Article

PERCEPTION

*Perception*
2020, Vol. 49(10) 1069–1089

DOI: 10.1177/0301006620952611
journals.sagepub.com/home/pec
SAGE

# The Uncanny Valley Phenomenon and the Temporal Dynamics of Face Animacy Perception

**Shensheng Wang, Yuk F. Cheong, Daniel D. Dilks and Philippe Rochat**
Emory University, United States

**Abstract**
Human replicas highly resembling people tend to elicit eerie sensations—a phenomenon known as the uncanny valley. To test whether this effect is attributable to people's ascription of mind to (i.e., *mind perception hypothesis*) or subtraction of mind from androids (i.e., *dehumanization hypothesis*), in Study 1, we examined the effect of face exposure time on the perceived animacy of human, android, and mechanical-looking robot faces. In Study 2, in addition to exposure time, we also manipulated the spatial frequency of faces, by preserving either their fine (high spatial frequency) or coarse (low spatial frequency) information, to examine its effect on faces' perceived animacy and uncanniness. We found that perceived animacy decreased as a function of exposure time only in android but not in human or mechanical-looking robot faces (Study 1). In addition, the manipulation of spatial frequency eliminated the decrease in android faces' perceived animacy and reduced their perceived uncanniness (Study 2). These findings link perceived uncanniness in androids to the temporal dynamics of face animacy perception. We discuss these findings in relation to the dehumanization hypothesis and alternative hypotheses of the uncanny valley phenomenon.

**Keywords**
uncanny valley, animacy, face perception, mind perception, dehumanization, temporal dynamics

Date Received: 23 November 2019; accepted: 4 August 2020

As humans, we are prone to attribute human characteristics, such as a mind, to objects, particularly those that bear a physical resemblance to us (K. Gray & Wegner, 2012; Guthrie, 1993; Martini et al., 2016). This tendency is broadly known as anthropomorphism (Epley & Waytz, 2010).

**Corresponding author:**
Philippe Rochat, Department of Psychology, Emory University, 36 Eagle Row, Atlanta, GA 30322, United States.
Email: psypr@emory.edu

For decades, scientists have capitalized on this human tendency to create robots with humanlike features to serve in healthcare, education, and various other domains (Broadbent, 2017; Damiano & Dumouchel, 2018). Although humanlike features, such as a face, prompt us to attribute minds to robots (DiSalvo et al., 2002; Fink, 2012), up to a certain point, increasingly humanlike features cease to render robots more human, but elicit feelings of unease, eeriness, and revulsion, creating the so-called uncanny valley phenomenon (Mori et al., 2012).

## *Mind Perception Hypothesis*

In the past decade and a half, researchers have debated over whether the uncanny valley exists (Bartneck et al., 2007; MacDorman & Ishiguro, 2006; Mathur & Reichling, 2016) and how to account for it (Wang et al., 2015). Although most accounts tend to focus on the physical properties of the human replicas (e.g., the inconsistency between humanlike eyes and a doll-like nose), K. Gray and Wegner (2012) focus on the cognitive processes of the human perceivers (see also Misselhorn, 2009). They argue that the uncanny valley is due to humans' ascription of minds to robots, which violates humans' expectations for machines lacking human capacities to feel and sense (H. M. Gray et al., 2007; *Mind Perception Hypothesis*). Nevertheless, the authors did not address why the ascription of mind would be associated with eerie feelings mostly in entities highly resembling humans, such as androids (Mathur & Reichling, 2016; Wang & Rochat, 2017), but not among pets, cartoon characters, or mechanical-looking robots (Misselhorn, 2009).

At the center of this controversy lies the issue regarding how perceptual cues allow us to perceive others as minded creatures (Deska & Hugenberg, 2017; Schroeder & Epley, 2016). Mind perception broadly entails both determining (a) whether others possess minds (*mind awareness*) and (b) what state of mind they might be in (*mental state awareness*; Epley & Waytz, 2010; Varga, 2017). Here, we focus on the *mind awareness* aspect of mind perception and investigate how faces signal the presence of minds (or the lack thereof), a process that cognitive psychologists and neuroscientists refer to as *face animacy perception* (Looser & Wheatley, 2010). In particular, we aim at demonstrating how the temporal dynamics of face animacy perception (Looser & Wheatley, 2010) might be inextricably linked to the uncanny valley in androids.

## *Dehumanization Hypothesis*

Philosophers have long argued that we perceive minds in others "through the senses" (see Varga, 2017, p. 788). Consistent with this proposition, research demonstrates distinct neural correlates in the brain involved in the perception of minds in faces: Only human faces elicit a sustained late positivity in the event-related potential beyond 400 ms following stimulus onset, whereas doll faces do not, despite both face types eliciting an early face-sensitive N170 component (Wheatley et al., 2011). These findings suggest that two stages of face perception—*face form* and *face animacy*—unfold in time. Although the first stage (before 400 ms after stimulus onset) allows perceivers to detect a face rapidly, the second stage (400 ms after stimulus onset) discerns whether the face possesses a mind (Looser et al., 2013).

Although the two-stage model seems to suggest a rigid dichotomy between *face* and *mind*, it does not rule out the possibility that humans would attribute minds to faces earlier than 400 ms. Rather, given that faces are a common signal for mind (Gao et al., 2010; Looser & Wheatley, 2010) and that face detection occurs as quickly as 100 ms (Crouzet et al., 2010), humans might attribute minds to faces as soon as they detect them. Importantly, this early attribution of mind might be characterized by an anthropomorphic tendency to overattribute mind to nonhuman agents (Guthrie, 1993). Therefore, Wang et al. (2015) propose the dehumanization hypothesis, arguing that the uncanny valley might be linked to a decrease in perceived animacy of android faces as a function of exposure time. In particular, the rapid face detection coupled with a natural tendency for anthropomorphism predicts an initial overattribution of minds to android faces, followed by a subtraction of minds from them when humans discriminate android from human faces beyond 400 ms after stimulus onset.

## *Pitting the Dehumanization Hypothesis Against the Mind Perception Hypothesis*

Although both the dehumanization and the mind perception hypotheses argue for the crucial role played by mind perception in eliciting the uncanny valley, the two hypotheses differ profoundly in their underlying assumptions, which lead to opposing predictions regarding how perceived face animacy would change as a function of exposure time. In particular, the dehumanization hypothesis assumes that people initially attribute high animacy to an android face, predicting a decrease in animacy over time following 400 ms after stimulus onset. In contrast, the mind perception hypothesis is based on the default assumption that an android face is inanimate; only with increasing exposure to the face do people attribute a mind to it, predicting an increase in animacy over time (Wang et al., 2015). Therefore, according to the *dehumanization hypothesis*, we predicted that perceived animacy should decrease in android faces. In contrast, we predicted that perceived animacy should remain stable in human and mechanical-looking robot faces.

## *Coarse Versus Fine Information in Face Animacy Perception*

By proposing the two-stage model of face processing, researchers assume that face form detection relies mainly on coarse information carried by *low spatial frequency* (LSF), whereas face animacy perception requires scrutinizing the fine details (Looser & Wheatley, 2010; Wheatley et al., 2011) carried by *high spatial frequency* (HSF; Peters et al., 2018).

Building upon this plausible assumption, we introduced a novel paradigm, in which faces are presented intact during the first 100 ms (form perception) but filtered to preserve either the HSF or LSF during later stages of processing (animacy perception; e.g., 100 ms–500 ms and 500 ms–1000 ms after stimulus onset). This design allows us to focus on the second stage of face processing most crucial for discerning face animacy (Wheatley et al., 2011) to examine how selectively interfering with this stage of processing might influence faces' perceived animacy. If android faces are naturally associated with a decrease in perceived animacy over time, according to the *dehumanization hypothesis,* this manipulation should predict a reduction in the decrease in perceived animacy, which allows us to examine its downstream effect of the perceived uncanniness of android faces. In particular, we predicted that removing either HSF or LSF after the first 100 ms of the presentation should mitigate the perceived uncanniness of android faces. In addition, we predicted that removing HSF should yield a more pronounced effect compared with the removal of LSF, given the assumption that HSF plays a dominant role in the perception of animacy during later stages of face processing.

## *Overview of Present Research*

The primary goal of the present research is to test the dehumanization hypothesis against the mind perception hypothesis. To do so, we focused on three types of faces varying in their level of uncanniness—human, android, and mechanical-looking robot[1] (Wang et al., 2020).

In Study 1, we systematically manipulated the exposure time of faces to examine how perceived face animacy might change as a function of exposure time, depending on face type. Following our dehumanization hypothesis, we predicted that only for android faces, perceived animacy would decrease over time (i.e., high-then-low animacy). In contrast, it would remain low in mechanical-looking robot faces and high in human faces. In Study 2, in addition to manipulating exposure time, we manipulated the spatial frequency of the face. Following our dehumanization hypothesis, we predicted that the manipulation of spatial frequency would both influence the temporal trajectory of face animacy perception and reduce the perceived uncanniness of android faces.

## Study 1

In Study 1, we systematically manipulated the exposure time of each type of face to examine its effect on participants' ratings of faces' perceived animacy.

### Methods

*Participants.* We recruited 58 undergraduate students ($M_{\text{age}} = 19.01$, $SD = 1.78$) from Emory University with normal or corrected-to-normal vision for our study. The sample size was determined based on previous research (Wang & Rochat, 2017). Participants received course credits and were debriefed after testing. All participants were naïve to the specifics of the experiment and the driving hypothesis.

*Materials.* The stimuli were 57 static, color, images of faces used in our previous research (Wang & Rochat, 2017; Wang et al., 2020) featuring humans ($n = 19$), androids ($n = 25$), and mechanical-looking robots ($n = 13$; see Appendix A). We cropped, resized, and converted these images to PNG format. The stimuli were presented on a 23-inch LCD computer screen, displayed on a white background. Each stimulus subtended a visual angle of approximately $5.4° \times 7.5°$.

*Procedure.* The task consisted of three blocks of trials (see Figure 1). The same 57 faces were presented one at a time and in a randomized order at 100 ms in the first, 500 ms in the second, and 1000 ms in the third block. We presented the three blocks in succession with increasing exposure time to control for the potential carryover effect of longer exposure time on short exposure time condition in the judgment of face animacy. Participants were told that faces would be presented briefly on the screen, and their task was to rate how alive each face looked. The term *alive* was chosen as a proxy of perceived animacy of the face following the procedure of Looser and Wheatley (2010).

Each trial started with a fixation cross (+) at the center of the screen for 500 ms. A face then appeared at the same location for the given exposure time. Immediately afterward, participants were prompted to rate the degree to which the face looked alive by using a 6-point Likert-type scale ranging from 1 (*not alive*) to 6 (*alive*) appearing on the screen with unlimited response time. Finally, participants provided demographic information on their age, gender, and ethnicity.

### Results and Discussion

In Study 1, we hypothesized that only for android faces, perceived animacy would decrease as a function of exposure time (i.e., high-then-low animacy). To test this hypothesis, we computed

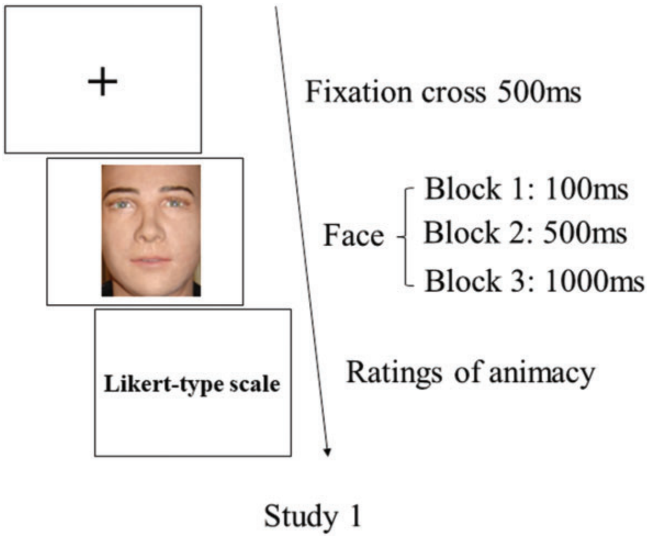


**Figure 1.** Schematic Representation of the Procedure in Study 1 in Which Faces Were Presented for 100 ms, 500 ms, or 1000 ms.
Following each presentation, participants rated their perceived animacy based on a 6-point Likert-type scale: 1 (*not alive*) to 6 (*alive*).

**Table 1.** Mean (Standard Deviation in Parentheses) Values of Perceived Animacy for Each Level of Exposure Time by Each Face Type ($N = 58$).

| | Face type | | |
|---|---|---|---|
| Exposure time (ms) | Human | Android | Mechanical |
| 100 | 5.07 (0.77) | 3.14 (0.59) | 1.13 (0.38) |
| 500 | 4.90 (0.82) | 2.44 (0.60) | 1.08 (0.29) |
| 1000 | 4.93 (0.81) | 2.37 (0.60) | 1.08 (0.25) |

the mean ratings of animacy for each of the three types of faces at each of the three exposure times (see Table 1 for descriptive statistics). Visual inspection of the normal quantile–quantile plots (hereafter referred to as normal Q–Q plots, see Appendix B) revealed that the distribution of mean ratings of animacy deviated from normality in human and mechanical-looking robot faces. Nevertheless, given that the sample was sufficiently large ($N = 58$), the sampling distribution of the mean would tend to be normally distributed. Therefore, the mean ratings of animacy were submitted to a 3 (face type [human, android, and mechanical-looking]) × 3 (exposure time [100 ms, 500 ms, and 1000 ms]) repeated-measures analysis of variance (ANOVA) using R's *afex* package (Singmann et al., 2018). For repeated-measures ANOVA with factors having more than two levels, when Mauchly's test indicated a violation of the assumption of sphericity, we employed Greenhouse–Geisser estimates of sphericity to correct the degrees of freedom. For each significant main effect or simple main effect, we conducted follow-up pairwise comparisons using a Bonferroni-adjusted alpha level of .017 per test (.05/3).

The ANOVA yielded a hypothesized significant interaction effect between face type and exposure time, $F(3.07, 174.82) = 26.65$, $p < .0001$, $\eta_p^2 = .32$. Corroborating our hypothesis, at an alpha level of .017, simple main effect analyses revealed a statistically significant difference between levels of exposure time in the ratings of animacy of android faces, $F(1.46, 83.36) = 55.31$, $p < .0001$, $\eta_p^2 = .49$, but not of mechanical-looking robot faces, $F(1.58, 89.98) = 3.63$, $p = .04$, $\eta_p^2 = .06$, or human faces, $F(1.56, 89.01) = 2.98$, $p = .07$, $\eta_p^2 = .05$. Follow-up pairwise comparisons revealed that for android faces, there were significant differences in ratings of animacy between 100 ms ($M = 3.14$, $SD = 0.59$) and 500 ms ($M = 2.44$, $SD = 0.60$), $t(57) = 8.44$, Cohen's $d = 1.11$, $p < .0001$, and between 100 ms ($M = 3.14$, $SD = 0.59$) and 1000 ms ($M = 2.37$, $SD = 0.60$), $t(57) = 7.76$, Cohen's $d = 1.02$, $p < .0001$, but there was no significant difference between 500 ms ($M = 2.44$, $SD = 0.60$) and 1000 ms ($M = 2.37$, $SD = 0.60$), $t(57) = 1.20$, Cohen's $d = .16$, $p = .23$. All $p$ values were two-tailed. The results of Study 1 are summarized in Figure 2.

To ensure that our results are not driven by particular images or participants, we additionally modeled item response data with crossed random effects for stimuli/images and participants (Locker et al., 2007) with face type and exposure time as predictors. The results showed that after controlling for the random stimulus and participant effects, the interaction between face type and exposure time remained statistically significant, $\chi^2$ (4) = 76.4, $p < .001$.

In summary, data of Study 1 demonstrate that perceived animacy decreases significantly as a function of exposure time for android but not for mechanical-looking robot or human faces. Importantly, in android faces, perceived animacy drops between 100 ms and 500 ms, consistent with the finding that humans begin to discern face animacy perception around 400 ms after stimulus onset (Wheatley et al., 2011).

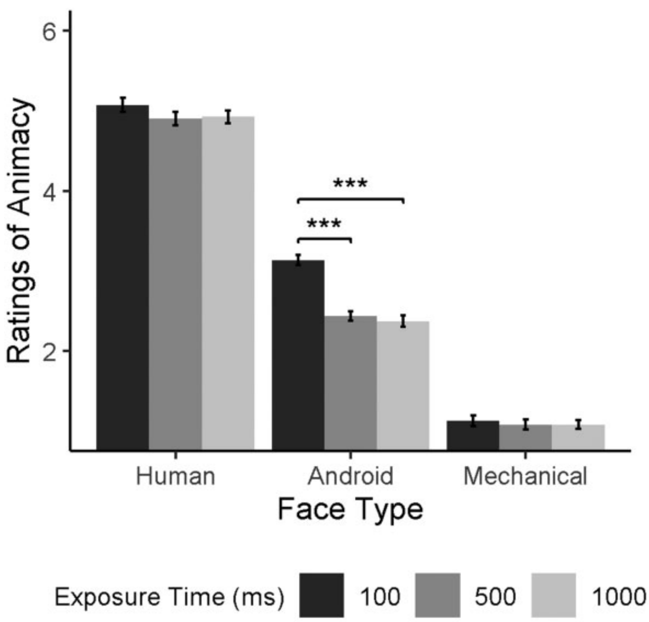


**Figure 2.** Mean Ratings of Animacy at Each Level of Exposure Time (100 ms, 500 ms, and 1000 ms) and for Each Face Type (Human, Android, and Mechanical-Looking).
***$p < .0001$.

## Study 2

In Study 2, we selectively preserved either the HSF (Study 2a) or LSF (Study 2b) of the faces during later stages of presentation (e.g., 100 ms–500 ms and 500 ms–1000 ms after stimulus onset) to examine their effects on changes in perceived animacy as a function of exposure time. In particular, based on our earlier discussion, we predicted that removing either HSF or LSF after the first 100 ms of the presentation should mitigate the eerie feelings by potentially reducing the magnitude of the decline in android faces' perceived animacy (see Study 1). In addition, if HSF plays a more significant role in face animacy perception than does LSF during later stages of face processing, we predicted that removing HSF (Study 2b) should create a similar, yet more pronounced effect compared with the removal of LSF (Study 2a).

### *Methods*

*Participants.* We conducted a power analysis using G*Power 3.1 (Faul et al., 2007) for a 2 (condition [BSF vs. BSF + HSF]) × 3 (exposure time [100 ms, 500 ms, and 1000 ms]) interaction, assuming a small effect ($\eta_p^2 = .05$), a type I error rate of .05, and a .8 correlation between measures. This suggested a target $N$ of 22 for 95% power in each study. Twenty-two undergraduate students ($M_{\text{age}} = 18.95$, $SD = 1.05$) participated in Study 2a, and another 22 ($M_{\text{age}} = 18.91$, $SD = 1.02$) participated in Study 2b.

*Materials.* The same 57 static face images of Study 1 were used but passed through a spatial frequency filter to preserve either HSF (Study 2a) or LSF (Study 2b). We used a cutoff frequency of 8 cycles per face width for LSF faces and a cutoff frequency of 32 cycles per face width for HSF faces based on a previous study (Goffaux et al., 2005).

The original images containing the broad spectrum of spatial frequency (BSF) and the filtered images (HSF and LSF) were equated on mean luminance using MATLAB (The Mathworks, 2017) and its SHINE toolbox (Willenbockel et al., 2010), resulting in 171 grayscale images (e.g., stimuli, see Appendix C).

#### *Procedure*

Study 2a. Participants completed the task in the control and experimental conditions in succession (see Figure 3). In each condition, face animacy ratings (Block 1 through 3) were followed by face uncanniness ratings (Block 4).

Control. In the control condition, the 57 faces were presented in BSF. Animacy ratings (Block 1–3) replicated the procedure of Study 1, except that the faces were grayscale after being matched on low-level visual properties (e.g., mean luminance). Uncanniness ratings (Block 4) replicated Block 3, in which each face was presented for 1000 ms; however, participants were instead prompted to rate face perceived uncanniness by using a 6-point Likert-type scale ranging from 1 (*least uncanny*) to 6 (*most uncanny*). The prompt read, "How much does this face make you feel uncanny (e.g., feelings of eeriness, unease, and/or repulsion)?" The terms for measuring uncanny feelings were adopted based on the literature (e.g., Ho et al., 2008; Mori et al., 2012).

Experimental. In the experimental condition, animacy ratings followed a similar procedure, except that at the 500 ms (Block 2) and 1000 ms (Block 3) levels of exposure time, instead of showing the faces in BSF throughout the presentation, the same face was presented twice in succession first in BSF and then in HSF. Specifically, BSF was first presented

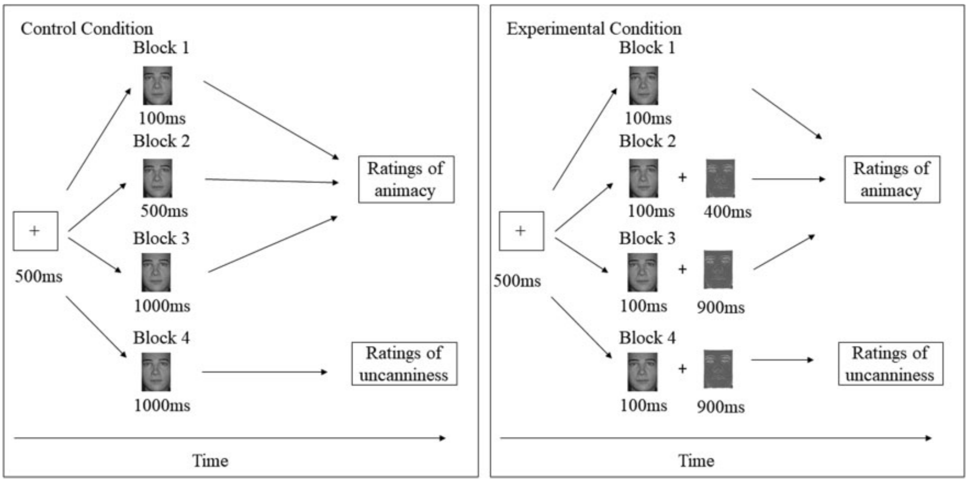


**Figure 3.** Schematic Representation of the Procedure in Study 2a, Which Consisted of a Control and an Experimental Condition, Each Comprising Four Blocks of Trials.
In both conditions, animacy rating appeared in the first three blocks and uncanniness rating followed in the fourth block. The procedure in Study 2b was the same as that in Study 2a, except that in the experimental condition, the filtered images presented in Blocks 2 to 4 preserved LSF rather than HSF.

for 100 ms and immediately followed by HSF for either 400 ms in Block 2 or 900 ms in Block 3. During the uncanniness rating task, each face was again first presented in BSF for 100 ms and immediately followed by HSF for another 900 ms; participants then rated the face on its perceived uncanniness.

Study 2b. Study 2b followed the same procedure of Study 2a, except that the filtered images preserved LSF rather than HSF.

## *Results and Discussion*

Because the perceived animacy of human and mechanical-looking robot faces did not significantly differ as a function of exposure time (Study 1), in Study 2, we focused on android faces to examine the effect of spatial frequency manipulation on faces' perceived animacy and uncanniness. In particular, according to the dehumanization hypothesis, we predicted that manipulating the spatial frequency of android faces should diminish the reduction in android faces' perceived animacy as a function of exposure time and reduce their perceived uncanniness.

*Study 2a.* We computed the mean ratings of animacy of android faces at each level of exposure time in both the control (BSF only) and experimental (BSF+HSF) conditions (see Table 2 for descriptive statistics for both Study 2a and 2b). Visual inspection of normal Q–Q plots indicated that the mean ratings of animacy and uncanniness for android faces were normally distributed for each combination of exposure time (100 ms, 500 ms, and 1000 ms) and condition (control and experimental; see Appendix D). Therefore, the mean ratings of animacy were submitted to a 2 (condition [BSF vs. BSF+HSF]) $\times$ 3 (exposure time [100 ms, 500 ms, and 1000 ms]) repeated-measures ANOVA. In the following analyses, we followed the same statistical procedures in Study 1.

**Table 2.** Mean (Standard Deviation in Parentheses) Values of Perceived Animacy for Each Level of Exposure Time by Condition and Study.

| Study | Condition | Exposure time (ms) | | |
|---|---|---|---|---|
| | | 100 | 500 | 1000 |
| Study 2a ($N = 22$) | Control | 3.46 (0.87) | 2.88 (0.86) | 2.65 (0.96) |
| | Experimental | 2.69 (0.90) | 3.03 (1.06) | 2.97 (1.10) |
| Study 2b ($N = 22$) | Control | 3.52 (0.79) | 2.90 (0.72) | 2.65 (0.82) |
| | Experimental | 2.80 (0.63) | 3.02 (0.70) | 3.00 (0.73) |

The ANOVA yielded a predicted interaction between condition and exposure time, $F(1.51, 31.79) = 34.82$, $p < .0001$, $\eta_p^2 = .62$. Corroborating our hypothesis, the findings in the control condition (BSF only) of Study 2a replicated those in Study 1, $F(1.53, 32.06) = 14.10$, $p = .0001$, $\eta_p^2 = .40$, and follow-up pairwise comparisons revealed significant differences in ratings of animacy between 100 ms ($M = 3.46$, $SD = 0.87$) and 500 ms ($M = 2.88$, $SD = 0.86$), $t(21) = 3.42$, Cohen's $d = 0.73$, $p = .003$, and between 100 ms ($M = 3.46$, $SD = 0.87$) and 1000 ms ($M = 2.65$, $SD = 0.96$), $t(21) = 4.38$, Cohen's $d = 0.93$, $p = .0003$, but not between 500 ms ($M = 2.88$, $SD = 0.86$) and 1000 ms ($M = 2.65$, $SD = 0.96$), $t(21) = 2.18$, Cohen's $d = 0.46$, $p = .04$.

In contrast, in the experimental condition (BSF+HSF), ratings of animacy differed among the three exposure time conditions, $F(1.44, 30.17) = 9.83$, $p = .001$, $\eta_p^2 = .32$. Follow-up pairwise comparisons revealed that ratings of animacy significantly differed between 100 ms ($M = 2.69$, $SD = 0.90$) and 500 ms ($M = 3.03$, $SD = 1.06$), $t(21) = -3.6$, Cohen's $d = -0.77$, $p = .002$, and between 100 ms ($M = 2.69$, $SD = 0.90$) and 1000 ms ($M = 2.97$, $SD = 1.10$), $t(21) = -3.0$, Cohen's $d = -0.64$, $p = .007$, but not between 500 ms ($M = 3.03$, $SD = 1.06$) and 1000 ms ($M = 2.97$, $SD = 1.10$) of exposure time, $t(21) = 1.24$, Cohen's $d = 0.26$, $p = .23$.

In addition to ratings of face animacy, we examined the effect of the spatial frequency manipulation (BSF+HSF) on participants' ratings of face uncanniness. We found that the mean ratings of uncanniness of android faces were significantly different between the control (BSF only; $M = 4.39$, $SD = 0.78$) and the experimental (BSF+HSF; $M = 3.99$, $SD = 0.90$) conditions, $t(21) = 3.70$, Cohen's $d = 0.79$, $p = .001$ (see Figure 4). All $p$ values were two-tailed.

*Study 2b.* Compared with the experimental condition of Study 2a in which images retaining only HSF were presented during the later stages of processing, we hypothesized that the manipulation in Study 2b would yield a similar or even stronger effect on ratings of face animacy and uncanniness. Like in Study 2a, we compared the changes in these ratings between the control (BSF only) and experimental (BSF+LSF) conditions.

The ANOVA yielded a predicted interaction between condition and exposure time, $F(1.28, 26.91) = 17.45$, $p = .0001$, $\eta_p^2 = .45$. Corroborating our hypothesis, the findings in the control condition (BSF only) of Study 2b replicated those in the control condition of Study 2a, $F(1.14, 23.92) = 18.14$, $p = .0002$, $\eta_p^2 = .46$, showing that ratings of animacy were significantly different between 100 ms ($M = 3.52$, $SD = 0.79$) and 500 ms ($M = 2.90$, $SD = 0.72$), $t(21) = 3.86$, Cohen's $d = 0.82$, $p < .001$; between 100 ms ($M = 3.52$, $SD = 0.79$) and 1000 ms ($M = 2.65$, $SD = 0.82$), $t(21) = 4.55$, Cohen's $d = 0.97$, $p < .001$; and between 500 ms ($M = 2.90$, $SD = 0.72$) and 1000 ms ($M = 2.65$, $SD = 0.82$), $t(21) = 3.98$, Cohen's $d = 0.85$, $p < .001$.

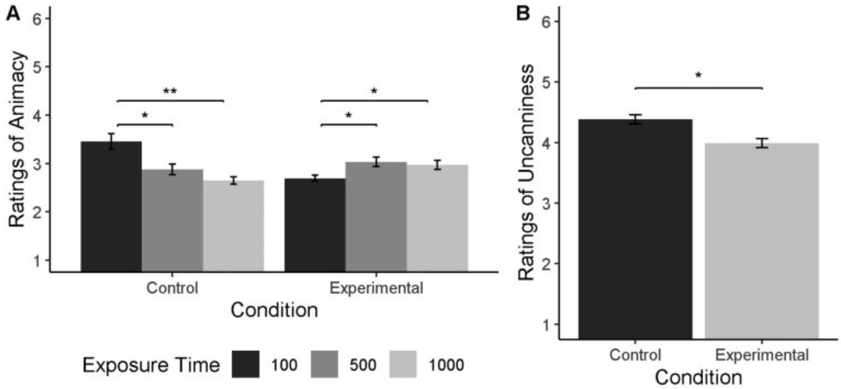


**Figure 4.** Mean Ratings of Animacy and Uncanniness of Android Faces in the Control (BSF) and Experimental (BSF + HSF) Conditions.
(a) Mean ratings of animacy of android faces at each level of exposure time (100 ms, 500 ms, and 1000 ms) in the control (BSF) and experimental (BSF + HSF) conditions. (b) Mean ratings of the uncanniness of android faces in the control and experimental conditions.
*$p < .01$. **$p < .001$.

Similar to the findings of Study 2a, presenting faces retaining only LSF during later stages of processing yield an overall main effect of exposure time on the perceived animacy of android faces, $F(1.69, 35.55) = 4.47$, $p = .024$, $\eta_p^2 = .18$. Nevertheless, follow-up pairwise comparisons revealed no significant differences in ratings of animacy between 100 ms ($M = 2.80$, $SD = 0.63$) and 500 ms ($M = 3.02$, $SD = 0.70$), $t(21) = –2.39$, Cohen's $d = –0.51$, $p = .027$; between 100 ms ($M = 2.80$, $SD = 0.63$) and1000 ms ($M = 3.00$, $SD = 0.73$), $t(21) = –2.31$, Cohen's $d = –0.49$, $p = .03$; and between 500 ms ($M = 3.02$, $SD = 0.70$) and 1000 ms ($M = 3.00$, $SD = 0.73$), $t(21) = .27$, Cohen's $d = 0.06$, $p = .79$ at a Bonferroni-adjusted alpha level of .017 per test (.05/3).

In addition to ratings of face animacy, we examined the effect of the spatial frequency manipulation (BSF+LSF) on participants' ratings of face uncanniness. We found that the mean ratings of uncanniness of android faces were significantly different between the control (BSF only; $M = 4.07$, $SD = 0.54$) and the experimental (BSF+LSF; $M = 3.83$, $SD = 0.52$) conditions, $t(21) = 2.86$, Cohen's $d = 0.61$, $p = .009$ (see Figure 5). All $p$ values were two-tailed.

To examine whether the magnitude of reduction in mean ratings of the uncanniness of android faces differs in terms of effect size between the experimental condition of Study 2b (Cohen's $d = 0.61$) and that of Study 2a (Cohen's $d = 0.79$), we conducted a mixed-design ANOVA treating the type of manipulation (BSF+HSF vs. BSF+LSF) as a between-subject and condition (control vs. experimental) as a within-subject factor. This analysis did not yield any significant interaction between manipulation type and condition, $F(1, 42) = 1.18$, $p = .28$, $\eta_p^2 = .03$. These findings suggest that removing either LSF or HSF from android faces during later stages of processing yields a similar effect in terms of reducing these faces' perceived uncanniness compared with the same faces containing the BSF in the control condition.

To ensure that our results are not driven by particular images or participants, we modeled item response data for the ratings of both animacy and uncanniness of android faces. The results showed that for both Study 2a and Study 2b (data were collapsed given that the type of manipulation was nonsignificant and excluded from the model), after controlling for the

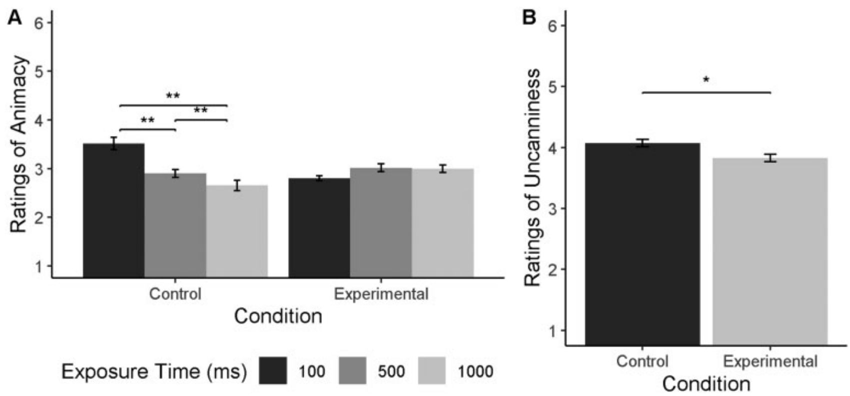


**Figure 5.** Mean Ratings of Animacy and Uncanniness of Android Faces in the Control (BSF) and Experimental (BSF+LSF) Conditions.
(A) Mean ratings of animacy of android faces at each level of exposure time (100 ms, 500 ms, and 1000 ms) in the control (BSF) and experimental (BSF+LSF) conditions. (B) Mean ratings of the uncanniness of android faces in the control and experimental conditions.
*$p < .01$. **$p < .001$.

random stimulus and participant effects, the interaction effect between condition and exposure time remained statistically significant for the ratings of animacy, $\chi^2$ (2) = 254.05, $p < .001$, and that the effect of condition remained statistically significant for the ratings of uncanniness, $\chi^2$ (1) = 45.54, $p < .001$.

In sum, in both Study 2a and Study 2b, we provided further evidence for the dehumanization hypothesis by showing that spatial frequency manipulation both eliminates the decrease in perceived animacy during later stages of processing for android faces and reduces their perceived uncanniness.

Contrary to our prediction, however, we found no statistically significant differences in the ratings of face uncanniness between the two spatial frequency manipulations in the two studies. One possibility is that although the scrutiny of details might characterize face animacy perception, coarse information might be indispensable for extracting fine details during face animacy perception. Alternatively, HSF might not play a more important role in face animacy perception than LSF. In fact, Goffaux et al. (2005) demonstrate that LSF is more important than HSF in supporting the configural processing of faces. Given the fact that configural processing contributes to the attribution of minds to faces (Deska et al., 2017) and the disruption of configural processing leads to dehumanization (Fincher & Tetlock, 2016), LSF might play a role as important as HSF during face animacy perception by facilitating configural face processing.

One might argue that the manipulation of spatial frequency, rather than influencing the perceived uncanniness of android faces via disrupting animacy perception, merely reduced faces' level of human likeness. If this hypothesis holds, the effect of spatial frequency manipulation on faces' perceived uncanniness should hold for android as well as human and mechanical-looking robot faces.

However, this alternative explanation is unlikely. An exploratory analysis rejected this prediction by showing that the effect of spatial frequency manipulation, unlike in androids, did not reach statistical significance in either human—Study 2a: $t(21) = -2.02$, Cohen's $d = -0.43$, $p = .06$; Study 2b: $t(21) = -0.33$, Cohen's $d = -0.07$, $p = .75$—or mechanical-looking robot faces, Study 2a: $t(21) = 1.05$, Cohen's $d = 0.22$, $p = .31$; Study 2b:

$t(21) = -0.50$, Cohen's $d = -0.11$, $p = .62$. Therefore, although spatial frequency manipulation might have reduced the level of human likeness of the faces, the differential effects of the manipulation among the different face types cannot be attributed to changes in the surface characteristics of the faces alone.

## General Discussion

Almost 50 years ago, Mori et al. (2012) predicted that as robots increasingly resemble humans, they may fall into an uncanny valley, eliciting unintended feelings of unease, repulsion, and eeriness, insofar as they are distinguishable from humans. Although researchers propose various hypotheses to account for this phenomenon, most focus on the surface characteristics of the human replicas. In contrast, few examine the cognitive underpinnings of humans' eerie feelings toward human replicas (for exceptions, see K. Gray & Wegner, 2012; Urgen et al., 2018). To our knowledge, the present study is the first that links the uncanny valley to the temporal dynamics of face animacy perception. Drawing upon evidence from a growing body of literature on face animacy perception (Looser & Wheatley, 2010; Wheatley et al., 2011), we conducted two studies to examine how perceived animacy changes as a function of exposure time across human, android, and mechanical-looking robot faces.

Collectively, our results demonstrate that the uncanny valley is inextricably linked to the temporal dynamics of face animacy perception. In Study 1, our results show that as a function of exposure time, perceived animacy undergoes a marked decrease for android faces, whereas it remains stable for both human and mechanical-looking robot faces. In Study 2, we further corroborate this link by showing that spatial frequency manipulation influences the decrease in perceived animacy of android faces and reduces their perceived uncanniness. Contrary to our prediction, however, removing HSF does not yield a stronger effect compared with the removal of LSF, suggesting that in addition to fine details, coarse information might contribute to face animacy perception during later stages of face processing.

Altogether, our findings support the dehumanization hypothesis against the mind perception hypothesis, showing that participants tend first to attribute and then deny[2] a mind to a face eliciting uncanny feelings. This dual process of anthropomorphism-then-dehumanization closely maps onto the two-stage (face-then-mind) model of face processing (Looser et al., 2013; Wheatley et al., 2011). In addition, our findings provide new insights into mind perception by demonstrating that humans perceive minds in faces differently depending on the particular type of faces, and such differences are characterized by the distinct temporal trajectories of faces' perceived animacy over time.

### *Dehumanization Hypothesis and Alternative Accounts of Uncanny Valley*

The dehumanization hypothesis is inextricably linked to other accounts of the uncanny valley (for a detailed discussion of this topic, see Wang et al., 2015). For example, it overlaps with the violation of expectation hypothesis and extends its contribution by specifying the nature of both the expectation and the violation thereof in relation to face animacy perception.

In this respect, the dehumanization hypothesis is particularly akin to the mind perception hypothesis (K. Gray & Wegner, 2012), which posits that the attribution of minds to nonhuman entities underpins androids' perceived uncanniness. This account implicitly assumes that observers perceive an android face as inanimate, and only with increasing exposure time do they start to attribute a mind to it. In the present research, however, we rejected this default assumption by demonstrating that the opposite process of dehumanization is more likely to account for the perceived uncanniness of android faces. Our findings, therefore, extend the contribution of K. Gray and Wegner (2012) on the basis of the temporal processes, now newly specified and by which mind perception might account for the uncanny valley phenomenon.

In addition, our findings demonstrate that mind perception does not occur in one shot but rather relies on the dynamic perception of face animacy over time. This dynamic perception might contribute to a sense of uncertainty, particularly linked to animacy perception (Carr et al., 2017; Jentsch, 1906/1997). A related but distinct hypothesis posits that the uncanny valley effect is attributable to categorical ambiguity that emerges at any categorical boundary (Ramey, 2006; Yamada et al., 2013). Although this account has received considerable attention in the literature, the evidence is mixed and beyond the scope of the present article (for a review, see Kätsyri et al., 2015).

Finally, the dehumanization and the perceptual mismatch hypotheses (e.g., Mitchell et al., 2011) are not mutually exclusive, but they rather focus on different levels of analyses (Marr, 1982). Whereas the perceptual mismatch hypothesis focuses on the computational level, mapping input (e.g., facial features of the human replica) to output (e.g., the uncanny feeling it elicits), the dehumanization hypothesis focuses on the algorithmic level of analysis, specifying the processes by which such mappings are achieved.

## *How Is Face Dehumanization Linked to the Uncanny Feelings?*

Although the present research links face animacy perception to the uncanny valley, it remains unclear how the dehumanization of a face might elicit the uncanny feelings.

Moral psychologists argue that moral emotions are evolved mechanisms that serve the function of facilitating social interactions (Sherman & Haidt, 2011). From a social functionalist perspective (Keltner & Haidt, 1999), the moral emotion of disgust serves as an affective mechanism for tracking negative social value in others, reducing social engagement (Harris & Fiske, 2007) and motivating social avoidance and dehumanization (Hodson & Costello, 2007; Rozin et al., 2008).

Given that humans and nonhuman agents are evaluated on the same continuum of humanness (Waytz et al., 2010), we argue that the dehumanization of both humans and anthropomorphized robots might activate evolved disgust mechanisms, motivating avoidance behavior and eliciting repulsion. In this respect, the dehumanization hypothesis extends the pathogen avoidance account[3] (MacDorman et al., 2009) by pointing to its proximate cause in terms of a cognitively/perceptually mediated process of dehumanization based on face animacy perception.

## *Limitations and Future Directions*

In the present research, we examined participants' ratings of animacy of human, android, and mechanical-looking robot faces, presenting each stimulus at three exposure times (i.e., 100 ms, 500 ms, and 1000 ms). In so doing, we fell short of fully capturing the temporal dynamics of face animacy perception on millisecond levels. Repeated exposure to the stimuli may also contribute to a distorted judgment of face uncanniness (Zlotowski et al., 2015). Given these limitations, future research should establish alternative measures of face animacy perception with a higher temporal resolution to examine further the cognitive and neural underpinnings of the uncanny valley. As such, event-related electroencephalography method (Wheatley et al., 2011), coupled with implicit measures of mind perception in faces (e.g., lexicon decision task; Hugenberg et al., 2015), should be a fruitful future direction.

In addition, although the perceptual roots of mind perception have been mainly investigated in the domain of face perception, other body parts, such as hands, may also convey valuable information regarding whether an agent possesses a mind. In fact, Mori et al. (2012) originally described the uncanny valley hypothesis by using prosthetic hands as an example. Although researchers have demonstrated that prosthetic hands elicit more eerie feelings than human or mechanical-looking hands, within prosthetic hands, increasing human likeness tends to reduce perceived eeriness (Poliakoff et al., 2013, 2018). Future research examining the temporal dynamics of animacy perception in hands should account for these differences between faces and hands and further illuminate the perceptual roots of mind perception.

## Conclusions

Mind perception lies at the center of interpersonal relations and might likely play a significant role in human–robot interactions. In the present study, we focus on mind perception by examining the temporal dynamics of face animacy perception across human, android, and mechanical-looking faces. Our findings indicate that humans' affective responses toward human replicas are determined by not only their physical features but also the temporal dynamics by which we perceive minds in their faces. By shifting attention from the human replicas' physical features to the perceivers' ascription of minds to them, researchers can provide new insights into the psychology of the uncanny valley and eventually shed light on how we perceive each other as humans.

### 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.

### Notes

1. Levels of uncanniness: human faces < mechanical-looking robot faces < android faces.
2. This denial of mind to android faces, rather than completely depriving them from the potential to possess a humanlike mind, reduces the degree of mind attributed to them.
3. Pathogen avoidance hypothesis is the idea that the detection of facial features (e.g., skins) in androids indicative of poor health or contagious disease elicits uncanny feelings by activating evolved pathogen avoidance mechanisms that underpin human disgust responses.

## ORCID iD

Shensheng Wang https://orcid.org/0000-0002-6472-3940

## References


Bartneck, C., Kanda, T., Ishiguro, H., & Hagita, N. (2007). *Is the uncanny valley an uncanny cliff?* [Paper presentation]. RO-MAN 2007 – The 16th IEEE International Symposium on Robot and Human Interactive Communication, Jeju, South Korea.

Broadbent, E. (2017). Interactions with robots: The truths we reveal about ourselves. *Annual Review of Psychology*, *68*, 627–652. https://doi.org/10.1146/annurev-psych-010416-043958

Carr, E. W., Hofree, G., Sheldon, K., Saygin, A. P., & Winkielman, P. (2017). Is that a human? Categorization (dis)fluency drives evaluations of agents ambiguous on human-likeness. *Journal of Experimental Psychology: Human Perception and Performance*, *43*, 651–666. https://doi.org/10.1037/xhp0000304

Crouzet, S. M., Kirchner, H., & Thorpe, S. J. (2010). Fast saccades toward faces: Face detection in just 100 ms. *Journal of Vision*, *10*, 11–17. https://doi.org/10.1167/10.4.16

Damiano, L., & Dumouchel, P. (2018). Anthropomorphism in human–robot co-evolution. *Frontiers in Psychology*, *9*, 468. https://doi.org/10.3389/fpsyg.2018.00468

Deska, J. C., Almaraz, S. M., & Hugenberg, K. (2017). Of mannequins and men: Ascriptions of mind in faces are bounded by perceptual and processing similarities to human faces. *Social Psychological and Personality Science*, *8*, 183–190. https://doi.org/10.1177/1948550616671404

Deska, J. C., & Hugenberg, K. (2017). The face-mind link: Why we see minds behind faces, and how others' minds change how we see their face. *Social and Personality Psychology Compass*, *11*, e12361. https://doi.org/10.1111/spc3.12361

DiSalvo, C., Gemperle, F., Forlizzi, J., & Kiesler, S. (2002). *All robots are not created equal*: *Design and the perception of humanness in robot heads* [Paper presentation]. DIS2002 Conference Proceedings, London, England.

Epley, N., & Waytz, A. (2010). Mind perception. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of social psychology* (Vol. *1*, pp. 498–541). John Wiley.

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). GPower 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175–191. https://doi.org/10.3758/BF03193146

Fincher, K. M., & Tetlock, P. E. (2016). Perceptual dehumanization of faces is activated by norm violations and facilitates norm enforcement. *Journal of Experimental Psychology: General*, *145*, 131–146. https://doi.org/10.1037/xge0000132

Fink, J. (2012). Anthropomorphism and human likeness in the design of robots and human-robot interaction. In Ge S. S., Khatib O., Cabibihan J. J., Simmons R., & Williams M. A. (Eds.), *Social robotics* (vol. 7621). ICSR 2012. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer.

Gao, T., McCarthy, G., & Scholl, B. J. (2010). The wolfpack effect: Perception of animacy irresistibly influences interactive behavior. *Psychological Science*, *21*, 1845–1853. https://doi.org/10.1177/0956797610388814

Goffaux, V., Hault, B., Michel, C., Vuong, Q. C., & Rossion, B. (2005). The respective role of low and high spatial frequencies in supporting configural and featural processing of faces. *Perception*, *34*, 77–86. https://doi.org/10.1068/p5370

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., & 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

Guthrie, S. (1993). *Faces in the clouds*. Oxford University Press.

Harris, L. T., & Fiske, S. T. (2007). Social groups that elicit disgust are differentially processed in mPFC. *Social Cognitive and Affective Neuroscience*, *2*, 45–51. https://doi.org/10.1093/scan/nsl037

Ho, C.-C., MacDorman, K. F., & Pramono, Z. A. D. D. (2008). *Human emotion and the uncanny valley: A GLM, MDS, and Isomap analysis of robot video ratings* [Paper presentation]. Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction, Amsterdam, the Netherlands.

Hodson, G., & Costello, K. (2007). Interpersonal disgust, ideological orientations, and dehumanization as predictors of intergroup attitudes. *Psychological Science*, *18*, 691–698. https://doi.org/10.1111/j.1467-9280.2007.01962.x

Hugenberg, K., Young, S., Rydell, R. J., Almaraz, S., Stanko, K. A., See, P. E., & Wilson, J. P. (2015). The face of humanity: Configural face processing influences ascriptions of humanness. *Social Psychological and Personality Science*, *7*, 167–175. https://doi.org/10.1177/1948550615609734

Jentsch, E. (1906/1997). On the psychology of the uncanny. *Angelaki*, *2*, 7–16. https://doi.org/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. https://doi.org/10.3389/fpsyg.2015.00390

Keltner, D., & Haidt, J. (1999). Social functions of emotions at four levels of analysis. *Cognition and Emotion*, *13*, 505–521. https://doi.org/10.1080/026999399379168

Locker, L., Hoffman, L., & Bovaird, J. A. (2007). On the use of multilevel modeling as an alternative to items analysis in psycholinguistic research. *Behavior Research Methods*, *39*, 723–730. https://doi.org/10.3758/BF03192962

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 Affective Neuroscience*, *8*, 799–805. https://doi.org/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. https://doi.org/10.1177/0956797610388044

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.

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

Marr, D. (1982). *Vision*. MIT Press.

Martini, M. C., Gonzalez, C. A., & Wiese, E. (2016). Seeing minds in others—Can agents with robotic appearance have human-like preferences? *PLoS One*, *11*, 1–23. https://doi.org/10.1371/journal.pone.0146310

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. https://doi.org/10.1016/j.cognition.2015.09.008

Misselhorn, C. (2009). Empathy with inanimate objects and the uncanny valley. *Minds and Machines*, *19*, 345–359. https://doi.org/10.1007/s11023-009-9158-2

Mitchell, W. J., Szerszen, K. A., 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. https://doi.org/10.1068/i0415

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

Peters, J. C., Goebel, R., & Goffaux, V. (2018). From coarse to fine: Interactive feature processing precedes local feature analysis in human face perception. *Biological Psychology*, *138*, 1–10. https://doi.org/10.1016/j.biopsycho.2018.07.009

Poliakoff, E., Beach, N., Best, R., Howard, T., & Gowen, E. (2013). Can looking at a hand make your skin crawl? Peeringÿintoÿthe uncanny valley for hands. *Perception*, *42*, 998–1000.

Poliakoff, E., O'Kane, S., Carefoot, O., Kyberd, P., & Gowen, E. (2018). Investigating the uncanny valley for prosthetic hands. *Prosthetics and Orthotics International*, *42*, 21–27. https://doi.org/10.1177/0309364617744083

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

Rozin, P., Haidt, J., & McCauley, C. R. (2008). Disgust. In M. Lewis, J. M. Haviland-Jones, & L. F. Barrett (Eds.), *Handbook of emotions* (3rd ed., pp. 757–776). The Guilford Press.

Schroeder, J., & Epley, N. (2016). Mistaking minds and machines: How speech affects dehumanization and anthropomorphism. *Journal of Experimental Psychology: General*, *145*, 1427–1437. https://doi.org/10.1037/xge0000214

Sherman, G. D., & Haidt, J. (2011). Cuteness and disgust: The humanizing and dehumanizing effects of emotion. *Emotion Review*, *3*, 245–251. https://doi.org/10.1177/1754073911402396

Singmann, H., Bolker, B., Westfall, J., & Aust, F. (2018). afex: Analysis of factorial experiments. R package version 0.19-1. https://CRAN.R-project.org/package = afex

The Mathworks Inc. (2017). *Version 9.3.0.713579 (R2017b)*.

Urgen, B. A., Kutas, M., & Saygin, A. P. (2018). Uncanny valley as a window into predictive processing in the social brain. *Neuropsychologia*, *114*, 181–185. https://doi.org/10.1016/j.neuropsychologia.2018.04.027

Varga, S. (2017). The case for mind perception. *Synthese*, *194*, 787–807. https://doi.org/10.1007/s11229-015-0994-8

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

Wang, S., & Rochat, P. (2017). Human perception of animacy in light of the uncanny valley phenomenon. *Perception*, *46*, 1386–1411. https://doi.org/10.1177/0301006617722742

Wang, S., Rochat, P., Pye, A., & Poliakoff, E. (2020). *"Dehumanized" faces and hands fall into the uncanny valley* [Manuscript in preparation]. Department of Psychology, Emory University. Division of Neuoscience and Experimental Psychology, Faculty of Biology, Medicine and Health, University of Manchester.

Waytz, A., Epley, N., & Cacioppo, J. (2010). Social cognition unbound: Insights into anthropomorphism and dehumanization. *Current Directions in Psychological Science*, *19*, 58–62. https://doi.org/10.1177/0963721409359302

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. https://doi.org/10.1371/journal.pone.0017960

Willenbockel, V., Sadr, J., Fiset, D., Horne, G., Gosselin, F., & Tanaka, J. N. (2010). The SHINE toolbox for controlling low-level image properties. *Journal of Vision*, *10*, 653. https://doi.org/10.1167/10.7.653

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. https://doi.org/10.1111/j.1468-5884.2012.00538.x

Zlotowski, 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. https://doi.org/10.3389/fpsyg.2015.00883

## Appendix A. Human, Android, and Mechanical-Looking Faces. (Study 1)

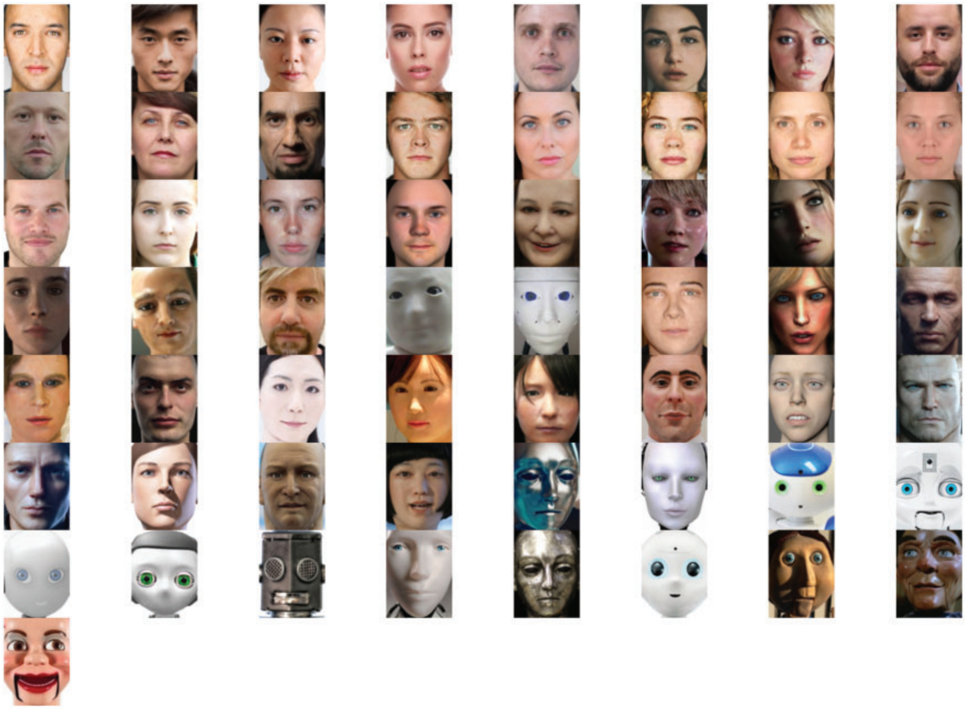

## Appendix B.

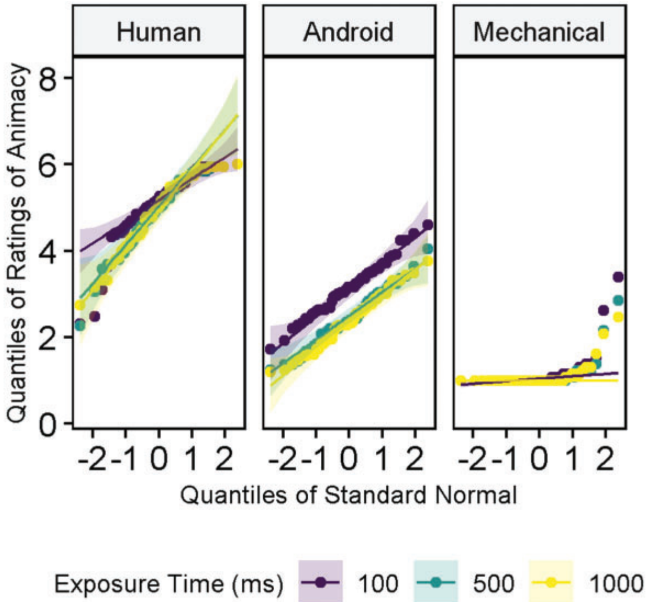


Normal Quantile–Quantile Plots (i.e., Normal Q–Q plots) of Mean Ratings of Animacy for Each Combination of Exposure Time (100 ms, 500 ms, and 1000 ms) and Face Type (Human, Android, and Mechanical) in Study 1.

## Appendix C. Example of Human, Android, and Mechanical-Looking Faces in BSF, HSF, and LSF. (Study 2)

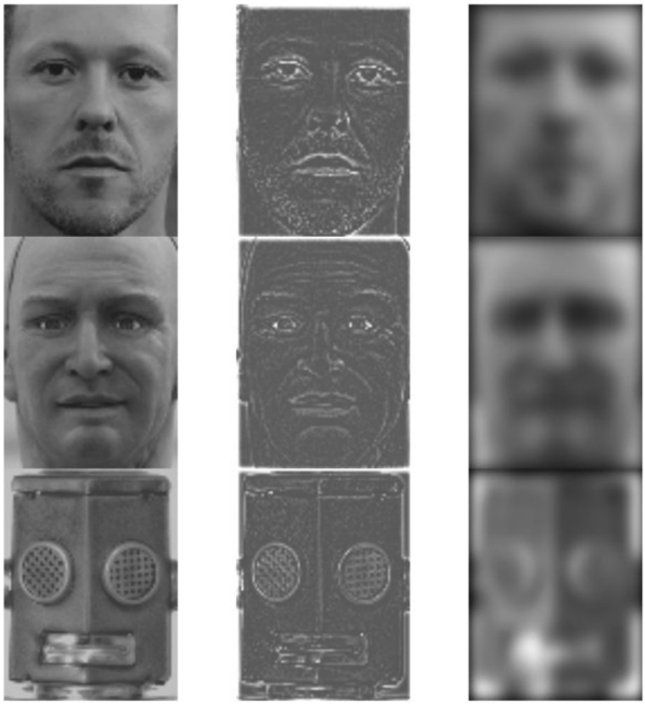

## Appendix D.

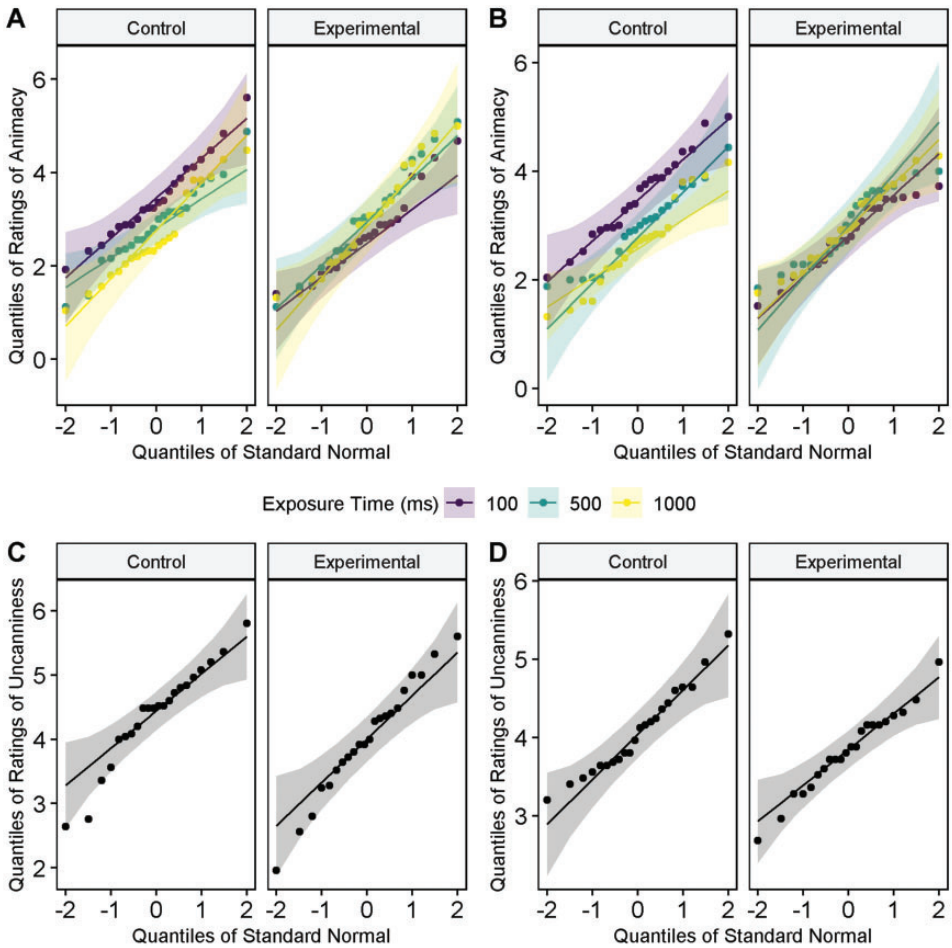


Normal Q–Q Plots of Mean Ratings of Animacy and Uncanniness in Study 2: (A) Mean ratings of animacy in Study 2a. (B) Mean ratings of animacy in Study 2b. (C) Mean ratings of uncanniness in Study 2a. (D) Mean ratings of uncanniness in Study 2b.