



# Rethinking the uncanny valley as a moderated linear function: Perceptual specialization increases the uncanniness of facial distortions

Alexander Diel<sup>a,b,c,\*</sup>, Michael Lewis<sup>a</sup>

<sup>a</sup> Cardiff University School of Psychology, UK

<sup>b</sup> University of Duisburg-Essen, Clinic for Psychosomatic Medicine and Psychotherapy, LVR University Hospital Essen, 45147, Essen, Germany

<sup>c</sup> Center for Translational Neuro- and Behavioral Sciences (C-TNBS), University of Duisburg-Essen, 45147, Essen, Germany

## ARTICLE INFO

Handling editor: Jen-Her Wu

### Keywords:

Face processing  
Inversion effect  
Moderated linear function  
Specialization  
Uncanny valley

## ABSTRACT

The relationship between artificial entities' human likeness and aesthetic preference is thought to be best modelled by an *N*-shaped cubic "uncanny valley" function, which however suffers from conceptual criticisms and lack of parsimony. Here it is argued that uncanniness effects may instead be modelled by a linear function of deviation moderated by perceptual specialization. The two models are compared in an experiment with five incrementally distorted face types (cartoon, CG, drawing, real, robot). Recognition performance for upright and inverted faces were used as a specialization measure. Specialization significantly moderated the linear effect of distortion on uncanniness, and could explain the data better than a conventional uncanny valley. The uncanny valley may thus be better understood as a moderated linear function of specialization sensitizing the uncanniness of deviating stimuli. This simpler yet more accurate model is compatible with neurocognitive theories and can explain uncanniness effects beyond the conventional uncanny valley.

## 1. Introduction

### 1.1. The uncanny valley

The acceptance of artificial humanlike entities gains growing importance with accelerated technological advancement. Social robots become increasingly widespread, for example in healthcare or hospitality (Broekens, Heerink, & Rosendal, 2009; Dawe, Sutherland, Barco, & Broadbent, 2019; Lu et al., 2020; Nakanishi et al., 2020), including during the COVID-19 pandemic (Aymerich-Franch & Ferrer, 2022). Similarly, realistic computer-generated (CG) characters find increasing use in commerce and healthcare (Ma, Sharifi, & Chatopadhyay, 2019; Scherer & Von Wangenheim, 2014). Yet a lack of acceptance of artificial humans due to their imperfectly humanlike appearance may reduce trust and likability (Appel, Weber, Krause, & Mara, 2016; Davies, 2016; Destephe et al., 2015; Mathur & Reichling, 2016; Olaronke, Rhoda, & Janet, 2017; Tinwell, Grimshaw, Nabi, & Williams, 2011). Understanding and overcoming issues caused by artificial entities' near humanlike appearance is essential to avoid material risks and improve smooth human-technology interaction.

Providing artificial entities with humanlike characteristics generally

increases likability (Mara, Appel, & Gnambs, 2022). However, approximating a certain level of humanlike appearance elicits strange, eerie or uncanny perceptions (Diel, Weigelt, & Macdorman, 2021; Ho & MacDorman, 2017; Mori, 2012). This *uncanny valley* effect is statistically defined as a nonlinear (quadratic or cubic) function between entities' human likeness and likability. Its statistical validity and underlying mechanisms have been a topic of human-computer interaction research for decades (Kätsyri et al., 2015; MacDorman and Ishiguro, 2006; Mori, 2012). Although the uncanny valley effect is a well-replicated phenomenon (Diel et al., 2021; Kätsyri et al., 2015; MacDorman and Ishiguro, 2006; Mori, 2012), conceptual limitations and a lack of parsimony (complex cubic functions are unusual in nature) beg the question of the validity of a cubic relationship between human likeness and likability or related ratings (here called the *contemporary uncanny valley model*). Here it is investigated whether this contemporary uncanny valley model can instead be rethought of as a moderated linear function of deviation, uncanniness, and specialization. This model of uncanniness is capable of explaining a broader range of observations beyond the contemporary uncanny valley model while not suffering from its limitations.

\* Corresponding author. University of Duisburg-Essen, Clinic for Psychosomatic Medicine and Psychotherapy, LVR University Hospital Essen, 45147 Essen, Germany.

E-mail address: [diela@cardiff.ac.uk](mailto:diela@cardiff.ac.uk) (A. Diel).

<https://doi.org/10.1016/j.chb.2024.108254>

Received 10 June 2023; Received in revised form 25 March 2024; Accepted 10 April 2024

Available online 12 April 2024

0747-5632/© 2024 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

### 1.1.1. Parsimony of a cubic relationship

Although simplicity is preferred in scientific explanations, interactions between variables do not always follow linear relationships. Polynomial degree reduction can help clarifying otherwise complex statistical patterns into simple laws: For example, the quadratic Yerkes-Dodson law of stress and performance can be reduced to a linear “deviation-from-optimum” relationship (Yerkes & Dodson, 1908). The non-monotonic nature of the uncanny valley requires a model that is cubic with a part that is concave up and a part that is concave down. Very few phenomena in nature follow a cubic function like this and so, if the uncanny valley is describing the simple relationship between two properties then its relationship would be fairly unique. A simpler alternative to a cubic model is to introduce a third variable that moderates the effect of the independent variable on the dependent variable depending on the degree of the former. In terms of the uncanny valley, a high level of human likeness may sensitize the effect of deviation from typical appearance on likability (or uncanniness), increasing the uncanniness if a stimulus is anomalous.

In this sense, the uncanny valley function can be understood as a complex example of the Simpson’s paradox: a statistical trend occurring over a whole set of data that disappears within sub-groups of the data. An example of how a third moderating variable may explain the otherwise cubic uncanny valley relationship is depicted in Fig. 1.

According to this view, the hypothetical moderating variable (see Fig. 1B) would depend on the stimulus category, and would correlate with the level of human likeness. This said moderator variable would increase the sensitivity to distortions, with the highest sensitivity for the most humanlike categories. If the rightmost group in Fig. 1B would be “realistic human face”, then the most humanlike stimuli would be real human faces, while the most uncanny stimuli would be realistic yet deviating human face, e.g., an android faces. Because the sensitivity to distortions in realistic faces is the highest, even slight errors in android faces would become apparent. The leftmost group may be “mechanical robot faces”, for which a wide range of face structure variation can be acceptable due to the low level of specialization for robot face structure.

Furthermore, if only “ideal” (i.e., non-distorted stimuli) would be considered, the plot would correspond to a simple linear relationship between human likeness and likeability (see Mara et al., 2022). In other words, human likeness increases likeability but also the sensitivity to potential distortions that would appear uncanny if detected. Because android faces are highly realistic, the chance of detecting design errors is especially high, leading to an uncanny valley effect in higher levels of human likeness.

If this were correct, then the sensitivity to distortions should be highest for the most realistic faces – which is consistent with previous literature (Diel & Lewis, 2022a; MacDorman, Green, Ho, & Koch, 2009; Mäkäräinen, Kätsyri, & Takala, 2014). The moderator variable should

then be able to explain how the sensitivity is higher for certain categories.

### 1.1.2. “Human likeness” and the uncanny valley

The contemporary uncanny valley model focuses on a human likeness dimension (Mori, 2012) which remains an essential part of the uncanny valley’s understanding today (Diel et al., 2021; Mara et al., 2022).

Human likeness is a multidimensional holistic impression of humanness (von Zitzewitz et al., 2013) including physical features like appearance (MacDorman & Ishiguro, 2006) and variables like behavior (Zlotowski et al., 2015). Mori (2012) himself suggested that Japanese *bunraku* puppets would lie beyond the uncanny valley due to their naturally human facial expressions and behaviour in theater, despite their clearly non-humanlike physical appearance. Uncanny valley research tends to focus on single-scale items of *human likeness* or *realism* to measure this variable, or on indices measuring multiple aspects of human likeness (Diel et al., 2021; Ho & MacDorman, 2017).

Despite the focus on *human* likeness, uncanny valley effects have been observed using animal stimuli (Diel & MacDorman, 2021; Löffler, Dörrenbächer, & Hassenzahl, 2020; MacDorman & Chattopadhyay, 2016; Sierra Rativa, Postma, & van Zaenen, 2022; Schwind, Wolf, & Henze, 2018a,b; Yamada, Kawabe, & Ihaya, 2013). Yet when including both human and non-human animal stimuli in one dataset, focusing only on a human likeness dimension would not sufficiently represent the animal-related uncanny valley effects. Furthermore, uncanny valley-like effects have been observed for inanimate categories like written text or physical places (Diel et al., 2021; Diel & Lewis, 2022c, 2022d). A two-variable model including only likability/uncanniness and human likeness would insufficiently account for uncanny valley effects beyond human stimuli.

In summary, while human likeness is considered a critical component in uncanny valley research, it is neither uniformly defined nor does it explain the total number of observations.

### 1.1.3. Likeability/uncanniness and the uncanny valley

While Mori (2012) originally called the dependent variables of the uncanny valley *shinwakan* and *bukimi*, often translated as likeability or affinity and eeriness or uncanniness (Bartneck et al., 2009; Ho & MacDorman, 2017), a wide range of terms have been used to measure the uncanny valley effect (Diel et al., 2021). Although likeability is a commonly used measure, it may be susceptible to confounding variables that may decrease the general appeal of a stimulus without being strictly uncanny (Diel et al., 2021). Meanwhile, several researchers have emphasized that the “uncanny” aspect of the uncanny valley is marked by a specific negative subjective experience of *eeriness*, *uncanniness*, or the *uncanny feeling* (Benjamin & Heine, 2023; Ho & MacDorman, 2017;

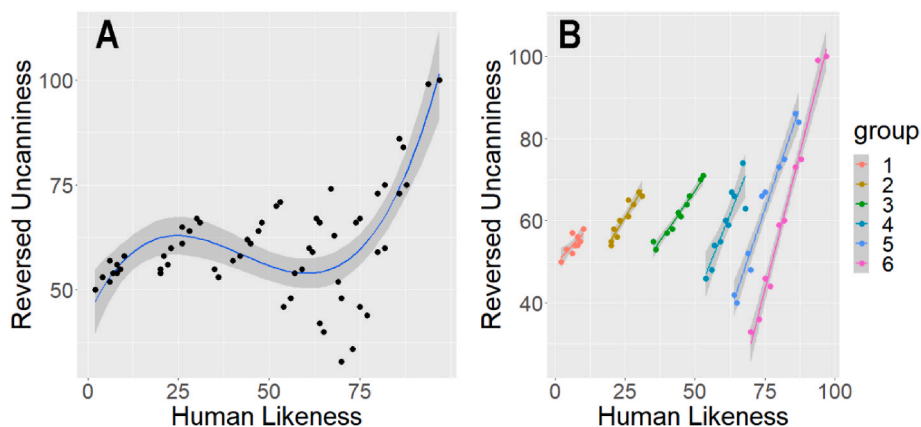


Fig. 1. An example representation of the uncanny valley fitted as a cubic relation between human likeness and uncanniness plotted over a set hypothetical data (1A), and a potential moderated linear function explaining the same set of data (1B).

Mangan, 2015), which may be related to the emotions of fear and disgust (Ho, Macdorman, & Pramono, 2008; MacDorman & Ishiguro, 2006). Hence, measuring a specific negative experience would have a higher discriminant validity compared to a general positive measure (e.g., likability).

Although behavioral, neural, and physiological measures have been used to investigate the uncanny valley (see Vaitonytė et al., 2023 Backspace; for a scoping review; see also Diel et al., 2021), a lack research concerning clearly valid measures makes it difficult to interpret neural correlates. It is unclear whether neural or physiological correlates would occur due to an experience of uncanniness or due to other factors, such as changes in cognitive processing due to manipulated stimulus properties. So far, there are no neural or physiological responses that are consistently elicited by different types of uncanny stimuli (e.g., faces, bodies, motion, animals, etc.) that can be used as indicators of an uncanny valley effect. Furthermore, neural or physiological measures may not capture the specific subjective experience that is critical to the uncanny valley, and may thus be better used to investigate the experience itself using already validated stimuli, rather than to measure the uncanny valley in order to validate stimuli.

In sum, a variety of methods have been used to measure the uncanny valley. As the uncanny valley is marked by a specific negative experience (eerie, uncanny, etc.), measures ought to capture this specific experience in order not to be confounded by other variables.

#### 1.1.4. Typicality/deviation and likability/uncanniness

A more general approach predicts changes in likability depending on a stimulus' typicality (or degree of deviation): deviating stimuli or patterns tend to be disliked across categories, and individual differences in the degree of aversion can be transferred across stimulus categories (Gollwitzer et al., 2017). Aversion caused by deviations may be related to increased processing disfluency (Winkielman, Schwarz, Fazendeiro, & Reber, 2003) or violations of expectations in predictive coding (Friston, 2010). Analogously, the uncanny valley is associated with a violation of cognitive structures (Lischetzke, Izydorczyk, Hüller, & Appel, 2017). Such mechanisms would not be bound to a human likeness dimension, and could explain uncanny valley effects across animal (e.g., Schwind, Wolf, & Henze, 2018a,b) and inanimate object (e.g., Diel & Lewis, 2022c) categories.

However, the effect of typicality (or deviation) on likability (or uncanniness) is not consistent across categories: Analogous distortions increase uncanniness more in human compared to cat faces, and cat faces compared to buildings (Diel & MacDorman, 2021). Furthermore, effects of facial distortion on likability are more pronounced in more realistic faces (Diel & Lewis, 2022a; Green, MacDorman, Ho, & Vasudevan, 2008; MacDorman et al., 2009; Mäkääinen et al., 2014). While the contemporary uncanny valley does not provide a clear solution on why the effect of deviation on uncanniness depends on the stimulus category, a redefined model may benefit from adding a moderating variable defining the strength of the effect of deviation on uncanniness.

#### 1.1.5. Specialization as a moderator variable

A high sensitivity to deviations in especially realistic human-related stimuli (e.g., faces) may be due to a high degree of processing specialization: Humans are highly specialized for upright human faces (Gauthier & Nelson, 2001; Maurer & Werker, 2014; Rhodes, Brake, Taylor, & Tan, 1989), enabling assessment of facial identity and aesthetics based on feature-relational information; a process that is disturbed when faces are presented inverted (*inversion effect*; Carbon & Leder, 2006; Mondloch, Le Grand, & Maurer, 2002). Perceptual specialization is not exclusive to faces (Gauthier & Nelson, 2001), and trained specialization for an otherwise novel category increases the uncanniness of distorted variants compared to non-distorted variants (Diel & Lewis, 2022b). As specialization is high in human stimulus categories, deviations would appear especially uncanny, creating an uncanny valley effect in artificial humans with slight design errors.

The inversion effect has been used as a measure of the degree of specialization. Face inversion effects are reduced for less realistic faces (e.g., computer-generated, virtual) compared to typical human faces (Balas & Pacella, 2015; Crookes et al., 2015; Di Natale, Simonetti, La Rocca, & Bricolo, 2023). Higher specialization for more realistic faces could thus explain a higher sensitivity to distortions in more realistic faces (Diel & Lewis, 2022a; Green et al., 2008; MacDorman et al., 2009; Mäkääinen et al., 2014). Furthermore, as specialization is less pronounced in less realistic faces like robots (Sacino et al., 2022; Zlotowski & Bartneck, 2013), tolerance for atypicalities or deviations should be higher in these categories. Together with a high deviation sensitivity in more realistic humanlike stimuli, the uncanny valley effect would thus emerge across the dimension of human likeness (Fig. 1).

Thus, the degree of specialization is a suitable third variable candidate for simplifying the uncanny valley into a moderated linear function. A moderated linear function of specialization, typicality/deviation, and likability/uncanniness may explain a wider range of data with higher accuracy than a nonlinear contemporary uncanny valley model while being statistically simpler and theoretically plausible (Mori, 2012).

### 1.2. Research question and hypotheses

The aim of this work is to investigate whether the uncanny valley can be better understood as a moderated linear function of deviation, uncanniness, and specialization. Specifically, it is investigated whether specialization in different face types (face inversion effect) moderates the effect of face distortion (incremental changes in face feature positions) on uncanniness: Uncanniness is expected to increase with facial distortions (up to a point at which a face is so distorted it is no longer perceived or categorized as a human face), and this effect should furthermore increase with specialization in the face group.

First, the face inversion effect is replicated for each face category, and it is replicated whether the effect is stronger for more realistic compared to less realistic faces (e.g., Crookes et al., 2015; Sacino et al., 2022).

1. A face inversion effect is stronger for more realistic human faces (human and cartoon faces) compared to less realistic faces (drawing and robot faces) (*inversion effect hypothesis*)

Second, the conventional uncanny valley is investigated by testing whether a polynomial (quadratic or cubic) function of human likeness ratings can explain uncanniness ratings (Mori, 2012).

2. A polynomial function of human likeness can explain the uncanniness across faces better than a linear function (*uncanny valley hypothesis*)

Third, it is suggested that a moderated linear function underlies the uncanny valley: a higher sensitivity to distortions is proposed for more humanlike or realistic faces, which would increase the relative uncanniness caused by deviations. As specialization sensitizes the detection of changes and distortions, it is tested whether recognition accuracy differences between upright and inverted faces for each condition (face inversion effect as a marker of expertise) predicts the effect of distortion on uncanniness.

3. Degree of inversion effect predicts uncanniness caused by distortion (*moderation hypothesis I*)

Finally, to investigate whether a moderated linear function is superior to an uncanny valley, it is tested whether specialization as a moderator for distortion and uncanniness can explain the data better than a polynomial function of uncanniness and human likeness.

4. A moderated linear function of face distortion level, uncanniness, and inversion effect can explain the data better than a nonlinear function of uncanniness and human likeness (*moderation hypothesis II*)

## 2. Methods

### 2.1. Participants

As previous research found odds ratio (OR) values of 0.62 (converted to a Cohen's  $d = 0.264$ ) for inversion effects in robot stimuli (Sacino et al., 2022), a power analysis with an effect size of  $d = 0.264$  revealed that 120 participants is sufficient for a power of  $1 - \beta = 0.8$ . Participants ( $M_{\text{age}} = 19.4$ ,  $SD_{\text{age}} = 0.84$ ) were 120 undergraduate Psychology students of the Cardiff University School of Psychology; 103 identified as female and 17 as male. Undergraduate Psychology students were recruited for convenience.

### 2.2. Material

Human faces were selected from the Chicago Face Database (Ma, Correll, & Wittenbrink, 2015). Different sets of 12 faces (three female, three male) were used for each of the first three levels of realism (real, cartoon, drawing). Cartoon and drawing faces were created using the cartoon character and sketch character tools of VanceAI toongineer (<https://vanceai.com/toongineer-cartoonizer/>). Realism level 4 (CG) faces were created using FACSGen. Finally, realism level 5 (robot) faces were selected from a previous study locating a wide range of robot faces before the uncanny valley (Mathur & Reichling, 2016).

All faces were distorted in the same manner: Distance between eyes were incrementally increased in five faces per group, and decreased in the other five faces, by 10% of the eyes' horizontal length. In addition, the position of the mouth was either incrementally increased or decreased (each in five faces per group) by 25% of the mouth's vertical length. A total of five distortion levels, including the original, were created. Distortions were manipulated in both directions (e.g., eye distances were either increased or decreased) to control for different types of facial distortions.

Finally, half of the undistorted base faces of each face realism group were inverted for the face recognition task. Stimuli divided by condition can be seen in Fig. 2.

### 2.3. Procedure

#### 2.3.1. Face recognition task

The face recognition task consisted of an encoding part and a recognition part. Only undistorted faces were used in the face recognition task. In the encoding part, participants viewed a total of 60 faces (12 faces per realism level; half upright, half inverted; half female, half male) sequentially in a random order. Participants were allowed to view each face for as long as they wanted. In the recognition part, an additional novel 60 faces (12 faces per realism level; half upright, half inverted) were shown to the participants together with the learnt faces, and participants were asked to indicate for each face whether they have seen the face in the encoding phase. Again, participants had an indefinite amount of time to decide for each face while simultaneously viewing the face.

#### 2.3.2. Face rating task

Both undistorted and distorted faces were used in the rating task. In the face rating task, each face was shown the participants in a random order, together with three scales: *uncanny/eerie*, *strange/weird*, and *realistic/humanlike*. Scales were shown in the same order as mentioned here, and the *strange/weird* scales were reversed. Participants were to rate each face on each scale ranging from 0 to 100. Each face was shown for the entire time until participants responded for each scale. A total of

300 faces (10 faces per 6 realism level and 5 distortion levels) were rated. Participants had an indefinite amount of time to rate each face while the face was presented.

#### 2.3.3. Data analysis and availability

RStudio and JASP were used for data analysis. Inversion effect index was calculated as 1 minus inverted recognition accuracy divided by upright recognition accuracy. Outlier removal (1.5 IQR from median) was performed on a stimulus-level for recognition accuracy (inversion effect index) and ratings. For the inversion effect index, 456 individual outlier values were removed. For the rating analysis, 389 individual human likeness ratings and 97 individual uncanniness ratings were removed. For the inversion effect, interactions between face realism and face orientation were investigated using within-subject within-base stimulus ANOVAs. For the prediction and comparison of the polynomial and moderating functions, linear mixed models with base stimulus and participants as random factors were used. The functions *lmer* from the R packages *lme4* and *lmerTest* were used for the linear mixed models (Bates, Mächler, Bolker, & Walker, 2015). Data and analysis are publicly available at <https://osf.io/xvh24>. The study was not preregistered.

## 3. Ethics statement

Research was conducted in accordance with the Declaration of Helsinki. The study was approved by the Cardiff University ethics committee board (EC.23.October 01, 6716).

## 4. Results

### 4.1. Face inversion effect

Recognition accuracy was calculated by averaging by-participant numbers of correct responses participant for each face condition. The face inversion effect was then tested by calculating quotient of inverted and upright face recognition accuracy for each face condition. The data is summarized in Fig. 3.

A within-subject and within-base stimulus ANOVA on face recognition accuracy with face type and orientation as factors found significant main effects of face type ( $F(1,118) = 28.57$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.001$ ) and orientation ( $F(4,115) = 37.383$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.01$ ) and a significant interaction ( $F(4,115) = 3.32$ ,  $p = 0.01$ ,  $\eta_p^2 = 0.001$ ). Results are indicative of a face inversion effect that differs between face types.

To investigate the face inversion effect per face type, post-hoc comparisons with Bonferroni-adjusted  $p$ -values between upright and inverted faces were performed for each face type. Upright faces were significantly better recognized than inverted faces for real faces ( $t(14271) = 4.6$ ,  $p_{\text{adj}} < 0.001$ ,  $d = 0.49$ ) and cartoon faces ( $t(14271) = 3.07$ ,  $p_{\text{adj}} = 0.006$ ,  $d = 0.32$ ), but not for face drawings ( $t(14271) = 2.11$ ,  $p_{\text{adj}} = 0.09$ ), CG faces ( $t(14271) = 0.16$ ,  $p_{\text{adj}} = 1$ ), or robot faces ( $t(14271) = 1.07$ ,  $p_{\text{adj}} = 1$ ). Thus, inversion effects were observed for real and cartoon faces, but not for the other face conditions. Thus, hypothesis 1 (*face inversion hypothesis*) is supported.

### 4.2. Uncanny valley

The uncanny valley hypothesis was investigated by testing whether a polynomial relationship between human likeness and uncanniness can explain the data better than a linear function. Linear mixed models with participants and base faces as random effects and linear, quadratic, and cubic function of human likeness as fixed effects were performed as predictors of uncanniness. Significant linear ( $t(1561) = -5.5$ ,  $p < 0.001$ ), quadratic ( $t(1561) = 9.29$ ,  $p < 0.001$ ), and cubic ( $t(1561) = -55.42$ ,  $p < 0.001$ ) functions of human likeness were found. Furthermore, the quadratic (AIC = 138004) model was a significantly better fit compared to the linear (AIC = 138897) model ( $\chi^2 = 894.58$ ,  $p < 0.001$ ), while the cubic (AIC = 137770) model was a better fit than the quadratic


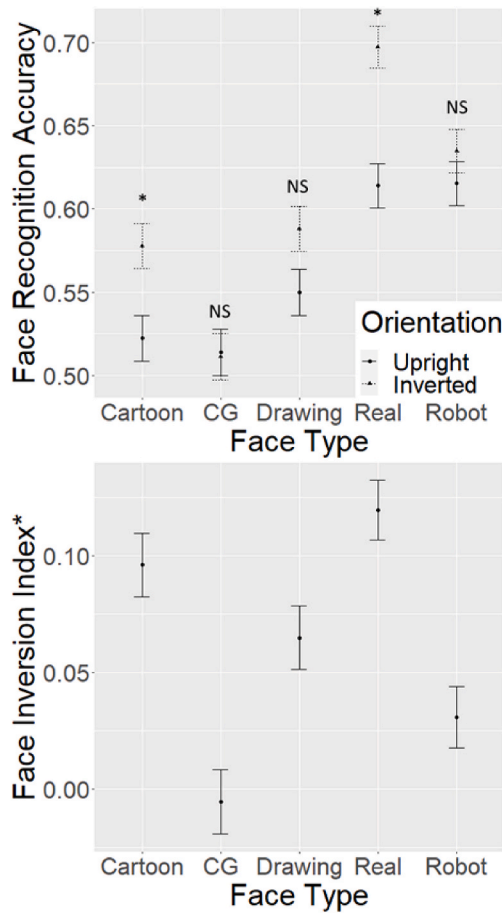


Fig. 2. Upright stimuli divided by distortion (horizontal axis; 0 to 4) and face type (vertical axis; real, cartoon, drawing, CG, robot) conditions. Note: Faces were also presented inverted real, cartoon, and drawing faces depicted here were not used in the experiment. The faces were artificially created by the StyleGAN generative network (Karras, Laine, Aittala, Hellsten, Lehrinen, & Aila, 2020).



Note. Face Inversion Index = 1 - (inverted face recognition accuracy / upright face recognition accuracy).

Fig. 3. Average face recognition accuracy across face type and condition (A) and level of face inversion effect (1 - inverted accuracy/upright accuracy) across face type (B). For 3A, asterisks show significantly higher recognition rates for upright compared to inverted faces while “NS” indicate no significant increases. Error bars indicate standard errors.

Note. Face Inversion Index = 1 - (inverted face recognition accuracy/upright face recognition accuracy).

Note. Face Inversion Index = 1 - (inverted face recognition accuracy/upright face recognition accuracy).

one ( $\chi^2 = 235.96, p < 0.001$ ). Thus, the cubic model of human likeness ( $R_c^2 = 0.43$ ) could best explain uncanniness. Thus, hypothesis 2 (*uncanny valley hypothesis*) is supported.

The fit is depicted in Fig. 4A..

#### 4.3. A moderated linear function

Hypothesis 3 predicts that a moderated linear function of distortion and expertise (face inversion index) explains uncanniness, and better than a polynomial function. Face inversion index as a proxy to expertise has been calculated as 1 minus inversion recognition rate divided by upright recognition rate. Linear mixed models with participant and base face as random factors and distortion and face inversion index as random effects found significant main effects of distortion ( $t(15496) = 2.83, p < .001$ ) and face inversion index ( $t(15559) = -13.41, p < .001$ ), as well as a significant interaction ( $t(15496) = 3.57, p < .001; R^2_c = .37$ ). The interaction is summarized in Fig. 5: face types are sorted based on their average Face Inversion Index (FII). Uncanniness levels are plotted for each stimulus distortion level. The plot and results show that

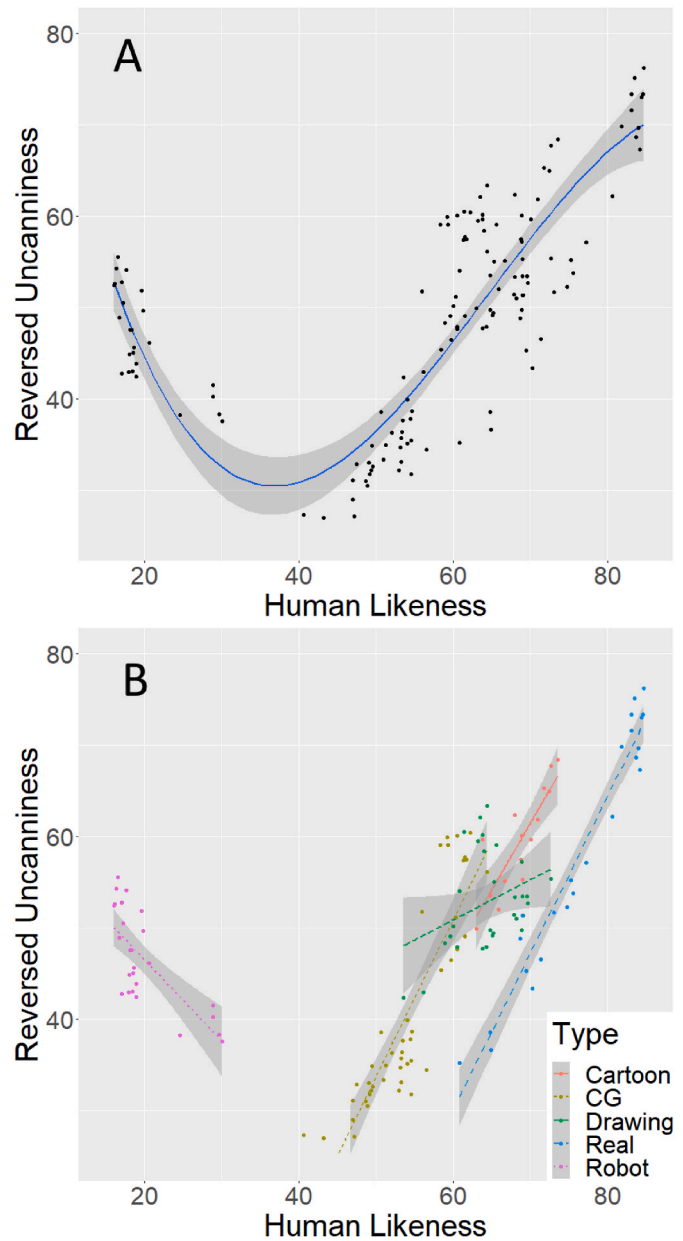


Fig. 4. Nonlinear (A) and moderated linear (B) fits on uncanniness. Gray areas represent standard errors. Dots show individual stimulus values averaged across participants, and are the same for both 4A and 4B. Fig. 4A depicts a nonlinear function of human likeness across all face types while Fig. 4B depicts linear relationships between human likeness and uncanniness for each base stimulus, categorized by types.

FII moderates the increase of uncanniness across distortion levels, supporting the notion of a moderated linear function.

In summary, with a higher degree of a face’s inversion effect, the effect of distortion on uncanniness increased. Thus, hypothesis 3 (moderation hypothesis I) was supported

Finally, to test whether a moderated linear function can best explain uncanniness, a linear mixed model with participants and base faces as random effects and actor type, distortion, face inversion index, and human likeness as fixed effects has been calculated and tested against the cubic function of human likeness. The moderated linear model ( $R_c^2 = 0.46; AIC = 137213$ ) could explain uncanniness better than the cubic ( $R_c^2 = 0.43; AIC = 137770$ ) model ( $\chi^2 = 628.96, p < 0.001$ ). Thus, hypothesis 4 (moderation hypothesis II) was supported.

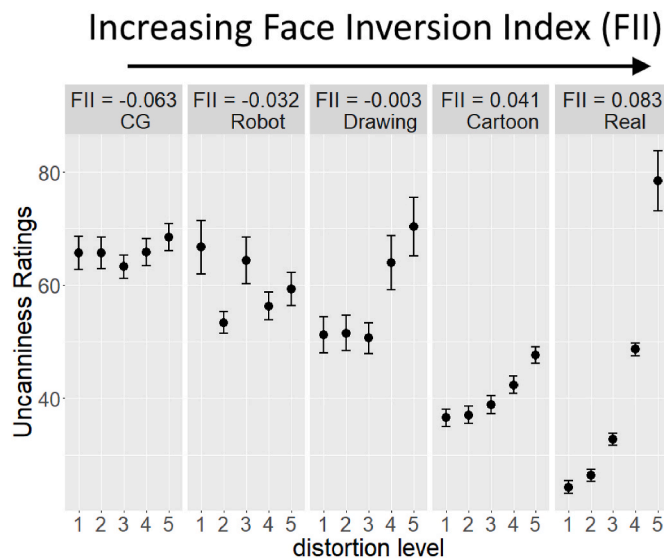


Fig. 5. Mean uncanniness ratings across face distortion levels divided by face realism type. Face realism types are sorted by level of FII. FII significantly predicted the effect of distortion on uncanniness ratings.

## 5. Discussion

### 5.1. Summary of results

The uncanny valley (Mori, 2012) is typically used to describe the relationship between artificial entities' human likeness and likability. Here it was investigated whether face inversion (a marker for specialized processing) can moderate the sensitivity of uncanniness to facial distortions, and whether such a moderated linear relationship can explain the data better than a traditional polynomial uncanny valley plot (Mori, 2012).

In accordance with previous research, the strength of inversion effects differed across face conditions (e.g., Di Natale et al., 2023; Sacino et al., 2022): Inversion effects were found for real human faces and cartoon faces, but were not found for drawing-style faces, CG faces, and robot faces (Fig. 3).

A function of human likeness and uncanniness indicative of an uncanny valley was found (Fig. 4A). However, the same data can also be plotted as a moderated linear function when data is divided by face type (Fig. 4B): changes of uncanniness across human likeness were linear in each face type with type-dependent slopes. These results are comparable to the hypothesized linear moderated relationship (Fig. 1).

A moderated linear function of face inversion index and deviation could significantly explain uncanniness: In faces with a higher face inversion effect, deviation caused stronger increases in uncanniness (Fig. 5). This model could explain the same data better than a contemporary uncanny valley model, indicating that a simpler moderated linear model is a more accurate representation than the nonlinear uncanny valley.

In summary, the uncanny valley can be rethought of as a moderated linear function: at lower human likeness levels (e.g., robots), specialization is relatively low, leading to a low sensitivity to deviations. Thus, a wide variation of near humanlike designs remains acceptable despite a spectrum of exaggerated or distorted face or body configurations. Higher levels of human likeness specialization cause higher sensitivity to distortions, increasing relative uncanniness. Slight deformations may be recognized through specialized processing that would be otherwise acceptable in less realistic entities. Thus, simple linear effects of deviation on uncanniness may be more pronounced in more specialized categories. When the data is however plotted only on the dimensions of human likeness and uncanniness, a nonlinear uncanny valley emerges

(Fig. 3).

### 5.2. Rethinking the uncanny valley

One advantage of a rethought moderated uncanniness function lies in the range of explainable data: the traditional uncanny valley model was restricted to a dimension of human likeness (Mori, 2012), complicating interpretations of an uncanny valley modelled for animal stimuli (e.g., Schwind, Wolf, & Henze, 2018a,b). A moderated linear function meanwhile can generalize predictions onto any stimulus category: Not only can a higher sensitivity towards distorted human faces compared to animal faces be explained by a higher level of specialization for the former; the theory also encompasses uncanniness in inanimate categories. The moderated linear model can also explain that the sensitivity to facial distortion is increased for more realistic faces (Diel & Lewis, 2022a; Green et al., 2008; MacDorman et al., 2009; Mäkäräinen et al., 2014), in familiar faces (Diel & Lewis, 2022a; Jung, Lee, & Choi, 2022), human compared to cat faces (Diel & MacDorman, 2021), or own-ethnicity faces (Saneyoshi, Okubo, Suzuki, Oyama, & Laeng, 2022).

### 5.3. Practical implications of a moderated linear function

The presented statistical reframing of the uncanny valley as a moderated linear function provides practical implications for the design of artificial humans. First, the new model reinforces previous suggestions that an uncanny valley occurs at higher levels of human likeness (due to a higher level of specialization), and can thus be avoided by designing artificial humans in a less realistic, perhaps cartoonish or stylistic manner in order to decrease the level of specialized processing (Schwind et al., 2018b). Furthermore, the uncanny valley is expected to depend on an individual's demographic: As face specialization is increased for faces of a demographic the viewer is used to (e.g., Meissner & Brigham, 2001), stronger uncanny valley effects are ought to occur when individuals are confronted with artificial humans of similar demographics (see also Saneyoshi et al., 2022). In a related vein, an uncanny valley may be avoided when deliberately designing artificial humans by deliberately choosing to design them in a demographic the viewers may be less familiar with.

Finally, the results suggests that an uncanny valley effect is not exclusive to physical appearance, but can also emerge in other modalities for which perceptual specialization is expected, such a biological motion (Tobin, Favelle, & Palermo, 2016). According to the new proposed model, uncanniness effects may occur due to the detection of anomalies or deviations in dynamic aspects (e.g., the motion of facial expressions) even if the appearance of an artificial human is acceptable (Diel, Sato, Hsu, & Minato, 2023). Hence, the new model provides cautionary predictions in that individuals would not only be sensitive to an artificial entity's appearance, but also humanlike motion and behavior. Care would thus need to be taken for the appropriate design of realistic human motion and behavior.

### 5.4. Limitations and future research

Colored images were used in this experiment. Different colors between the realism levels (e.g., colorful real human and cartoon faces, less or not colorful sketch or robot faces) may affect recognition ability between realism levels. However, as the colors would be the same for upright and inverted faces within the realism levels, different inversion effects should have not occurred due to different color schemes even though general recognition may have been improved for more colorful images.

This study's methodology focused on self-assessment research. While self-reported rating scales are the most commonly used measure in uncanny valley research (Diel et al., 2021), they may be limited reporting effects: for examples, participants may be urged to provide a positive response after repeatedly giving negative responses, although stimulus

randomization should control for such repetition effects. In addition, as the methods used here are consistent with those used in previous research, the same limitations would apply to those as well. Neural or physiological measures may provide measures without such limitations, although there are currently no such methods that can reliably measure an uncanny valley effect in different uncanny stimulus categories. Nevertheless, additional measures can be used in future research to investigate the moderated linear function proposed here: For example, perceptual specialization has been linked to activity in specialized brain areas, such as the fusiform gyrus for face stimuli (Kanwisher & Moscovitch, 2000). Distortions in specialized categories may correlate with increased activity in specialized areas, potentially due to an increased processing need for distorted stimuli. Analogously, Kim et al. (2016) found increased face-sensitive neural activity for uncanny faces. Future research may further investigate the association between specialization, uncanniness caused by distortion, and activity in specialized brain areas.

Finally, it would be interesting to investigate the existence of a moderating linear function in stimulus categories beyond faces. Perceptual specialization has been observed for body stimuli (Keye, Mingming, Tiantian, Wenbo, & Weiqi, 2017; Reed, Stone, Bozova, & Tanaka, 2003). More realistic artificial bodies may fall into an uncanny valley because they elicit stronger specialized processing, making potential design errors more apparent. Similarly, humans are specialized to dynamic facial expressions (Martinez, 2017; Tobin et al., 2016), which may be a cause of uncanniness in artificial humans attempting to replicate such expressions (Tinwell et al., 2011). In fact, recent evidence suggests that specialized processing for dynamic emotion expression increases uncanniness effects caused by distortions in face expression motion, supporting the notion of the linear moderated function in dynamic facial expressions (Diel et al., 2023). Future research may further investigate the interaction between specialization, deviation, and uncanniness across different categories.

## 6. Conclusion

The uncanny valley describes a nonlinear, *N*-shaped relationship between entities' human likeness and likability. Here it was observed that, rather than being a polynomial function, the uncanny valley can be better understood as a linear function between typicality (or deviation) and likability (or uncanniness), moderated by the degree of specialization of the relevant stimulus' category. Such a statistical model can explain a wider range of previous data compared to the "traditional" uncanny valley, including uncanniness observed in inanimate objects and moderating effects of a stimulus' realism level or species. Thus, rethinking the uncanny valley as a moderated linear function serves well to improve the understanding of uncanniness caused by anomalous or deviating stimuli.

## CRedit authorship contribution statement

**Alexander Diel:** Writing – review & editing, Writing – original draft, Visualization, Validation, Investigation, Formal analysis, Data curation, Conceptualization. **Michael Lewis:** Writing – review & editing, Visualization, Supervision, Resources, Conceptualization.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Data availability

The link to the data (OSF) is shared in the manuscript.

## References

- Appel, M., Weber, S., Krause, S., & Mara, M. (2016). On the eeriness of service robots with emotional capabilities. In *2016 11<sup>th</sup> ACM/IEEE international conference on human-robot interaction (HRI)* (pp. 411–412). New Zealand: Christchurch. <https://doi.org/10.1109/HRI.2016.7451781>.
- Aymerich-Franch, L., & Ferrer, I. (2022). Liaison, safeguard, and well-being: Analyzing the role of social robots during the COVID-19 pandemic. *Technology in Society*, 70, Article 101993.
- Balas, B., & Pacella, J. (2015). Artificial faces are harder to remember. *Computers in Human Behavior*, 52, 331–337. <https://doi.org/10.1016/j.chb.2015.06.018>
- 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(1), 71–81. <https://doi.org/10.1007/s12369-008-0001-3>.
- Benjamin, R., & Heine, S. J. (2023). From Freud to Android: Constructing a Scale of Uncanny Feelings. *Journal of personality assessment*, 105(1), 121–133. <https://doi.org/10.1080/00223891.2022.2048842>.
- Broekens, J., Heerink, M., & Rosendal, H. (2009). Assistive social robots in elderly care: A review. *Gerontechnology*, 8(2), 94–103. <https://doi.org/10.4017/gt.2009.08.02.002.00>
- Carbon, C.-C., & Leder, H. (2006). When faces are heads: View-dependent recognition of faces altered relationally or compositionally. *Swiss Journal of Psychology / Schweizerische Zeitschrift für Psychologie / Revue Suisse de Psychologie*, 65(4), 245–252. <https://doi.org/10.1024/1421-0185.65.4.245>
- Crookes, K., Ewing, L., Gildenhuis, J. D., Kloth, N., Hayward, W. G., Oxner, M., et al. (2015). How well do computer-generated faces tap face expertise? *PLoS One*, 10(11), Article e0141353. <https://doi.org/10.1371/journal.pone.0141353>
- Davies, N. (2016). Can robots handle your healthcare? *Engineering & Technology*, 11(9), 58–61. <https://doi.org/10.1049/et.2016.0907>
- Dawe, J., Sutherland, C., Barco, A., & Broadbent, E. (2019). Can social robots help children in healthcare contexts? A scoping review. *BMJ Paediatrics Open*, 3(1), Article e000371. <https://doi.org/10.1136/bmjpo-2018-000371>
- Destephe, M., Brandao, M., Kishi, T., Zecca, M., Hashimoto, K., & Takanishi, A. (2015). Walking in the uncanny valley: Importance of the attractiveness on the acceptance of a robot as a working partner. *Frontiers in Psychology*, 6, 204. <https://doi.org/10.3389/fpsyg.2015.00204>
- Di Natale, A. F., Simonetti, M. E., La Rocca, S., & Bricolo, E. (2023). Uncanny valley effect: A qualitative synthesis of empirical research to assess the suitability of using virtual faces in psychological research. *Computers in Human Behavior*, 10, Article 100288. <https://doi.org/10.1016/j.chbr.2023.100288>
- Diel, A., & Lewis, M. (2022a). Familiarity, orientation, and realism increase face uncanniness by sensitizing to facial distortions. *Journal of Vision*, 22(4), 14. <https://doi.org/10.1167/jov.22.4.14>
- Diel, A., & Lewis, M. (2022b). The deviation-from-familiarity effect: Expertise increases uncanniness of deviating exemplars. *PLoS One*, 17(9), Article e0273861. <https://doi.org/10.1371/journal.pone.0273861>
- Diel, A., & Lewis, M. (2022c). Structural deviations drive an uncanny valley of physical places. *Journal of Environmental Psychology*, 82, Article 10844. <https://doi.org/10.1016/j.jenvp.2022.101844>
- Diel, A., & Lewis, M. (2022d). The uncanniness of written text is explained by configural deviation and not by processing disfluency. *Perception*, 3010066221114436. <https://doi.org/10.1177/03010066221114436>. Advance online publication.
- Diel, A., & MacDorman, K. F. (2021). Creepy cats and strange high houses: Support for configural processing in testing predictions of nine uncanny valley theories. *Journal of Vision*, 21(4). <https://doi.org/10.1167/jov.21.4.1>. Article 1.
- Diel, A., Sato, W., Hsu, C. T., & Minato, T. (2023). Asynchrony enhances uncanniness in human, android, and virtual dynamic facial expressions. *BMC Research Notes*, 16(1), 368. <https://doi.org/10.1186/s13104-023-06648-w>
- Diel, A., Weigelt, S., & Macdorman, K. F. (2021). A meta-analysis of the uncanny valley's independent and dependent variables. *ACM Transactions on Human-Robot Interaction (THRI)*, 11, 1–33. <https://doi.org/10.1145/3470742>
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. <https://doi.org/10.1038/nrn2787>
- Gauthier, I., & Nelson, C. A. (2001). The development of face expertise. *Current Opinion in Neurobiology*, 11(2), 219–224. [https://doi.org/10.1016/S0959-4388\(00\)00200-2](https://doi.org/10.1016/S0959-4388(00)00200-2)
- Gollwitzer, A., Marshall, J., Wang, Y., & Bargh, J. A. (2017). Relating pattern deviancy aversion to stigma and prejudice. *Nature human behaviour*, 1(12), 920–927. <https://doi.org/10.1038/s41562-017-0243-x>.
- Green, R. D., MacDorman, K. F., Ho, C.-C., & Vasudevan, S. (2008). Sensitivity to the proportions of faces that vary in human likeness. *Computers in Human Behavior*, 24(5), 2456–2474. <https://doi.org/10.1016/j.chb.2008.02.019>
- Ho, C., Macdorman, K. F., & Pramono, Z. A. (2008). Human emotion and the uncanny valley: A glm, MDS, and isomap analysis of robot video ratings. *2008 3rd ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 169–176. <https://doi.org/10.1145/1349822.1349845>
- Ho, C.-C., & MacDorman, K. F. (2017). Measuring the uncanny valley effect: Refinements to indices for perceived humanness, attractiveness, and eeriness. *International Journal of Social Robotics*, 9(1), 129–139. <https://doi.org/10.1007/s12369-016-0380-9>.
- Jung, N., Lee, M., & Choi, H. (2022). The uncanny valley effect for celebrity faces and celebrity-based avatars. *Science of Emotion and Sensibility*, 25(1), 91–102. <https://doi.org/10.14695/KJSOS.2022.25.1.91>
- Kanwisher, N., & Moscovitch, M. (2000). The cognitive neuroscience of face processing: An introduction. *Cognitive Neuropsychology*, 17, 1–11. <https://doi.org/10.1080/026432900380454>



- 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>.
- Keye, Z., Mingming, Z., Tiantian, L., Wenbo, L., & Weiqi, H. (2017). The inversion effect of body recognition. *Advances in Psychological Science*, 27(1), 27–36. <https://doi.org/10.3724/SP.J.1042.2019.0002>
- Kim, D.-G., Kim, H.-Y., Kim, G., Jang, P.-S., Jung, W.-H., & Hyun, J.-S. (2016). Exploratory understanding of the uncanny valley phenomenon based on event-related potential measurement. *Korean Society for Emotion and Sensibility*, 19, 95–110. <https://doi.org/10.14695/KJSOS.2016.19.1.95>
- Lischetzke, T., Izdorczyk, D., Hüller, C., & Appel, M. (2017). The topography of the uncanny valley and individuals' need for structure: A nonlinear mixed effects analysis. *Journal of Research in Personality*, 68, 96–113. <https://doi.org/10.1016/j.jrp.2017.02.001>
- Löffler, D., Dörrenbächer, J., & Hassenzahl, M. (2020). The uncanny valley effect in zoomorphic robots: The U-shaped relation between animal likeness and likeability. *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 261–270. <https://doi.org/10.1145/3319502.3374788>
- Lu, V. N., Wirtz, J., Kunz, W. H., Paluch, S., Gruber, T., Martins, A., et al. (2020). Service robots, customers and service employees: What can we learn from the academic literature and where are the gaps? *Journal of Service Theory and Practice*, 19(3), 361–391. <https://doi.org/10.1108/JSTP-04-2019-0088>
- Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, 47(4), 1122–1135.
- Ma, T., Sharif, H., & Chattopadhyay, D. (2019). Virtual humans in health-related interventions: A meta-analysis. *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–6. <https://doi.org/10.1145/3290607.3312853>
- MacDorman, K. F., & Chattopadhyay, D. (2016). Reducing consistency in human realism increases the uncanny valley effect; increasing category uncertainty does not. *Cognition*, 146, 190–205. <https://doi.org/10.1016/j.cognition.2015.09.019>
- 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(3), 695–710. <https://doi.org/10.1016/j.chb.2008.12.026>
- MacDorman, K. F., & Ishiguro, H. (2006). The uncanny advantage of using androids in cognitive and social science research. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, 7(3), 297–337. <https://doi.org/10.1075/is.7.3.03mac>
- Mäkäräinen, M., Kätsyri, J., & Takala, T. (2014). Exaggerating facial expressions: A way to intensify emotion or a way to the uncanny valley? *Cognitive Computation*, 6, 708–721. <https://doi.org/10.1007/s12559-014-9273-0>, 2014.
- Mangan, B. (2015). The uncanny valley as fringe experience. *Interaction Studies*, 16(2), 193–199. <https://doi.org/10.1075/is.16.2.05man>.
- Mara, M., Appel, M., & Gnambs, T. (2022). Human-like robots and the uncanny valley: A meta-analysis of user responses based on the godspeed scales. *Zeitschrift für Psychologie*, 230(1), 33–46. <https://doi.org/10.1027/2151-2604/a000486>
- Martinez, A. M. (2017). Visual perception of facial expressions of emotion. *Current opinion in psychology*, 17, 27–33. <https://doi.org/10.1016/j.copsyc.2017.06.009>
- 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>
- Maurer, D., & Werker, J. F. (2014). Perceptual narrowing during infancy: A comparison of language and faces. *Developmental Psychobiology*, 56(2), 154–178. <https://doi.org/10.1002/dev.21177>
- Meissner, C. A., & Brigham, J. C. (2001). Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review. *Psychology, Public Policy, and Law*, 7, 3–35.
- Mondloch, C. J., Le Grand, R., & Maurer, D. (2002). Configural face processing develops more slowly than featural face processing. *Perception*, 31(5), 553–566. <https://doi.org/10.1068/p3339>
- Mori, M. (2012). The uncanny valley (K. F. MacDorman & norri kageki, trans.). *IEEE Robotics and Automation*, 19(2), 98–100 (Original work published in 1970).
- Nakanishi, J., Kuramoto, I., Baba, J., Ogawa, K., Yoshikawa, Y., & Ishiguro, H. (2020). Continuous hospitality with social robots at a hotel. *SN Applied Sciences*, 2, 452. <https://doi.org/10.1007/s42452-020-2192-7>
- Olaronke, I., Rhoda, I., & Janet, O. (2017). A framework for avoiding uncanny valley in healthcare. *IJBHTM*, 7(1), 1–10.
- Reed, C. L., Stone, V. E., Bozova, S., & Tanaka, J. (2003). The bodyinversion effect. *Psychological Science*, 14(4), 302–308. <https://doi.org/10.1111/1467-9280.14431>
- Rhodes, G., Brake, S., Taylor, K., & Tan, S. (1989). Expertise and configural coding in face recognition. *British Journal of Psychology*, 80(3), 313–331. <https://doi.org/10.1111/j.2044-8295.1989.tb02323.x>
- Sacino, A., Cocchella, F., De Vita, G., Bracco, F., Rea, F., Sciutti, A., et al. (2022). Human-or object-like? Cognitive anthropomorphism of humanoid robots. *PLoS One*, 17(7), Article e0270787. <https://doi.org/10.1371/journal.pone.0270787>
- Saneyoshi, A., Okubo, M., Suzuki, H., Oyama, T., & Laeng, B. (2022). The other-race effect in the uncanny valley. *International Journal of Human-Computer Studies*, 166, Article 102871. <https://doi.org/10.1016/j.ijhcs.2022.102871>
- Scherer, A., & Von Wangenheim, F. (2014). Service with a smile or screen? How replacing personnel with machines affects customers' satisfaction with a service. *Advances in Consumer Research*, 42, 662–663.
- Schwind, V., Leicht, K., Jäger, S., Wolf, K., & Henze, N. (2018a). Is there an uncanny valley of virtual animals? A quantitative and qualitative investigation. *International Journal of Human-Computer Studies*, 111, 49–61. <https://doi.org/10.1016/j.ijhcs.2017.11.003>
- Schwind, V., Wolf, K., & Henze, N. (2018b). Avoiding the uncanny valley in virtual character design. *Interactions*, 25(5), 45–49. <https://doi.org/10.1145/3236673>
- Sierra Rativa, A., Postma, M., & van Zaanen, M. (2022). The uncanny valley of a virtual animal. *Computer Animations and Virtual Worlds*, 33(2), e2043. <https://doi.org/10.1002/cav.2043>
- 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(2), 741–749. <https://doi.org/10.1016/j.chb.2010.10.018>
- Tobin, A., Favelle, S., & Palermo, R. (2016). Dynamic facial expressions are processed holistically, but not more holistically than static facial expressions. *Cognition & Emotion*, 30(6), 1208–1221. <https://doi.org/10.1080/02699931.2015.1049936>
- Vaitonytė, J., Alimardani, M., & Louwerse, M. M. (2023). Scoping review of the neural evidence on the uncanny valley. *Computers in Human Behavior Reports*, 9, Article 100263. <https://doi.org/10.1016/j.chbr.2022.100263>
- von Zitzewitz, J., Boesch, P. M., Wolf, P., & Riener, R. (2013). Quantifying the Human Likeness of a Humanoid Robot. *International Journal of Social Robotics*, 5, 263–276. <https://doi.org/10.1007/s12369-012-0177-4>
- Winkielman, P., Schwarz, N., Fazendeiro, T., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. *The Psychology of Evaluation: Affective Processes in Cognition and Emotion*, 189, 217.
- Yamada, Y., Kawabe, T., & Ihaya, K. (2013). Categorization difficulty is associated with negative evaluation in the "uncanny valley" phenomenon. *Japanese Psychological Research*, 55(1), 20–32. <https://doi.org/10.1111/j.1468-5884.2012.00538.x>
- Yerkes, R. M., & Dodson, J. D. (1908). The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative Neurology and Psychology*, 18(5), 459–482. <https://doi.org/10.1002/cne.920180503>
- Zlotowski, J., & Bartneck, C. (2013). The inversion effect in HRI: Are robots perceived more like humans or objects? *Tokyo, Japan: 8th ACM/IEEE international conference on Human-robot interaction, 3-6-Mar 2013. Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction*, 365–372. <https://doi.org/10.1109/HRI.2013.6483611>
- 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>.