ($t_{61} = 2.19$, $p > .09$).

Human likeness ("real" ratings). Furthermore, a one-way repeated measures ANOVA yielded a significant difference in human likeness ("real" ratings) among the three groups of faces established based on the sorting task—$F(2, 122) = 23.46$, $p < .01$, partial $\eta^2 = .28$. Pairwise comparisons showed that the real ($M = 0.09$; $SD = 0.11$; $t_{61} = 8.81$, $p < .01$) and the in-between faces ($M = -0.03$; $SD = 0.37$; $t_{61} = 4.91$, $p < .01$) both obtained significantly higher ratings of human likeness ("real" rating) than the unreal faces ($M = -0.23$; $SD = 0.18$), whereas the difference between the real and the in-between faces was not statistically significant ($t_{61} = 1.87$, $p > .20$). These findings showed that the in-between faces were rated as closer to real people than to nonhuman entities and suggest that this subtle difference between the real and the in-between faces may be crucial for inducing the categorical uncertainty and be associated with stronger negative emotional responses.

Aliveness. Finally, a one-way repeated measures ANOVA yielded a significant difference in aliveness ratings among the three groups of faces—$F(2, 122) = 77.69$, $p < .01$, partial $\eta^2 = .56$. Pairwise comparisons showed that the unreal ($M = -0.20$; $SD = 0.17$; $t_{61} = -10.75$, $p < .01$) and the in-between faces ($M = -0.27$; $SD = 0.21$; $t_{61} = -12.56$, $p < .01$) had significantly lower ratings of aliveness compared with the real faces ($M = 0.11$; $SD = 0.06$). In contrast, the difference between the unreal and the in-between faces did not reach statistical significance

($t_{61} = 1.99$, $p > .015$). These findings suggest that although the in-between faces appeared humanlike in terms of their similarity to real people in physical appearance, they were rated as inanimate, if not more inanimate than the unreal faces.

Discussion

The uncanny valley hypothesis (Mori, 1970/2005) predicted that human replicas closely resembling humans tend to elicit strong repulsion in observers. In Study 1 and Study 2, we plotted participants' emotional responses against perceived human likeness (real rating of Study 1), based on either explicit judgment (emotional ratings of Study 1) or implicit (reaction times of Study 2) measures. The results corroborated Mori's hypothesis yielding a nonlinear relation that resembled the uncanny valley. In Study 3, using a reaction time-based sorting task, we obtained data providing further information on the foundation to the uncanny valley. Consistent with previous findings that negative affects are associated with categorically ambiguous stimuli (Ferrey et al., 2015; Yamada et al., 2013), these data suggest that categorical uncertainty associated with discrimination of real and unreal faces might contribute to the uncanny feeling, providing further empirical support to what was proposed over a century ago by Jentsch (1906/1997). We found that the in-between faces that elicited a sense of categorical uncertainty as indexed by longer reaction times in the sorting task were also those rated more negatively in the emotional ratings of Study 1. By linking categorical uncertainty using a sorting task to ratings of emotional responses, our results supported the idea that difficulties in deciding whether an entity belongs either to the animate or inanimate category contribute to eliciting the uncanny feeling toward faces of human replicas that closely resemble real humans.

One might argue that the present findings are attributable to perceptual mismatch (e.g., Different facial features possess different degrees of animacy) as suggested by MacDorman et al. (2005). Nevertheless, we argue that perceptual mismatch alone is unlikely to account for the results. In particular, all kinds of human replica, including robots, androids, mannequins, dolls, and computer-generated characters, possess more or less degrees of perceptual mismatch. Otherwise, they would be indistinguishable from real people for human observers. Importantly, not all human replicas elicit uncanny feelings as some lifelike androids do (Bartneck et al., 2007). Taking the example of characters in animations, for example, their physical appearance often does not match their body movements and facial expressions on the level of human likeness. Nevertheless, they are appealing to most people. We argue that although perceptual mismatch may be a necessary condition for the uncanny feeling to emerge, it is insufficient in explaining the uncanny feeling without considering the categorical uncertainty associated with animacy perception.

The findings in Study 3 are broadly in line with the dehumanization hypothesis, which suggests that human replicas' lack of humanness may lie at the core of the uncanny feeling they elicit (see Wang et al., 2015 for further discussion). The dehumanization hypothesis gains support from two lines of research: the theory of anthropomorphism (e.g., Guthrie, 1993), which explains humans' propensity to attribute human characteristics to the nonhuman as stemming from a perceptual strategy to search for and often tend to overestimate significance, and the two-stage model of face processing (e.g., Looser et al., 2013), which shows that humans first perceive a coarse face form and then perceive animacy in the face. In particular, our results show that the in-between faces were characterized by both an ambiguous sense of animacy as indexed by longer reaction times during the sorting task and perceived lack of "life" (i.e., animacy) as indexed by lower ratings of aliveness compared with the real and the unreal faces. We suggest that that the in-between faces may first be anthropomorphized, and then ''dehumanized'' upon recognizing their inanimate nature, thereby eliciting the uncanny feeling. Noteworthy is also the fact that the categorical uncertainty hypothesis and the dehumanization hypothesis are not mutually exclusive in interpreting the present findings, given that both of them are closely linked to the perception of animacy. Given the correlational nature of the study, future research should systematically manipulate the degree of categorical uncertainty to examine the link between categorical uncertainty and the uncanny feeling to draw causal inferences.

General Discussion

The perception of animacy plays a critical role in humans' social perception and navigation of the social world, particularly in relation to face perception (Koldewyn, Hanus, & Balas, 2014; Looser & Wheatley, 2010). With modern technologies rendering the faces of computer-generated characters and robots increasingly lifelike, human observers, instead of liking them, tend to be repulsed by their high degrees of human likeness, ending up experiencing an uncanny feeling. The existence of such uncanny feeling might index the fact that humans are particularly tuned to detecting face animacy. This led us to adopt the uncanny phenomenon as a vehicle to enrich our understanding of human animacy perception in general.

According to the uncanny valley hypothesis (Mori, 1970/2005), the uncanny feeling might derive from the high level of human likeness of human replicas. Decades of research on this hypothesis, however, has generated inconsistent findings regarding whether there exists a nonlinear relation between emotional responses and the human likeness of human replicas. As it stands, it is yet unclear whether the perception of animacy might indeed be linked to human likeness, itself supposedly eliciting the uncanny feeling.

In the existing literature, different hypotheses exist trying to interpret the uncanny feeling as deriving from either evolved mechanisms of pathogen avoidance, mortality salience, and mate selection, or categorical uncertainty associated with face animacy perception. Researchers disagree on which hypothesis might be the best candidate to explain the uncanny feeling. In the present research, we had three aims: (a) validating the uncanny valley hypothesis with new experimental paradigms, (b) probing the link between threat detection and animacy detection, (c) validating the categorical uncertainty hypothesis of the uncanny feeling to shed a new light on animacy perception. In other words, we aimed at testing the role of human likeness in the perception of animacy using the uncanny feeling experience as a proxy, as well as probing the degree to which categorical uncertainty associated with animacy perception might be related to such feeling.

In three studies, we first corroborated the uncanny valley hypothesis and then supported both threat detection and categorical uncertainty as mechanisms potentially contributing to the uncanny feeling and the uncanny phenomenon. Results of Study 1 show that the relation between participants' emotional ratings and ratings of the human likeness of faces followed a function matching the uncanny valley, as originally proposed by Mori (1970/2005). In Study 2, using a novel visual looming task to measure objectively the human likeness of artificial and human faces, results show that plotting participants' estimated, relative to actual time-to-contact (an implicit, presumably fearful or threat avoidant response) against their ratings of human likeness yielded features once again matching the uncanny valley. Finally, in Study 3, we corroborate the hypothesis that negative emotional responses (potentially the so-called ''uncanny feeling'') to the 89 faces is linked to categorical uncertainty regarding the relative animacy of these faces when participants are asked to sort the faces as quickly as possible as either real or unreal. We found that faces with least consensus on their real versus unreal status elicited significantly longer reaction times in the sorting task, therefore indicating higher categorical uncertainty. More importantly, because these faces were also those rated more negatively in the survey of Study 1, we interpret these results as evidence that categorical uncertainty is linked to the uncanny feeling.

In all, these findings support Mori (1970/2005)'s idea that human likeness plays an important role in determining the emotional reactions human replicas elicit in human observers. There is indeed a nonlinear relation between human likeness judgments and the emotional responses of eeriness, disgustingness, unsettlingness, attractiveness, threateningness, and likableness, all contributing to the uncanny feeling (Study 1). In addition, Study 2 suggests that perceived threat as measured by estimated relative to actual time-to-contact in the visual looming task contributes to animacy detection and the associated uncanny phenomenon. Finally, Study 3 also demonstrates the role of categorical uncertainty in face animacy detection and the associated uncanny phenomenon. In short, our results suggest that the uncanny phenomenon is linked to the detection of a boundary between animate and inanimate face categories, as well as a sensitivity to the potential threat of entities crossing such boundary.

Because we used diverse experimental paradigms to capture the uncanny feeling as proxy of human animacy detection, questions remain whether our dependent variables are actually measuring the same thing, that is, the same presumed uncanny feeling experience of our participants. This point is particularly valid in light of the fact that there were no strict correspondences across studies regarding which of the faces were the most "uncanny." Furthermore, criteria in the determination of the uncanny valley were arbitrary, a limitation that is not unique to our research but of all those that preceded it, which do not predict the exact shape of the theoretical curve originally plotted by Mori (1970/2005). We did not commit to a precise shape of the uncanny valley, simply highlighting the fact that all the graphs displayed a nonmonotonic trend when plotting emotional responses as a function of human likeness of the faces. Future research and interpretation of face animacy detection in relation to the uncanny phenomenon would benefit from making precise predictions regarding the exact shape of the uncanny valley.

In general, the reported three studies built on and provide further empirical support to existing hypotheses concerning the nature and causes of the uncanny feeling, in particular the mortality salience hypothesis (Study 2) and the categorical uncertainty hypothesis (Study 3). These hypotheses, however, await further rigorous testing to ascertain valid measurements of the uncanny feeling.

Note that in the present studies, like in the majority of previous experimental investigations, static facial images were used to test the uncanny valley hypothesis in its most basic form. Future research should also include full body human replicas including motion and voice features to enhance validity by including richer information typically processed in the natural social environment. This would allow for more precise insights on how humans respond and relate to artificial entities that may potentially come to life in a more ecologically valid environment.

Finally, aside from physical features, future research should explore participants' varying proclivities to anthropomorphize and dehumanize human replicas. Researchers have proposed that the uncanny feeling associated with animacy detection might be linked to mind perception or anthropomorphism, particularly the strong automatic propensity to attribute human experience (e.g., ability to sense and feel) to nonhuman characters (Gray & Wegner, 2012). Other researchers proposed that the uncanny feeling might also be linked to a dehumanization process (Angelucci, Bastioni, Graziani, & Rossi, 2014; Wang et al., 2015;

Złotowski, Proudfoot, & Bartneck, 2013). Along these lines, future research could examine individual differences in terms of variable predispositions toward anthropomorphism and dehumanization in perceiving human replicas (Waytz, Cacioppo, Epley, 2010). Such variations might be an important moderator of the individual differences in the sensitivity to the uncanny valley (MacDorman & Entezari, 2015). They could indeed predict degrees of the uncanny feeling toward most closely resembling human replicas. These issues await further empirical scrutiny, which would help in further understanding what determines human animacy detection in light of the associated uncanny phenomenon.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

References

Aja, A. (Director). (2006). *The hills have eyes* [film]. Los Angeles, CA: Fox Searchlight Pictures.

Angelucci, A., Bastioni, M., Graziani, P., & Rossi, M. G. (2014). A philosophical look at the uncanny valley. In J. Seibt, R. Hakli, & M. Nørskov (Eds.), *Sociable robots and the future of social relations: Proceedings of robo-philosophy 2014* (pp. 165–169). doi:10.3233/978-1-61499-480-0-165

Ball, W., & Tronick, E. (1971). Infant responses to impending collision: Optical and real. *Science*, *171*, 818–820. doi:10.1126/science.171.3973.818

Bartneck, C., Kanda, T., Ishiguro, H., & Hagita, N. (2007). Is the uncanny valley an uncanny cliff? In *Proceedings of the 16th IEEE international symposium on robot and human interactive communication,* RO-MAN 2007, pp. 368–373. doi:10.1109/ROMAN.2007.4415111

Bartneck, C., Kulić, D., Croft, E., & Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, *1*, 71–81. doi:10.1007/s12369-008-0001-3

Brendel, E., DeLucia, P. R., Hecht, H., Stacy, R. L., & Larsen, J. T. (2012). Threatening pictures induce shortened time-to-contact estimates. *Attention, Perception, & Psychophysics*, *74*, 979–987. doi:10.3758/s13414-012-0285-0

Brenton, H., Gillies, M., Ballin, D. & Chatting, D. (2005). *The uncanny valley: Does it exist?* Presented at the 19th British HCI Group Annual Conference: Workshop on Human-Animated Character Interaction, 2005.

Burleigh, T. J., & Schoenherr, J. R. (2015). A reappraisal of the uncanny valley: Categorical perception or frequency-based sensitization? *Frontiers in Psychology*, *5*, 1488. doi:10.3389/fpsyg.2014.01488

Burleigh, T. J., Schoenherr, J. R., & Lacroix, G. L. (2013). Does the uncanny valley exist? An empirical test of the relationship between eeriness and the human likeness of digitally created faces. *Computers in Human Behavior*, *29*, 759–771. doi:10.1016/j.chb.2012.11.021

Cheetham, M., Suter, P., & Jäncke, L. (2011). The human likeness dimension of the ‘‘uncanny valley hypothesis’’: Behavioral and functional MRI findings. *Frontiers in Human Neuroscience*, *5*, 126. doi:10.3389/fnhum.2011.00126

Cheetham, M., Suter, P., & Jäncke, L. (2014). Perceptual discrimination difficulty and familiarity in the uncanny valley: More like a ‘‘Happy Valley’’. *Frontiers in Psychology*, *5*, 1219. doi:10.3389/fpsyg.2014.01219

Cheetham, M., Wu, L., Pauli, P., & Jäncke, L. (2015). Arousal, valence, and the uncanny valley: Psychophysiological and self-report findings. *Frontiers in Psychology*, *6*, 981. doi:10.3389/fpsyg.2015.00981

Coker-Appiah, D. S., White, S. F., Clanton, R., Yang, J., Martin, A., & Blair, R. J. R. (2013). Looming animate and inanimate threats: The response of the amygdala and periaqueductal gray. *Social Neuroscience*, *8*, 621–630. doi:10.1080/17470919.2013.839480

Collet-Serra, J. (Director). (2005). *House of wax* [film]. Los Angeles, CA: Warner Bros.

Curtis, V., Aunger, R., & Rabie, T. (2004). Evidence that disgust evolved to protect from risk of disease. *Proceedings of the Royal Society of London: Biological Sciences*, *271*, S131–S133. doi:10.1098/rsbl.2003.0144

de Borst, A. W., & de Gelder, B. (2015). Is it the real deal? Perception of virtual characters versus humans: An affective cognitive neuroscience perspective. *Frontiers in Psychology*, *6*, 576. doi:10.3389/fpsyg.2015.00576

Ferrey, A. E., Burleigh, T. J., & Fenske, M. J. (2015). Stimulus-category competition, inhibition, and affective devaluation: A novel account of the uncanny valley. *Frontiers in Psychology*, *6*, 249. doi:10.3389/fpsyg.2015.00249

Freud, S. (1919). The uncanny. In J. Strachey, & A. Freud (Eds.), *The standard edition of the complete psychological works of Sigmund Freud* (pp. 217–256). London, England: Hogarth Press.

Gawronski, B. (2002). What does the implicit association test measure? A test of the convergent and discriminant validity of prejudice-related IATs. *Experimental Psychology*, *49*, 171–180. doi:10.1026/1618-3169.49.3.171

Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition*, *125*, 125–130. doi:10.1016/j.cognition.2012.06.007

Greenberg, J., Pyszczynski, T., Solomon, S., Simon, L., & Breus, M. (1994). Role of consciousness and accessibility of death-related thoughts in mortality salience effects. *Journal of Personality and Social Psychology*, *67*, 627–637. doi:10.1037/0022-3514.67.4.627

Groom, V., Nass, C., Chen, T., Nielsen, A., Scarborough, J. K., & Robles, E. (2009). Evaluating the effects of behavioral realism in embodied agents. *International Journal of Human–Computer Studies*, *67*, 842–849. doi:10.1016/j.ijhcs.2009.07.001

Guthrie, S. (1993). *Faces in the clouds*. New York, NY: Oxford University Press.

Hanson, D. (2005, December). *Expanding the aesthetic possibilities for humanoid robots*. Presented at IEEE-RAS International Conference on Humanoid Robots, Tsukuba, Japan.

Ho, C.-C., & MacDorman, K. F. (2010). Revisiting the uncanny valley theory: Developing and validating an alternative to the Godspeed indices. *Computers in Human Behavior*, *26*, 1508–1518. doi:10.1016/j.chb.2010.05.015

Ho, C.-C., MacDorman, K. F., & Pramono, Z. A. (2008). Human emotion and the uncanny valley: A GLM, MDS, and Isomap analysis of robot video ratings. *In Proceedings of the 3rd ACM/IEEE international conference on human robot interaction, HRI '08* (pp. 169–176). doi:10.1145/1349822.1349845

Hoffmann, E. T. A. (1817/1994). *The Sandman* (19th Century German Stories; J. Oxenford, Trans.; R. Godwin-Jones, Ed.). Richmond: Virginia Commonwealth University.

Jentsch, E. (1906/1997). On the psychology of the uncanny (R. Sellars, Trans.). *Angelaki, 2*, 7–16. doi:10.1080/09697259708571910

Kätsyri, J., Förger, K., Mäkäräinen, M., & Takala, T. (2015). A review of empirical evidence on different uncanny valley hypotheses: Support for perceptual mismatch as one road to the valley of eeriness. *Frontiers in Psychology*, *6*, 390. doi:10.3389/fpsyg.2015.00390

Koldewyn, K1., Hanus, P., & Bala,s B. (2014). Visual adaptation of the perception of "life": animacy is a basic perceptual dimension of faces. *Psychonomic bulletin & review*, *21*, 969–975. doi: 10.3758/s13423-013-0562-5

Looser, C. E., Guntupalli, J. S., & Wheatley, T. (2013). Multivoxel patterns in face-sensitive temporal regions reveal an encoding schema based on detecting life in a face. *Social Cognitive and Affective Neuroscience*, *8*, 799–805. doi:10.1093/scan/nss078

Looser, C. E., & Wheatley, T. (2010). The tipping point of animacy how, when, and where we perceive life in a face. *Psychological Science*, *21*, 1854–1862. doi:10.1177/0956797610388044

MacDorman, K. (2006). Subjective ratings of robot video clips for human likeness, familiarity, and eeriness: An exploration of the uncanny valley. In *ICCS/CogSci-2006 long symposium: Toward social mechanisms of android science* (pp. 26–29). Vancouver, BC, Canada.

MacDorman, K. F., & Entezari, S. O. (2015). Individual differences predict sensitivity to the uncanny valley. *Interaction Studies*, *16*, 141–172. doi:10.1075/is.16.2.01mac

MacDorman, K. F., Green, R., Ho, C.-C., & Koch, C. (2009). Too real for comfort? Uncanny responses to computer generated faces. *Computers in Human Behavior*, *25*, 695–710. doi:10.1016/j.chb.2008.12.026

MacDorman, K., & Ishiguro, H. (2006). The uncanny advantage of using androids in cognitive and social science research. *Interaction Studies*, *7*, 297–337.

MacDorman, K. F., Minato, T., Shimada, M., Itakura, S., Cowley, S. J., & Ishiguro, H. (2005). Assessing human likeness by eye contact in an android testbed. In *Proceedings of the XXVII Annual Meeting of the Cognitive Science Society* (pp. 1373–1378). Stresa, Italy, 21–23 July 2005.

Mangan, B. (2015). The uncanny valley as fringe experience. *Interaction Studies*, *16*, 193–199. doi:10.1075/is.16.2.05man

Mathur, M. B., & Reichling, D. B. (2016). Navigating a social world with robot partners: A quantitative cartography of the uncanny valley. *Cognition*, *146*, 22–32. doi:10.1016/j.cognition.2015.09.008

Mitchell, W. J., Szerszen, K. A. Sr., Lu, A. S., Schermerhorn, P. W., Scheutz, M., & MacDorman, K. F. (2011). A mismatch in the human realism of face and voice produces an uncanny valley. *i-Perception*, *2*, 10–12. doi:10.1068/i0415

Mori, M. (1970/2005). The uncanny valley (K. F. MacDorman & T. Minato, Trans.). *Energy*, *7*, 33–35.

Mori, M., MacDorman, K. F., & Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Transanctions on Robotics & Automation Magazine*, *19*, 98–100. doi:10.1109/MRA.2012.2192811

Poliakoff, E., Beach, N., Best, R., Howard, T., & Gowen, E. (2013). Can looking at a hand make your skin crawl? Peering into the uncanny valley for hands. *Perception*, *42*, 998–1000. doi:10.1068/p7569

Ramey, C. H. (2006, July). *An inventory of reported characteristics for home computers, robots, and human beings: Applications for android science and the uncanny valley*. Presented at the Proceedings of the ICCS/CogSci-2006 Long Symposium 'Toward Social Mechanisms of Android Science', Vancouver, Canada.

Rhodes, G., Zebrowitz, L. A., Clark, A., Kalick, S. M., Hightower, A., & McKay, R. (2001). Do facial averageness and symmetry signal health? *Evolution and Human Behavior*, *22*, 31–46. doi:10.1016/S1090-5138(00)00060-X

Rozin, P., & Fallon, A. E. (1987). A perspective on disgust. *Psychological Review*, *94*, 23–41. doi:10.1037/0033-295X.94.1.23

Saygin, A. P., Chaminade, T., Ishiguro, H., Driver, J., & Frith, C. (2012). The thing that should not be: Predictive coding and the uncanny valley in perceiving human and humanoid robot actions. *Social Cognitive and Affective Neuroscience*, *7*, 413–422. doi:10.1093/scan/nsr025

Schiff, W., Caviness, J. A., & Gibson, J. J. (1962). Persistent fear responses in rhesus monkeys to the optical stimulus of "looming". *Science*, *136*, 982–983. doi:10.1126/science.136.3520.982

Seyama, J., & Nagayama, R. S. (2007). The uncanny valley: Effect of human likeness on the impression of artificial human faces. *Presence: Teleoperators and Virtual Environments*, *16*, 337–351.

Strait, M., & Scheutz, M. (2014). Measuring users' responses to humans, robots, and human-like robots with functional near infrared spectroscopy. In *The 23rd IEEE international symposium on robot and human interactive communication*, pp.1128–1133. doi:10.1109/ROMAN.2014.6926403

Tinwell, A., Grimshaw, M., Nabi, D. A., & Williams, A. (2011). Facial expression of emotion and perception of the uncanny valley in virtual characters. *Computers in Human Behavior*, *27*, 741–749. doi:10.1016/j.chb.2010.10.018

Tinwell, A., Grimshaw, M. N., & Williams, A. (2010). Uncanny behavior in survival horror games. *Journal of Gaming and Virtual Worlds*, *2*, 3–25. doi:10.1386/jgvw.2.1.3_1

Vagnoni, E., Lourenco, S. F., & Longo, M. R. (2012). Threat modulates perception of looming visual stimuli. *Current Biology*, *22*, R826–R827. doi:10.1016/j.cub.2012.07.053

Wang, S., Lilienfeld, S. O., & Rochat, P. (2015). The uncanny valley: Existence and explanations. *Review of General Psychology*, *19*, 393–407. doi:10.1037/gpr0000056

Waytz, A., Cacioppo, J., & Epley, N. (2010). Who sees human? The stability and importance of individual differences in anthropomorphism. *Perspectives on Psychological Science*, *5*, 219–232. doi:10.1177/1745691610369336

Wheatley, T., Weinberg, A., Looser, C., Moran, T., & Hajcak, G. (2011). Mind perception: Real but not artificial faces sustain neural activity beyond the N170/VPP. *PLoS ONE*, *6*, e17960. doi:10.1371/journal.pone.0017960

Yamada, Y., Kawabe, T., & Ihaya, K. (2013). Categorization difficulty is associated with negative evaluation in the "uncanny valley" phenomenon. *Japanese Psychological Research*, *55*, 20–32. doi:10.1111/j.1468-5884.2012.00538.x

Zemeckis, R. (Director). (2004). *The polar express* [film]. Los Angeles, CA: Warner Bros.

Złotowski, J., Proudfoot, D., & Bartneck, C. (2013). More human than human: Does the uncanny curve really matter? In *Proceedings of the HRI2013 workshop on design of humanlikeness in HRI from uncanny valley to minimal design*, Tokyo, pp. 713. .

Złotowski, J. A., Sumioka, H., Nishio, S., Glas, D. F., Bartneck, C., & Ishiguro, H. (2015). Persistence of the uncanny valley: The influence of repeated interactions and a robot's attitude on its perception. *Frontiers in Psychology*, *6*, 883. doi:10.3389/fpsyg.2015.00883.

Appendices

Table A1. Eight statement survey questionnaire used for rating of faces in Experiment 1: Participants that they were going to see some faces and instructed them to rate each face on eight statements on a 7-point Likert scale (1 ¼1/4Not at all and 7 ¼Very much,).

Slide						
This figure makes me feel eerie?						
1	2	3	4	5	6	7
Not at all						*Very much*
This figure makes me feel disgusted?						
1	2	3	4	5	6	7
Not at all						*Very much*
This figure makes me feel unsettling?						
1	2	3	4	5	6	7
Not at all						*Very much*
This figure looks attractive?						
1	2	3	4	5	6	7
Not at all						*Very much*
This figure looks threatening?						
1	2	3	4	5	6	7
Not at all						*Very much*
This figure looks likable?						
1	2	3	4	5	6	7
Not at all						*Very much*
This figure looks real?						
1	2	3	4	5	6	7
Not at all						*Very much*
This figure looks alive?						
1	2	3	4	5	6	7
Not at all						*Very much*

Table A2. Descriptive Statistics of Reaction Times on the Six Presentation Conditions.

Presentation condition		Estimated time-to-contact	
Initial size	Actual time-to-contact	*M*	*SD*
400 pixel	2500 ms	3004 ms	1022 ms
	3500 ms	3541 ms	1332 ms
	4500 ms	3986 ms	1757 ms
500 pixel	2500 ms	2809 ms	901 ms
	3500 ms	3375 ms	1539 ms
	4500 ms	3636 ms	1529 ms

Table A3. Mean z Scores of Emotional Ratings Across Real, In-Between, and Unreal Face Categories.

	Face category					
	Real		In-between		Unreal	
Emotional rating	*M, SD*	95% CI	*M, SD*	95% CI	*M, SD*	95% CI
Eeriness	−0.10, 0.08	[−0.12, −0.08]	0.30, 0.26**	[0.24, 0.37]	0.15, 0.20	[0.10, 0.20]
Disgustingness	−0.09, 0.08	[−0.11, −0.07]	0.40, 0.25*	[0.33, 0.46]	0.09, 0.18	[0.05, 0.14]
Unsettlingness	−0.03, 0.12	[−0.06, 0.00]	0.15, 0.29*	[0.07, 0.22]	0.01, 0.26	[−0.05, 0.08]
Threateningness	−0.06, 0.08	[−0.08, −0.04]	0.37, 0.28*	[0.30, 0.44]	0.02, 0.19	[−0.03, 0.07]
Attractiveness	0.08, 0.07	[0.06, 0.10]	−0.30, 0.23**	[−0.36, −0.24]	−0.10, 0.17	[−0.15, −0.06]
Likableness	0.05, 0.06	[0.04, 0.07]	−0.32, 0.19**	[−0.37, −0.27]	−0.01, 0.15	[−0.05, 0.03]
Human likeness	0.09, 0.11	[0.06, 0.11]	−0.03, 0.37[a]	[−0.12, 0.07]	−0.23, 0.18	[−0.28, −0.18]
Aliveness	0.11, 0.06	[0.10, 0.13]	−0.27, 0.21[b]	[−0.33, −−0.22]	−0.20, 0.17	[−0.24, −0.15]

Note. Asterisks denote significant differences in mean z scores of emotional ratings between the in-between and the other two face categories (i.e., real and unreal). ***p* < .001. **p* < .016.

[a]The in-between faces were significantly more human like ("real" rating) than the unreal faces, whereas they were as humanlike as the real faces.

[b]The in-between faces obtained significantly lower ratings of aliveness compared with the real faces, whereas their ratings of aliveness did not significantly differ from those of the unreal faces.