Chapter 10
Persistence of the Uncanny Valley

Jakub A. Złotowski, Hidenobu Sumioka, Shuichi Nishio, Dylan F. Glas, Christoph Bartneck and Hiroshi Ishiguro

Abstract In recent years, the uncanny valley theory has been heavily investigated by researchers from various fields. However, the videos and images used in these studies do not permit any human interaction with the uncanny objects. Therefore, in the field of human–robot interaction, it is still unclear what impact, if any, an uncanny-looking robot will have in the context of an interaction. In this paper, we describe an exploratory empirical study using a live interaction paradigm that involves repeated interactions with robots that differ in embodiment and their attitude toward humans. We find that both components of uncanniness investigated here (likeability and eeriness) can be affected by an interaction with a robot. The likeability of a robot is mainly affected by its attitude, and this effect is especially prominent for a machine-like robot. Merely repeating interactions is sufficient to reduce the degree of eeriness, irrespective of a robot's embodiment. As a result, we urge other researchers to investigate the uncanny valley theory in studies that involve actual human–robot interactions in order to fully understand the changing nature of this phenomenon.

Keywords Uncanny valley · Anthropomorphism · Human–robot interaction Multiple interactions · Eeriness · Likeability · Dehumanization

---

This chapter is a modified version of a previously published paper [1], edited to be comprehensive and fit with the context of this book.

J. A. Złotowski
CITEC Center of Excellence Cognitive Interaction Technology, Bielefeld University, Bielefeld, Germany
e-mail: jakub.zlotowski@pg.canterbury.ac.nz

H. Sumioka · S. Nishio (✉) · D. F. Glas · H. Ishiguro
Advanced Telecommunications Research Institute International, Keihanna Science City, Kyoto, Japan
e-mail: nishio@botransfer.org

C. Bartneck
Human Interface Technology Laboratory New Zealand, University of Canterbury, Christchurch, New Zealand

H. Ishiguro
Department of Systems Innovation, Graduate School of Engineering Science, Osaka University, Osaka, Japan

H. Ishiguro and F. Dalla Libera (eds.), *Geminoid Studies*,
https://doi.org/10.1007/978-981-10-8702-8_10

10.1 Introduction

The uncanny valley theory was originally presented by Mori [2] in relation to a prosthetic arm. In recent years, it has gathered considerable attention in the fields of robotics, virtual agents, and cognitive sciences, as well as in mass media. The uncanny valley hypothesis suggests a nonlinear relationship between a robot's anthropomorphism and affinity. It proposes that, by increasing the humanlike appearance of a robot, we can increase our affinity for it. However, when a robot's appearance becomes sufficiently humanlike, but still distinguishable, people's emotional reaction becomes strongly negative. Once the appearance of a robot becomes indistinguishable from a real human, the affinity reaches its optimum at the same level as for human beings. Furthermore, Mori suggested that the movement of a prosthetic arm compared with a static arm will amplify the emotional response.

The uncanny valley is often used to explain people's rejection of anthropomorphic robots and virtual agents both in science and popular media, where it was given as a reason for the failure of the computer-animated movie The Polar Express. However, despite its wide adoption, there is relatively little empirical proof supporting the existence of the uncanny valley [3], e.g., the initial empirical work by [4] and [5] indicated that humanlike-ness might not be the only factor influencing our perception of an object as eerie. Rendering style could be related to the uncanny valley for virtual agents [6]. Moreover, it might be necessary to consider the effects of not only realism, but also the abnormality of artificial human appearance in investigating the uncanny valley phenomenon [7, 8]. It has been found that a mismatch between appearance and voice can result in the uncanny valley effect [9]. Furthermore, a mismatch between the appearance and movement of an android leads to stronger brain activation in the anterior portion of the intraparietal sulcus [10], which could provide a neurological explanation for the uncanny valley. In contrast, [11] reported that realistic motion can improve acceptability, especially of characters classified in the deepest point of the valley. This goes against the original theory of [2], who suggested that motion will increase the uncanny effect. The uncanny valley has been reported for other primates, with monkeys looking at real faces and unrealistic synthetic faces longer than at realistic synthetic monkey faces [12].

10.1.1 Related Work

Several potential explanations for the uncanny valley have been proposed. Apart from the neurological explanation [10], other factors include empathy [13], perception of experience [14], threat avoidance [2], and terror management [15]. A mathematical model using a Bayesian representation of categorical perception has been developed to explain how stimuli containing conflicting cues can give rise to a perceptual tension at category boundaries that leads to the uncanny feeling [16]. However, empirical investigations of these categorical boundaries suggest that ambiguous morphs close to the human endpoint induce a positive affect rather than the negative reaction suggested by the uncanny valley hypothesis [17, 18]. Furthermore, [19] found that images of prosthetic hands with intermediate humanlike-ness produced the strongest feelings of eeriness, whereas within different categories of images, increased humanlike-ness was related with the lowest degree of eeriness.

Vast research efforts have been dedicated to studying the dimensions of the uncanny valley. In particular, the original Japanese term used by [2]—*Shinwankan*—is difficult to translate to English. Various studies have used different translations, such as familiarity [5], likeability [20], affinity [21], eeriness [22], and empathy [23]. This variation in terms might affect the comparability of the results. Moreover, the humanlike-ness axis of Mori's graph has been the subject of empirical investigation [24].

The shape of the graph representing the uncanny valley is disputed. In one study, toy robots and humanoids were preferred over humans [25]. The authors proposed that the relationship between humanlike-ness and likeability resembles a cliff rather than a valley, where even perfectly realistic anthropomorphic robots are liked less than toy robots or mechanoids. These results imply that building highly humanlike androids might be futile, as their chances of acceptance are worse than for machine-like robots. Another study [20] reported that a highly realistic robot (android) was liked as much as a human. Furthermore, they reported that an android's realistic motion did not decrease its likeability and questioned the existence of the uncanny valley. This result is in line with a study using virtual agents [11]. However, [22] pointed out that the scales used by Bartneck and colleagues were correlated with warmth and, as a result, with each other, which might have affected the results. Overall, there is a lack of agreement between different studies regarding the dimensions and shape of the uncanny valley, with indications that Mori's theory could be too simplistic to accurately depict the relationship between humanlike-ness and perception of a robot or virtual agent. Moreover, it is not clear whether this theory has any actual consequences for interaction.

10.1.2 Does the Uncanny Valley Affect Human–Robot Interaction?

Despite being a common research theme, the effect of the uncanny valley hypothesis on human–robot interaction (HRI) is unknown. Previous studies on the uncanny valley used either images or videos of different targets that were supposed to induce the uncomfortable, eerie feeling (the exception is the work of [20], which involved short-term HRI). However, these studies did not permit any interaction between participants and robots or virtual agents. To understand how the uncanny valley affects HRI, it is necessary to involve physically collocated robots, as their physical presence could be an important mediating factor [26]. Previous work suggests that people's attitudes toward robots change during interaction [27], but it has never been empirically shown that the uncanny feeling will persist.

Little is known about the lasting effect of the uncanny valley. It is implicitly assumed that this negative emotional response toward anthropomorphic technology will have enduring consequences and lead people to reject androids that are distinguishable from humans. As this assumption has never been verified, it is important to consider an alternative hypothesis in which the uncanny valley only leads to a negative emotional response when the target is novel, and where the feeling of eeriness disappears during the course of HRI. It is possible that the affective habituation caused by repeated interactions will allow people to become accustomed to a machine that looks almost human, but is still not a perfect copy. Furthermore, the uncanny valley effect might decrease when an android interacts with a human in a friendly way. If this is the case, the effects of the uncanny valley on HRI might be limited to the pre-interaction phase.

10.1.3 Research Questions

There is some empirical evidence suggesting only a short-term effect of the uncanny valley. In a study conducted during the ARS Electronica festival, visitors who had interacted with an android were interviewed afterward. The majority did not report an uncanny feeling [28, 29]. As this study had the form of an open interview that allowed people to talk freely about their experience, only a qualitative analysis was possible. Therefore, it is important to quantitatively show whether the uncanny feeling is experienced less during and after interaction with an android. Secondly, the analysis of the uncanny valley phenomenon with virtual agents indicates that there could be a relation between knowing an agent (previous exposure) and the uncanny discomfort experienced by people exposed to it [30]. Lower levels of previous exposure to an agent were related with higher discomfort.

Moreover, there are psychological theories that can suggest a relation between repeated exposure to a stimuli and the uncanny valley hypothesis: mere exposure effect and affective habituation. It has been shown that mere exposure to a neutral stimulus leads to an increased positive affect toward it [31]. On the other hand, for strongly positive or negative stimuli, the intensity of the reaction decreases after multiple exposures. This process is called affective habituation [32].

The relationship between attraction and familiarity in interpersonal relations is well documented. Positive relationships are a result of frequent face-to-face contact [33]. However, if the person was initially disliked, greater familiarity would lead to greater dislike of that person [33]. This finding is consistent with other studies [34] that found repeated exposure to an unpleasant stimulus does not increase its likeability. Moreover, people are rated more positively by those whom they see more frequently [35] and express a stronger liking toward those whose ideas they have been exposed to for longer [36].

Four explanations have been proposed for the familiarity principle of attraction. Firstly, repeated exposure leads to increased processing fluency [37], which on its own is affectively positive [38]. Secondly, novel stimuli can produce uncertainty and negative reactions that diminish after a stimulus is found to be harmless [39]. Thirdly, as a result of classical conditioning, because most interactions are not aversive and mildly positive, others with whom people interact more often become paired with a positive affect [40, 41]. Fourthly, building on the previous explanation, repeated exposure creates an opportunity for interaction, and these interactions are more likely to lead to rewarding social experiences [41, 42].

The mere exposure effect does not require interaction, but exposure is sufficient for it to occur and it has been reported for various types of stimuli [43]. Although [44] proposed that familiarity leads to dislike in real interpersonal relations, because additional information about others makes them less similar to oneself, a live interaction paradigm showed that two previously unacquainted people exhibit a positive affect with increased familiarity [42].

In relation with the uncanny valley, it is possible that extreme stimuli weaken the affective reaction as people become more familiar because of affective habituation. However, for stimuli that are initially neutral, increased exposure could produce a more positive affect as a result of the mere exposure effect.

This study is the first exploratory work to investigate the effect of a robot's attitude and multiple interactions on the uncanny valley phenomenon by applying a live interaction paradigm in which actual HRI occurs. In particular, we focus on two aspects of interaction that could affect the uncanniness of a robot: (i) the number of interactions and (ii) the robot's attitude toward humans. Moreover, we have chosen two of the most common components representing the y-axis of the uncanny valley graph, *likeability* and *eeriness*, as they could be differently influenced by different aspects of HRI.

Likeability is an important factor affecting human–human relationships. For long-term HRI, it is expected to play an equally important role. Multiple factors affect human–human liking. One of the most important factors is a history of interaction with a specific person. In particular, we tend to like those with whom we have positive rather than negative interactions [45]. Moreover, the perception of a robot can be affected by its behavior [46]. Both positive and negative behavior have been anthropomorphized in robots, but people had more mechanistic conceptions for an impolite robot than for a positively behaving robot [27]. A robot that has a positive attitude toward humans could increase its likeability, as the classical conditioning explanation of the mere exposure effect would suggest. Similarly, an unfriendly robot could be liked less than it was before an interaction began. However, it is possible that an embodiment of a robot will play a role in affecting how strong an effect its behavior will have on its likeability. Thus, we hypothesize that:

H_{1a}: *A friendly robot's likeability will increase with repeated interactions.*

H_{1b}: *An unfriendly robot's likeability will decrease with repeated interactions.*

We also believe that previous exposure to a robot, irrespective of its behavior, is more important to its perceived eeriness. Eerie robots could produce affective habituation, and the initial strong negative emotional response will weaken with increased exposure. Similarly, for a robot that was initially perceived as neutral, repeated interactions may positively increase the affective perception according to the mere exposure effect.

In addition to looking at explicit measures such as self-reports, we investigate implicit attitudes toward humanlike robots. Implicit measures assess automatic reactions that are not consciously controllable [47] and are incrementally valid [48]. In addition, implicit measures complement rather than replace explicit measures, as they measure different aspects of the investigated attitude [49, 50]. Therefore, we also measured the perceived eeriness of the robots implicitly. Thus, our next hypotheses are:

H_{2a}: *Repeated interactions with a robot will reduce its explicit perceived eeriness.*

H_{2b}: *Repeated interactions with a robot will reduce its implicit perceived eeriness.*

Recent work in HRI indicates that it might be necessary to consider anthropomorphism as a multidimensional rather than unidimensional phenomenon [51]. These dimensions come from work on dehumanization—the process of depriving others of human qualities. It has been proposed that there are two distinct senses of humanness [52]: Human Uniqueness (HU) and Human Nature (HN). HU characteristics reflect socialization and distinguish humans from animals, e.g., intelligence, intentionality, or secondary emotions. HN covers inborn biological dispositions that distinguish humans from automata, e.g., warmth, sociability, or primary emotions. Anthropomorphism of a robot is not fixed and changes during an interaction [27]. To date, it has not been determined whether HU and HN dimensions of humanness attributed to a robot are also affected by the number of interactions or remain constant. In addition, previous work has indicated that the dimensions of mind attribution might be responsible for the uncanny valley phenomenon [14]. In particular, machines that are perceived as capable of experience but not agency are more uncanny. The dimensions of mind attribution and humanness are closely related [53]: agency reflects HU and experience reflects HN. Thus, our final hypothesis is:

H_3: *HN traits, but not HU traits, are related to a robot's perceived eeriness and likeability.*

10.2 Materials and Methods

Our study was conducted using a $2 \times 2 \times 3$ mixed experimental design in which a robot's embodiment (humanlike vs. machinelike) and attitude (positive vs. negative) were between-subjects factors and the number of interactions (Interaction I vs. Interaction II vs. Interaction III) was a within-subjects factor. We have explicitly measured a robot's perceived eeriness, anthropomorphism, likeability, and HN and HU dimensions of humanness. Furthermore, we used the Brief Implicit Association Test (BIAT) [54] as an implicit measurement tool of eeriness. This is a computer-based program that requires participants to classify a series of words into specified categories and measures the strength of the association between these concepts and attributes using participants reaction times.

10.2.1 Participants

Sixty native Japanese speakers were recruited by a recruitment agency for the study. The recruitment agency for part- and full-time student jobs posted a message on its website about the possibility of participating in a study involving a robot. Participants were paid ¥2000 for their time. All participants were undergraduate students of various universities and departments located in Kansai. Participants who had previously participated in a study involving one of the robots were excluded from selection. A software failure meant that the data from two participants were corrupted. Therefore, we had to exclude these data from the analysis. Of the remaining 58 participants, 26 are female and 32 are male. Their ages ranged from 18–36 years with a mean of 21.47. The study took place at the Advanced Telecommunications Research (ATR) Institute International. Adequate ethical approval was obtained from the ATR Ethics Committee and informed consent forms were signed by the participants.

10.2.2 Materials and Apparatus

All the implicit and explicit measurements were conducted using PsychoPy v1.78 running on a laptop. Participants interacted either with Geminoid HI-2 or Robovie R2. Geminoid HI-2 is the second generation of androids built as a copy of a real human (see Fig. 10.1). Geminoid is indistinguishable from a human being for several seconds. Once people realize its slight imperfections, they have a negative feeling [55, 56]. Robovie R2 is a machinelike robot that has some human features, such as a head and hands. Therefore, Geminoid HI-2 represents a robot that is near the deepest point of the uncanny valley, whereas the humanlike features of Robovie R2 should make it highly likeable [56]. Furthermore, as the uncanny valley can also be caused by a mismatch between appearance and voice or movement (e.g., [9, 10]), we ensured that the Geminoid HI-2 fell into the valley by using a synthetic childlike voice and machinelike jerky movement that does not fit the appearance of a male adult. The same movements and voice were used for Robovie R2, in which no mismatch occurs. During HRI, both robots expressed idle motions that were added to increase their degree of animation. Geminoid HI-2 exhibited movement resembling blinking and breathing, as well as idle movements of its hands and synchronization of its lips to its speech. As Robovie R2 does not have a mouth, identical idle behavior was not possible. Therefore, we implemented slight head and hand motions during speech.

The experiment took place in a room that had been divided into two parts separated by a folding screen to block the view (see Fig. 10.2). A robot was placed in the experimental space, and all HRIs occurred there. In the measurement space, participants watched an introduction video that explained the order of the experiment and filled out all the questionnaires on a laptop. This ensured that participants did not need to judge the robot in its presence, which could have affected the results.

Fig. 10.1 Geminoid HI-2 and a participant

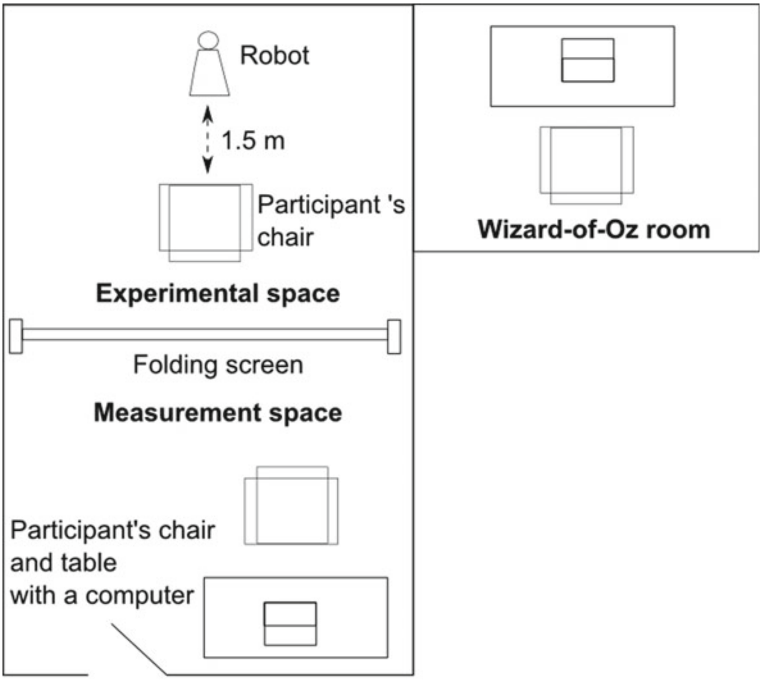

Fig. 10.2 Diagram of the experimental and measurement spaces and Wizard-of-Oz room

The experimental space was equipped with cameras and the robot's behaviors were controlled by a Wizard-of-Oz sitting in another room.

10.2.3 Procedure

We used a live interaction paradigm. Participants were first shown an introduction video that explained the experimental procedure. They were told that the study involved creative and persuasive talking and that they would need to convince a robot to give them a job based on a CV (an identical CV was given to all participants). The experimenter ensured that participants understood the instructions and sat them at a computer. During all HRIs and questionnaires, the experimenter left the participant alone in the room. The experiment was divided into four phases: pre-interaction video, Interaction I, Interaction II, and Interaction III.

Although we ensured that none of the participants had previously interacted or participated in an experiment with the specific robot to which they were assigned, it was still possible that they had seen the robot elsewhere. In particular, in Japan, it is common to see robots used in experiments on various TV programs. Therefore, to minimize the differences in potential prior exposure, participants were asked to watch a short video (∼15 s) in which a robot (either Robovie R2 or Geminoid HI-2) introduced itself and its capabilities. The dialogue was identical for both robots. After the video, participants performed the BIAT and filled out all a questionnaire.

During Interaction I, participants were taken to the experimental room and sat 1.5 m in front of a robot (see Fig. 10.3). They were told to have a small conversation with the robot to become familiar before the actual job interview began. The robot was introduced as *Robo*. During this conversation, the robot asked participants three neutral questions (e.g., "Is it cold today?" or "Where do you come from?"). After this short conversation, participants were asked to fill out the same questionnaire as before.

In Interaction II, the experimenter provided a short job description for which the participant was instructed to apply. Participants were asked to apply for positions as an engineer and bank manager. The order of interviews was counterbalanced between Interaction II and III. Furthermore, a participant received a CV of a person whom she was supposed to imitate during the interview. The CVs were identical for all participants, but the gender of the applicant was always the same as the real gender of the participant. Participants were asked to use this CV as the basis for their responses, but they could invent the information required to answer the questions. To motivate participants to perform the task as well as possible, they were informed by the experimenter that they would be paid an additional sum if they secured the job. They were given 5 min to prepare for the interview, and then the experimenter collected the CVs and job description sheets and brought the participant to the robot.

The interview began with the robot briefly describing the company and job position for which the participant was applying. After the introduction, the participant was asked three job interview questions. The questions were generic and common

Fig. 10.3 Experimental setup. Participant sitting in front of Robovie R2 during interaction

for job interviews, e.g., "Please tell me about yourself" or "What is your biggest weakness?" While the participant was responding, the robot provided feedback using nonlexical conversation sounds and nonverbal communication. In the positive condition, it either nodded or nodded and uttered "Un" (a Japanese expression signaling agreement with the speaker). In the negative condition, it either shook its head or nodded its head and uttered "Asso" (a Japanese expression indicating lack of interest in what the speaker says that is rather rude). This feedback was initiated by the Wizard as appropriate in the natural flow of the conversation, e.g., when a participant paused to think about her response.

After each question, the robot thanked the participant and asked the next question. After the third question, the robot informed the participant that it would announce its decision about whether to give the job to the participant later (in fact, the decision was never announced). Although the outcome was not announced directly to the participants, some hints were provided in each condition. In the positive condition, the robot hinted approval of what the participant said during the interview. In the negative condition, it was not particularly pleased with a participant's responses, suggesting they consider applying elsewhere. At that point, participants were asked to fill in a questionnaire for the third time. This time, multiple dummy questions regarding the interview were included. Interaction III was identical to Interaction II, but the CVs, job positions, and questions asked by the robot were different. participants were permitted to answer each of the questions freely and we did not measure the duration of the interactions. The whole procedure took approximately 1 h.

10.2.4 Measurements

In the experiment, we used several questionnaires and BIAT [54] as dependent measures. We explicitly measured the robots' perceived eeriness and anthropomorphism on 5-point Likert scales derived from [22]. Moreover, their likeability was measured using the corresponding Godspeed scale [57] (range 1–5). To establish the relationship between multidimensional anthropomorphism and the uncanny valley, we measured two dimensions of anthropomorphism, HN and HU, on scales developed by [58]. Both dimensions had 10 items and were measured on a scale from 1 (not at all) to 7 (very much) (e.g., "The *Robo* is… shallow"). This experiment was part of a bigger study that involved additional self-report scales that were collected at the same time and are not reported here. We used a validated version of the likeability scale in Japanese. Perceived eeriness, anthropomorphism, HN, and HU were only available in English. Therefore, we conducted a back-translation process to obtain their Japanese versions. We calculated the reliability of each scale separately for each interaction round using Cronbach's α. According to [59], Cronbach's $\alpha > 0.6$ is acceptable for newly developed scales for research purposes. Based on this threshold, all the scales (apart from HU) were adequately reliable. The lowest Cronbach's α values during any of the three measurements were as follows: likeability $\alpha = 0.83$, perceived eeriness $\alpha = 0.62$, anthropomorphism $\alpha = 0.88$, HN $\alpha = 0.65$, and HU $\alpha = 0.54$. The low reliability of the HU scale indicates that the HU results should be interpreted with great caution.

Furthermore, we used BIAT [54] as a computer-based implicit measurement tool of eeriness. BIATs involve participants classifying a series of words into superordinate categories. The task involved combining a concept classification (*"Robo"* vs. "Human") with an attribute classification ("Eeriness" vs. "Non-eeriness"). We were interested in measuring the strength of association between *"Robo"* and "Eeriness."

In the BIAT, only two categories at a time were displayed on the screen and a total of three categories were evaluated ("Interview Robot *Robo*," "Human," and "Eeriness"). The fourth category ("Non-eeriness") is called non-focal and was used as a distractor (attribute word that does not belong to the categories being evaluated in a specific block) for "Eeriness." The other two categories ("Interview Robot *Robo*" and "Human") were used as distractors for each other. There were two blocks of 16 trials each that were repeated four times. The following stimuli were used: "Interview Robot *Robo*" (Automaton, Machine, Robot, Artificial), "Human" (Person, Natural, Mankind, Real), "Eeriness" (Eerie, Freaky, Spine-tingling, Shocking), and "Non-eeriness" (Reassuring, Numbing, Uninspiring, Boring).

At the beginning of BIAT, participants were presented with two of the categories being evaluated (e.g., "Interview Robot *Robo*" and "Eeriness") and the words that belong to each of these categories. During the actual classification task, these categories were displayed in the top part of the screen. At the center of the screen, a series of words were shown that either belong to these categories or do not (see Fig. 10.4). Participants were asked to respond as quickly as possible by pressing "K" if the word belongs to either of the categories or "D" if not. As an example, if the

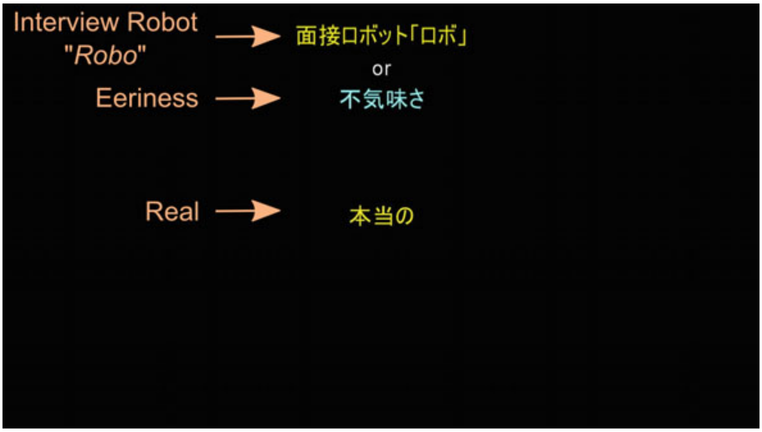

Fig. 10.4 Screenshot from the BIAT with English annotations. Two classification concepts ("Interview Robot *Robo*" and "Eeriness") and an attribute word ("Real") are being classified by a participant

categories were "Human" and "Eeriness," a participant should press "K" if the target word is "Mankind" or "Freaky," but should press "D" if the word is "Artificial" or "Reassuring." If a participant misclassified a word, a red cross appeared on the screen until the correct key was pressed.

The total time from the word appearing until the correct key was pressed was calculated with millisecond precision. These times were used to establish the strength of association between the categories. When an association between two categories is stronger, participants should be able to make their choices faster than for a pair of categories that are implicitly not associated with each other. The order of the BIATs was randomized and the order of the blocks was counterbalanced.

10.3 Results

In the first step of the analyses, we looked at the explicit and implicit measures. We then looked at the relationship between these different dependent measures. To analyze the data, we conducted a series of three-way ANOVAs with embodiment and attitude as between-subjects factors and number of interactions as a within-subjects factor. The assumptions of all statistical tests were met unless otherwise specified.

10.3.1 Likeability

First, we looked at the likeability, and in particular how a robot's attitude can affect it during HRI (Fig. 10.5). Because the assumption of a normal distribution was violated for parametric testing for anthropomorphism, we employed a permutation test with three factors using the aovp function with 1000 iterations from the lmPerm package [60] in R [61]. Likeability was significantly affected by the robots' attitude, $p = 0.001$. Positively behaving robots (M = 3.82, $\sigma = 0.67$) were liked more than negatively behaving robots (M = 3.24, $\sigma = 0.9$). Moreover, we found a statistically significant effect of embodiment with probability $p = 0.01$. Robovie R2 (M = 3.7, $\sigma = 0.88$) was liked more than Geminoid HI-2 (M = 3.37, $\sigma = 0.78$). In addition, we found a marginally significant interaction effect between embodiment and attitude, $p = 0.07$. Robovie R2 was more liked when it behaved positively (M = 4.15, $\sigma =$ 0.54) than negatively (M = 3.26, $\sigma = 0.94$), $p < 0.001$. However, the attitude of Geminoid HI-2 did not significantly affect its perceived likeability.

Furthermore, we found a statistically significant interaction effect between the robots' attitude and number of interactions, $p < 0.001$. During Interaction I, a robot's attitude did not affect its likeability. However, during Interaction II, a robot's positive attitude (M = 3.86, $\sigma = 0.66$) increased its likeability compared with the negative attitude (M = 2.93, $\sigma = 0.98$), $p < 0.001$. Similarly, during Interaction III, a robot's positive attitude (M = 3.97, $\sigma = 0.69$) resulted in higher likeability than a negatively behaving robot (M = 3.2, $\sigma = 0.94$), $p < 0.001$. The interaction effect between embodiment and measurement was significant with $p < 0.001$. The difference was only observed during Interaction I, when Robovie R2 (M = 3.9, $\sigma = 0.56$) was liked more than Geminoid HI-2 (M = 3.34, $\sigma = 0.61$).

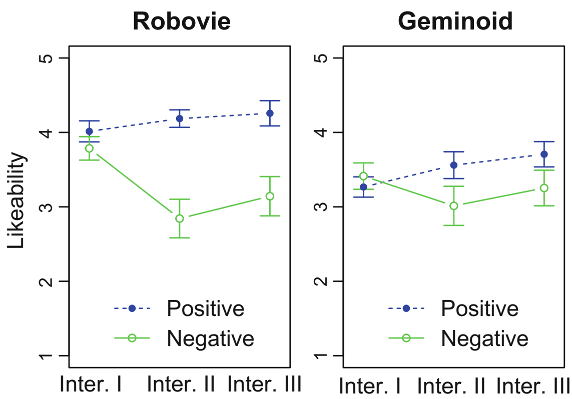

Fig. 10.5 Effect of three factors on likeability. The rating of likeability based on attitude and interaction round, grouped by robot type

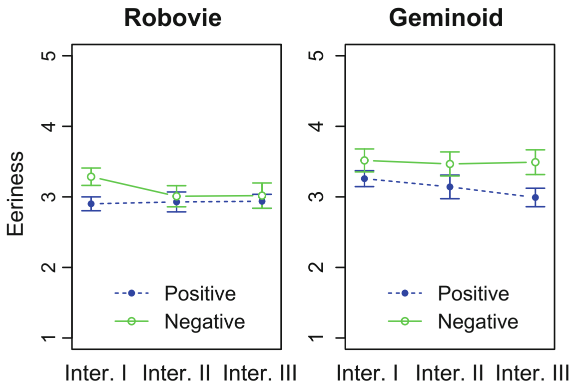

Fig. 10.6 Effect of three factors on explicit eeriness. The rating of explicit eeriness based on attitude and interaction round, grouped by robot type

10.3.2 Eeriness

The second component of the uncanny valley—eeriness—was measured explicitly and implicitly (Fig. 10.6). We are interested in establishing the effect of repeated interactions on a robot's perceived eeriness. The explicit measure of eeriness showed the main effect of embodiment to be statistically significant, $F(1{,}54) = 5.14$, $p = 0.03$, $\eta_G^2 = 0.07$. Geminoid HI-2 ($M = 3.31$, $\sigma = 0.62$) was perceived as significantly more eerie than Robovie R2 ($M = 3.01$, $\sigma = 0.51$). Moreover, there was a significant main effect relating to attitude, $F(1{,}54) = 4.27$, $p = 0.04$, $\eta_G^2 = 0.06$. A robot behaving negatively ($M = 3.3$, $\sigma = 0.64$) was perceived as more eerie than one behaving positively ($M = 3.03$, $\sigma = 0.49$). In addition, there was a main effect relating to the number of interactions, $F(2{,}108) = 3.1$, $p = 0.05$, $\eta_G^2 = 0.01$. Post-hoc tests using the Bonferroni correction revealed that, with marginal significance, participants rated the robots as more eerie after Interaction I ($M = 3.25$, $\sigma = 0.52$) than after Interaction III ($M = 3.11$, $\sigma = 0.6$), $p = 0.08$.

Apart from the explicit eeriness, we also measured implicit eeriness. In the BIAT, shorter response times indicate a stronger association between categories. Thus, an increased time would indicate a weaker association between a robot and eeriness. However, the reduced response time with the increased number of interactions could also be due to participants improving at the task itself. Therefore, we have transformed the reaction times to z-scores within each interaction round, enabling the comparison of results between interactions. A three-way ANOVA with embodiment and attitude as between-subjects factors and the number of interactions as a within-subjects factor did not indicate any statistically significant main or interaction effects.

10.3.3 Anthropomorphism

We then looked at one- and two-dimensional measures of anthropomorphism. It was expected that there would be a main effect relating to a robot's embodiment, and in particular that Geminoid HI-2 would be perceived as more humanlike than Robovie R2. As the assumption of a normal distribution was violated for parametric testing for anthropomorphism, we employed a permutation test with three factors using the aovp function with 1000 iterations from the lmPerm package [60] in R [61]. We found a marginally statistically significant main effect of embodiment with probability $p = 0.08$ (see Fig. 10.7). Geminoid HI-2 (M = 2.47, $\sigma = 1.1$) was more anthropomorphic than Robovie R2 (M = 2.17, $\sigma = 0.92$). Moreover, we found a significant interaction effect between the robots' attitude and number of interactions with probability $p < 0.001$. Only during Interaction III did a robot's positive attitude (M = 2.63, $\sigma = 1.07$) result in higher likeability compared with a negatively behaving robot (M = 2.11, $\sigma = 1.02$), $p = 0.05$.

We then proceeded to the two-dimensional measurement of anthropomorphism to investigate its relation with the uncanny valley. The results related to the model of anthropomorphism proposed by [51] will be discussed in another paper. In line with previous research, we did not find statistically significant main or interaction effects for the HU dimension (see Fig. 10.8).

However, we found a main effect relating to embodiment, with $F(1,54) = 5.13$, $p = 0.03$, $\eta_G^2 = 0.07$ on the HN dimension (see Fig. 10.9). Robovie R2 (M = 3.16, $\sigma = 0.77$) was attributed with more HN traits than Geminoid HI-2 (M = 2.74, $\sigma = 0.85$). In addition, there was a significant main effect of attitude, $F(1,54) = 8.46$, $p = 0.005$, $\eta_G^2 = 0.12$. Robots with positive attitude (M = 3.21, $\sigma = 0.74$) were attributed as having more HN than those with a negative attitude (M = 2.67, $\sigma = 0.85$). There was also a significant main effect relating to the number of interactions, $F(2,108) = 7.39$, $p = 0.001$, $\eta_G^2 = 0.02$. Post-hoc tests using the Bonferroni correction for the family-wise error revealed that the robots were attributed more HN traits after Interaction I (M = 3.4, $\sigma = 0.77$) than after Interactions II (M = 2.88, $\sigma = 0.87$), $p = 0.02$, or

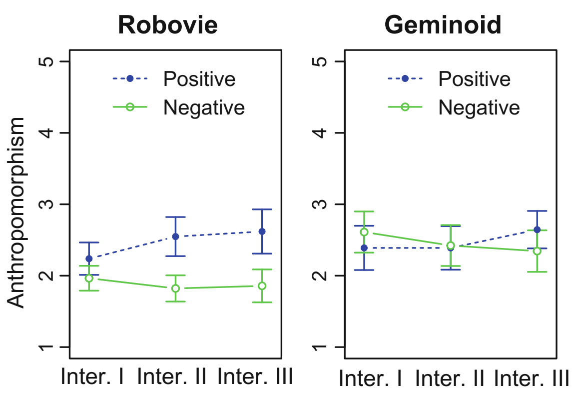

Fig. 10.7 Effect of three factors on anthropomorphism. The rating of anthropomorphism based on attitude and interaction round, grouped by robot type

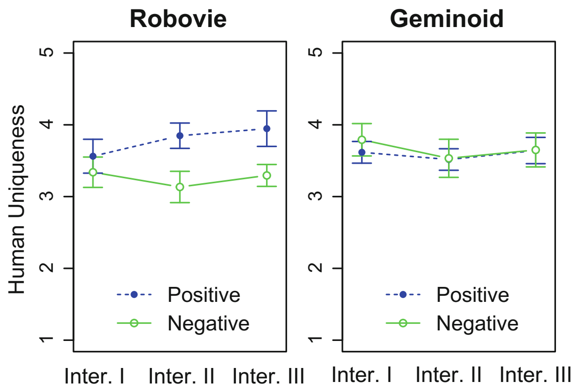

Fig. 10.8 Effect of three factors on Human Uniqueness. The rating of Human Uniqueness based on attitude and interaction round, grouped by robot type

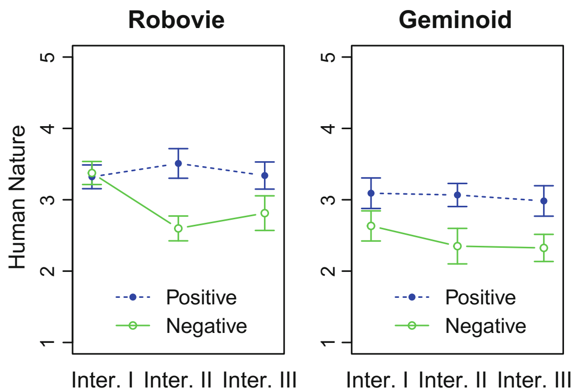

Fig. 10.9 Effect of three factors on Human Nature. The rating of Human Nature based on attitude and interaction round, grouped by robot type

III (M = 2.86, $\sigma = 0.86$), $p = 0.02$. Furthermore, there was a significant interaction effect between attitude and number of interactions, F(2,108) = 9.8, $p < 0.001$, $\eta_G^2 = 0.03$. Attitude only produced a significant effect in Interactions II [F(1,56) = 15.82, $p < 0.001$, $\eta_G^2 = 0.22$] and III [F(1,56) = 7.75, $p = 0.007$, $\eta_G^2 = 0.12$].

10.3.4 Relationship Between the Uncanny Valley and HRI Factors

Next, we looked at the relationship between the different dependent variables used in this study to establish how the uncanny valley is related to factors that are important for HRI. The correlations between likeability, eeriness, one- and two-dimensional anthropomorphism are presented in Table 10.1.

The following convention was used to determine the effect size of Pearson's r coefficient: small ($0.1 \le |r| < 0.3$), medium ($0.3 \le |r| < 0.5$), large ($0.5 \le |r|$).

Table 10.1 Correlation between dependent measures using Pearson's r coefficient

	Likeability	Eeriness	Anthropomorphism	HU	HN
Likeability		−0.11	0.4***	0.31***	0.51***
Eeriness	−0.11		0.05	0.15*	0.13
Anthropomorphism	0.4***	0.05		0.15*	0.15*
HU	0.31***	0.15*	0.15*		0.35***
HN	0.51***	0.13	0.15*	0.35***	•

*$p < 0.05$, ***$p < 0.001$

A correlation with a large effect size was observed between likeability and HN. Furthermore, likeability had a medium effect size correlation with anthropomorphism and HU. Eeriness and likeability were not correlated. Eeriness was significantly correlated with HU with a small effect size.

10.4 Discussion

In this study, we investigated the effect of repeated interactions and a robot's attitude on the uncanny valley phenomenon using a live interaction paradigm. In particular, we investigated the impact of these factors on a robot's likeability, as well as explicit and implicit measures of perceived eeriness. Explicit eeriness and likeability were not significantly correlated, which indicates that they measure different aspects of the uncanny valley. Although that might initially seem like an unexpected and counterintuitive finding, there are examples showing that a negative correlation between eeriness and likeability is not necessary. People can dislike other people, but at the same time do not perceive them as eerie. However, there are also cases when eeriness is desirable, e.g., people who like to watch horror movies that might involve eerie creatures. Therefore, measuring both of aspects offers a richer picture than considering only one.

The analysis of likeability showed that the more machinelike robot (Robovie R2) was liked more than the highly humanlike Geminoid HI-2. Moreover, a robot's attitude toward a human interaction partner could be used to affect its likeability, with friendly robots being liked more than unfriendly robots. However, the effect of a robot's attitude is not independent of its embodiment. The interaction effect between embodiment and attitude is especially profound in the case of a more machinelike robot. Although Robovie R2's positive behavior resulted in a small increase in likeability, its negative attitude resulted in a drop in likeability to a level similar to that observed for the negatively behaving Geminoid HI-2. In case of the latter robot, its attitude did not significantly affect its likeability. Thus, H_{1a} and H_{1b} are not supported.

These results seem to indicate that a robot that is perceived as uncanny is not able to affect its likeability through positive or negative interactions. In that sense, its

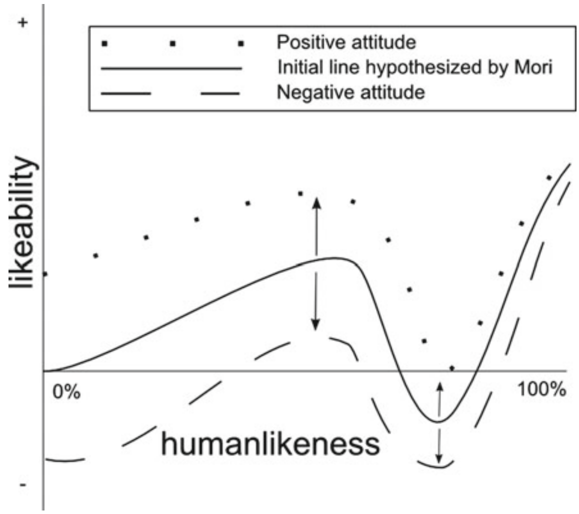

Fig. 10.10 Hypothesized effect of robots attitude on the uncanny valley. Likeability of a robot will increase with its positive attitude toward a human interaction partner or decrease with its negative attitude. The less humanlike the robot, the stronger the effect

lower likeability is persistent. On the other hand, the impact of a machinelike robot's attitude is much greater, especially when it behaves negatively. The less humanlike a robot is, the stronger that effect could be. A hypothetical relationship between humanlike-ness and the effect of a robot's attitude on its likeability is presented in Fig. 10.10.

These findings on likeability provide a new perspective on the psychological theories related with the effect of familiarity. In particular, the results are more consistent with the mere exposure effect than affective habituation. As suggested by the work of [33, 34], greater familiarity with an unpleasant stimulus did not enhance the likeability of Geminoid HI-2, which contradicts the affective habituation theory. However, in the case of the more neutral stimulus (Robovie R2), its behavior during interactions affected its likeability. This supports the explanation of the familiarity effect proposed by [41, 42], in which repeated exposure creates opportunities for interaction and those interactions that are positive will lead to a favorable impression of a person, or in this case a robot. Therefore, in live HRI, mere exposure to a robot is insufficient to induce a positive affect toward it, and a positively toned interaction is required. However, in the case of a strongly unpleasant robot, even positive behavior can be insufficient to enhance its likeability.

Looking at the eeriness aspect of the uncanny valley, we found that Geminoid HI-2 was rated as more eerie than Robovie R2. However, more interestingly, we observed that after the last interaction, both robots were perceived as being less eerie than after the first interaction. This indicates that perceived eeriness decreases with exposure to a robot. Moreover, this reduction is the same between robots that initially had different levels of eeriness, thus supporting H_{2a}. Therefore, although perceived eeriness of a highly anthropomorphic robot can be decreased by merely increasing

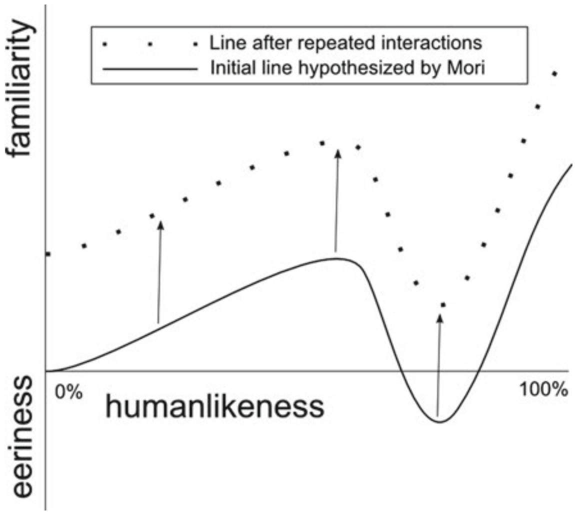

Fig. 10.11 Hypothesized effect of repeated HRIs on the uncanny valley. The reduction of a robot's eeriness is relatively constant, regardless of the level of its humanlike-ness

the number of HRIs, the gap between machinelike and humanlike robots remains relatively constant (see Fig. 10.11).

As both robots were perceived as less eerie after multiple interactions, it is possible that the mere exposure effect [31] and affective habituation [32] were both involved. Geminoid HI-2 was initially perceived as an extremely eerie robot. In this case, it is possible that affective habituation occurred and the affective reaction became weaker with increased exposure. On the other hand, for an initially neutral-looking robot (Robovie R2), additional exposure could not decrease its eeriness, irrespective of behavior. Therefore, the effect of familiarity on the perceived eeriness worked differently than for likeability, where a robot's positive behavior led to a favorable impression. If the familiarity effect of attraction also affects perceived eeriness, the explanation that a positive interaction is required is not supported. The more probable explanations for the results obtained for Robovie R2 are that a novel stimulus that initially fosters wary reactions is found to be benign after repeated interactions [39] or that additional exposure increases a robot's processing fluency [37] as its appearance becomes more familiar. As increased processing fluency has a positive effect, it is possible that this processing affect is then transferred to the robot, leading to a decrease in perceived eeriness. However, as the goal of this study was not to investigate which processes are responsible for the familiarity impact on attraction, future studies should verify whether any of these explanations are correct.

Our findings for both likeability and perceived eeriness are relevant for HRI designers. The likeability of robot is affected by its behavior. However, the effect is much stronger in the case of a more machinelike robot. In particular, a machinelike robot can swiftly stop being liked, despite its appearance, by employing negative behavior. It is much harder to increase the likeability of a robot that initially falls into the uncanny valley, as a friendly attitude is not sufficient make it likeable.

However, people quickly become used to the unfamiliar appearance of a robot. In our study, three short interactions were sufficient to reduce its perceived eeriness. However, this reduction was not found to be stronger for the more anthropomorphic robot. Therefore, the relative difference in perceived eeriness between the robots remained at the same level. Nevertheless, in this study, we enhanced the eeriness of Geminoid HI-2 by creating a mismatch between its appearance, speech, and movement. If the only source of eeriness was its embodiment, it is possible that the effect of multiple interactions with Geminoid HI-2 would be more profound. It is also noteworthy that the perceived eeriness of Geminoid HI-2 after Interaction III reached the level of Robovie R2 after Interaction I. Therefore, Geminoid HI-2 only remained more eerie because the perceived eeriness of Robovie R2 had decreased. It is possible that a higher number of interactions, after the machinelike robot had reached its maximum familiarity, may allow a highly humanlike robot to reach the same level.

We also found that a negatively behaving robot was rated as more eerie than a positively behaving robot. However, this finding could be explained as a result of the HRI context used in this experiment. In Japanese culture, it is not typical for an interviewer to express lack of interest during a job interview in such an explicit and rude way as the robot did in this experiment. Such an attitude could have led the robot to be perceived as more eerie than if it had behaved in a way that is common during human–human job interviews.

The analysis of implicit eeriness using BIAT did not show any significant differences. Thus, H_{2b} is not supported. Therefore, in the current form, BIAT might not be optimally suited to measuring eeriness. We speculate that this result could be due to the weak association between a robot's category ("Interview Robot *Robo*") that was displayed on a screen and the specific robot with which the participants interacted. As implicit attitudes tend to change more slowly than explicit attitudes, it is possible that our manipulation was too weak to modify attitudes toward a specific robot. As a result, participants might have responded to the robot's category as being merely a representation of robots in general rather than their specific robotic interaction partner. In future studies, it might be beneficial to use a picture of a robot instead of a name as a representation of its category.

In line with previous research, the HU dimension of anthropomorphism was not significantly affected by the embodiment of a robot. Furthermore, the attribution of HN traits was affected by the embodiment and is therefore more relevant to the uncanny valley. Thus, H_3 is supported. However, in contrast with previous work [14], it was the less uncanny robot (Robovie R2) that was attributed with more HN. Despite this dimension having more impact on the uncanny valley, the relationship looks to be more complex than initially proposed. The biggest difference between the work of [14] and ours is the robots used in the experiments. In [14], a single robot that either had the back of its head visible or a humanlike face cover was used. The HN dimension is closely related with emotions, and a robot that has no face is not capable of expressing emotions with facial expressions. Therefore, it will be attributed with less capability to experience (HN). In our experiment, the default and fixed appearance of Robovie R2's face could be perceived as a smile. However, Geminoid HI-2 has a highly humanlike face that suggests that it can exhibit facial expressions. As a result, participants might have had higher expectations, but the robot's facial expression remained the same (and was rather stern) during the interactions. This could have been perceived as emotional coldness in the robot, leading participants to attribute less HN to Geminoid HI-2. Nevertheless, more research is needed to establish the relationship between HN and the uncanny valley. Furthermore, considering the inadequate reliability of the HU dimension, it is necessary to interpret these results with special care. It is possible that HU is a different construct in Japan than in Western cultures.

10.4.1 Limitations and Future Work

In our experiment, we used only two robots that differed in their level of anthropomorphism. An alternative explanation for our results could be that it is a robot's friendliness in appearance that is more important than humanlike-ness in fostering likeability. We cannot exclude the possibility that there are differences along some other dimensions reflected by appearance. It is possible that the interaction between embodiment and attitude would be reversed if we used a different pair of robots. In particular, Geminoid HI-2 has a stern-looking facial expression, whereas the design of Robovie R2 could be perceived as cute and friendly with its big, childlike head. The appearance of Robovie R2 could invoke expectations for it to behave positively, and the mismatch between these expectations and the actual behavior of the robot could result in a strong decrease in likeability. If a friendlier-looking android, e.g., Geminoid F, was used in the experiment instead of Geminoid HI-2, it is possible that we would have observed a similar pattern of reactions to its unfriendly behavior as for Robovie R2. However, the question remains as to why the opposite trend was not observed in the case of Geminoid HI-2's mismatched positive attitude. Therefore, future studies should also include qualitative data that could help to understand why people perceive robots as eerie or likeable. Moreover, there could be demographic factors such as age, gender, or educational background that work as moderators. The role of these factors on the uncanny valley has not yet been explored in sufficient detail.

The scale used for measuring anthropomorphism [22] in the uncanny valley experiments was developed in a study that involved only static images of robots. However, contrary to expectations, Robovie R2 and Geminoid HI-2 differ only marginally in terms of perceived humanlike-ness. Because previous work indicates that androids are perceived as more humanlike than machinelike robots (e.g., [22]), the small difference between the two robots in this study must be due to other factors than merely embodiment. To increase the uncanniness of Geminoid HI-2, we used a voice and movements that did not match its embodiment. However, the humanlike-ness scale may have been affected by this manipulation, as certain items do not only apply to the embodiment, e.g., items rated by the participants include "Artificial"-"Lifelike" or "Fake"-"Natural." As a result, our manipulation not only made Geminoid HI-2 more eerie, but also less humanlike than if only its embodiment was evaluated.

This finding also indicates that a robot's behavior can be a more important factor in anthropomorphism than its embodiment. A potential solution involves the development of a new scale of anthropomorphism that is not affected by the potential mismatch between a robot's embodiment and speech or movement. Alternatively, before investigating the uncanny valley in interactions, it would be possible to first rate a robot's humanlike-ness by presenting a static robot with no HRI.

Another limitation of this study is that participants were allowed to freely interact with the robot for as long as they wanted. Therefore, we did not consider the interaction duration in this study, but only the number of interactions. It is possible that participants who interacted with a positively behaving robot were encouraged by its positive feedback to provide more detailed answers for their questions, and interacted for longer as a result. This extended interaction could have increased the familiarity of the robot and reduced its eeriness. It is also possible that the duration of interactions was insufficient to lead to the affective habituation effect of an uncanny robot. The perceived eeriness of both robots was reduced as a result of repeated interactions. However, it is still possible that, after a higher number of interactions, the affective habituation effect would become stronger for the more eerie robot. A long-term study with highly anthropomorphic robots could answer this question. In particular, future experiments should involve longer interactions with a robot, with sessions spread over multiple days.

Future work should also consider the dynamic nature of anthropomorphism. The complexity and multifaceted nature of anthropomorphism highlights a potential challenge for investigating the uncanny valley in actual, long-term HRI, rather than using images or videos that focus only on a robot's embodiment. Previous work on the uncanny valley treated it as a static feature of a robot or virtual agent. However, [27] showed that a robot's anthropomorphism changes during HRI. The results of this study also illustrate that, at least in the case of Robovie R2, attitude affects perceived humanlike-ness. Mori's hypothesis does not accommodate for such a finding. Studies of the uncanny valley should recognize that both anthropomorphism and uncanniness can vary during HRI, and they might consider whether the uncanny valley should be investigated using the pre-interaction level of anthropomorphism based only on a robot's appearance or the level of anthropomorphism measured in HRI at the same point of time as measures of uncanniness.

This study was an exploratory work that, for the first time, investigated the uncanny valley in repeated HRIs. Our results show the potential benefits of researching the complexity of this phenomenon in studies that involve human interaction with a collocated robot. Nevertheless, at the same time, our results indicate that, to understand the impact of the uncanny valley on HRI, future research must go beyond picture- and video-based studies and enable people to interact with robots. The great majority of studies have tried to find the origin of this phenomenon. This is a worthy goal. However, until we can show that Mori's theory has any significant (long-term) impact on HRI, we risk spending resources on research into an artificial problem. In the end, it matters very little whether a picture of a robot is perceived as eerie or disliked if, during an actual interaction with the robot, this effect vanishes as a result of behavior or interaction-context factors being more prominent.

Acknowledgements This work was partially supported by Grant-in Aid for Scientific Research (S), KAKENHI (25220004) and JST CREST (Core Research of Evolutional Science and Technology) research promotion program "Creation of Human-Harmonized Information Technology for Convivial Society" Research Area. The authors would like to thank Kaiko Kuwamura, Daisuke Nakamichi, Junya Nakanishi, and Kurima Sakai for their help with data collection.

References

1. Złotowski, J.A., H. Sumioka, S. Nishio, D.F. Glas, C. Bartneck, and H. Ishiguro. 2015. Persistence of the uncanny valley: The influence of repeated interactions and a robot's attitude on its perception. *Frontiers in Psychology* 6: 883.
2. Mori, M. 1970. The uncanny valley. *Energy* 7 (4): 33–35.
3. Blow, M., K. Dautenhahn, A. Appleby, C. Nehaniv, and D. Lee. 2006. Perception of robot smiles and dimensions for human-robot interaction design. In *The 15th IEEE international symposium on robot and human interactive communication, 2006. ROMAN 2006*, 469–474.
4. Hanson, D. 2006. Exploring the aesthetic range for humanoid robots. In *Proceedings of the ICCS/CogSci-2006 long symposium: Toward social mechanisms of android science*, 39–42.
5. MacDorman, K.F. 2006. Subjective ratings of robot video clips for human likeness, familiarity, and eeriness: An exploration of the uncanny valley. In *ICCS/CogSci-2006 long symposium: Toward social mechanisms of android science*, 26–29.
6. McDonnell, R., M. Breidt, and H.H. Bälthoff. 2012. Render me real?: Investigating the effect of render style on the perception of animated virtual humans. *ACM Transactions on Graphics* 31 (4): 91:1–91:11.
7. MacDorman, K.F., R.D. Green, C.-C. Ho, and C.T. Koch. 2009. Too real for comfort? uncanny responses to computer generated faces. *Computers in Human Behavior* 25 (3): 695–710.
8. Seyama, J., and R.S. Nagayama. 2007. The uncanny valley: Effect of realism on the impression of artificial human faces. *Presence: Teleoperators and Virtual Environments* 16 (4): 337–351.
9. Mitchell, W.J., K.A. Szerszen, A.S. Lu, P.W. Schermerhorn, M. Scheutz, and K.F. MacDorman. 2011. A mismatch in the human realism of face and voice produces an uncanny valley. *i-Perception* 2 (1): 10–12.
10. Saygin, A.P., T. Chaminade, H. Ishiguro, J. Driver, and C. Frith. 2012. The thing that should not be: Predictive coding and the uncanny valley in perceiving human and humanoid robot actions. *Social Cognitive and Affective Neuroscience* 7 (4): 413–422.
11. Piwek, L., L.S. McKay, and F.E. Pollick. 2014. Empirical evaluation of the uncanny valley hypothesis fails to confirm the predicted effect of motion. *Cognition* 130 (3): 271–277.
12. Steckenfinger, S.A., and A.A. Ghazanfar. 2009. Monkey visual behavior falls into the uncanny valley. *Proceedings of the National Academy of Sciences* 106 (43): 18362–18366.
13. MacDorman, K.F., P. Srinivas, and H. Patel. 2013. The uncanny valley does not interfere with level 1 visual perspective taking. *Computers in Human Behavior* 29 (4): 16711685.
14. Gray, K., and D. Wegner. 2012. Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition* 125 (1): 125–130.
15. MacDorman, K.F., and H. Ishiguro. 2006. The uncanny advantage of using androids in cognitive and social science research. *Interaction Studies* 7 (3): 297–337.
16. Moore, R.K. 2012. A bayesian explanation of the 'Uncanny valley' effect and related psychological phenomena. *Scientific Reports* 2.
17. Cheetham, M., P. Suter, and L. Jancke. 2014. Perceptual discrimination difficulty and familiarity in the uncanny valley: More like a happy valley. *Frontiers in Psychology* 5.

18. Looser, C.E., and T. Wheatley. 2010. The tipping point of animacy how, when, and where we perceive life in a face. *Psychological Science* 21 (12): 1854–1862.
19. Poliakoff, E., N. Beach, R. Best, T. Howard, and E. Gowen. 2013. Can looking at a hand make your skin crawl? peering into the uncanny valley for hands. *Perception* 42 (9): 998–1000.
20. Bartneck, C., T. Kanda, H. Ishiguro, and N. Hagita. 2009. My robotic doppelganger—A critical look at the uncanny valley theory. In *18th IEEE international symposium on robot and human interactive communication, RO-MAN2009*, 269–276. IEEE.
21. Mori, M., K.F. MacDorman, and N. Kageki. 2012. The uncanny valley. *IEEE Robotics and Automation Magazine* 19 (2): 98–100.
22. Ho, C., and K. MacDorman. 2010. Revisiting the uncanny valley theory: Developing and validating an alternative to the godspeed indices. *Computers in Human Behavior* 26 (6): 1508–1518.
23. Misselhorn, C. 2009. Empathy with inanimate objects and the uncanny valley. *Minds and Machines* 19 (3): 345–359.
24. Cheetham, M., P. Suter, and L. Jancke. 2011. The human likeness dimension of the "Uncanny valley hypothesis": Behavioral and functional MRI findings. *Frontiers in Human Neuroscience* 5.
25. Bartneck, C., T. Kanda, H. Ishiguro, and N. Hagita. 2007. Is the uncanny valley an uncanny cliff? In *Proceedings—IEEE international workshop on robot and human interactive communication*, 368–373, Jeju, Republic of Korea.
26. Kiesler, S., A. Powers, S.R. Fussell, and C. Torrey. 2008. Anthropomorphic interactions with a robot and robot-like agent. *Social Cognition* 26 (2): 169–181.
27. Fussell, S.R., S. Kiesler, L.D. Setlock, and V. Yew. 2008. How people anthropomorphize robots. In *HRI 2008—Proceedings of the 3rd ACM/IEEE international conference on human-robot interaction: Living with robots*, 145–152, Amsterdam, Netherlands.
28. Becker-Asano, C., K. Ogawa, S. Nishio, and H. Ishiguro. 2010. Exploring the uncanny valley with geminoid HI-1 in a real-world application. In *Proceedings of the IADIS international conference on interfaces and human computer interaction 2010, IHCI, proceedings of the IADIS international conference on game and entertainment technologies 2010, part of the MCCSIS 2010*, 121–128, Freiburg, Germany.
29. von der Pütten, A.M., N.C. Krämer, C. Becker-Asano, and H. Ishiguro. 2011. An android in the field. In *Proceedings of the 6th international conference on Human-robot interaction, HRI '11*, 283–284, New York, NY, USA. ACM.
30. Dill, V., L.M. Flach, R. Hocevar, C. Lykawka, S.R. Musse, and M.S. Pinho. 2012. Evaluation of the uncanny valley in CG characters. In *12th international conference on intelligent virtual agents, IVA 2012, September 12, 2012–September 14, 2012*, vol. 7502 LNAI, 511–513. Springer.
31. Zajonc, R.B. 1968. Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology* 9 (2p2): 1.
32. Dijksterhuis, A., and P.K. Smith. 2002. Affective habituation: Subliminal exposure to extreme stimuli decreases their extremity. *Emotion* 2 (3): 203.
33. Ebbesen, E.B., G.L. Kjos, and V.J. Konečni. 1976. Spatial ecology: Its effects on the choice of friends and enemies. *Journal of Experimental Social Psychology* 12 (6): 505–518.
34. Perlman, D., and S. Oskamp. 1971. The effects of picture content and exposure frequency on evaluations of negroes and whites. *Journal of Experimental Social Psychology* 7 (5): 503–514.
35. Brockner, J., and W.C. Swap. 1976. Effects of repeated exposure and attitudinal similarity on self-disclosure and interpersonal attraction. *Journal of Personality and Social Psychology* 33 (5): 531–540.
36. Brickman, P., P. Meyer, and S. Fredd. 1975. Effects of varying exposure to another person with familiar or unfamiliar thought processes. *Journal of Experimental Social Psychology* 11 (3): 261–270.
37. Bornstein, R.F., and P.R. D'Agostino. 1994. The attribution and discounting of perceptual fluency: Preliminary tests of a perceptual fluency/attributional model of the mere exposure effect. *Social Cognition* 12 (2): 103–128.

38. Reber, R., P. Winkielman, and N. Schwarz. 1998. Effects of perceptual fluency on affective judgments. *Psychological Science* 9 (1): 45–48.
39. Lee, A.Y. 2001. The mere exposure effect: An uncertainty reduction explanation revisited. *Personality and Social Psychology Bulletin* 27 (10): 1255–1266.
40. Clark, L.A., and D. Watson. 1988. Mood and the mundane: Relations between daily life events and self-reported mood. *Journal of Personality and Social Psychology* 54 (2): 296–308.
41. Denrell, J. 2005. Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review* 112 (4): 951–978.
42. Reis, H.T., M.R. Maniaci, P.A. Caprariello, P.W. Eastwick, and E.J. Finkel. 2011. Familiarity does indeed promote attraction in live interaction. *Journal of Personality and Social Psychology* 101 (3): 557–570.
43. Bornstein, R.F. 1989. Exposure and affect: Overview and meta-analysis of research, 1968–1987. *Psychological Bulletin* 106 (2): 265–289.
44. Norton, M.I., J.H. Frost, and D. Ariely. 2007. Less is more: The lure of ambiguity, or why familiarity breeds contempt. *Journal of Personality and Social Psychology* 92 (1): 97–105.
45. Smith, E.R., and D.M. Mackie. 2007. *Social psychology*, 3rd ed, Jan. New York: Psychology Press.
46. Goetz, J., S. Kiesler, and A. Powers. 2003. Matching robot appearance and behavior to tasks to improve human-robot cooperation. In *ROMAN 2003. The 12th IEEE international workshop on robot and human interactive communication*, 55–60.
47. De Houwer, J., S. Teige-Mocigemba, A. Spruyt, and A. Moors. 2009. Implicit measures: A normative analysis and review. *Psychological Bulletin* 135 (3): 347–368.
48. Steffens, M.C., and S. Schulze. 2006. König. Predicting spontaneous big five behavior with implicit association tests. *European Journal of Psychological Assessment* 22 (1): 13–20.
49. Admoni, H., and B. Scassellati. 2012. A multi-category theory of intention. In *Proceedings of COGSCI 2012*, 1266–1271, Sapporo, Japan.
50. Gawronski, B. 2002. What does the implicit association test measure? a test of the convergent and discriminant validity of prejudice-related IATs. *Experimental Psychology (formerly Zeitschrift für Experimentelle Psychologie)* 49 (3): 171–180.
51. Złotowski, J., E. Strasser, and C. Bartneck. 2014. Dimensions of anthropomorphism: From humanness to humanlikeness. In *Proceedings of the 2014 ACM/IEEE international conference on human-robot interaction, HRI '14*, 66–73, New York, NY, USA. ACM.
52. Haslam, N. 2006. Dehumanization: An integrative review. *Personality and Social Psychology Review* 10 (3): 252–264.
53. Haslam, N., B. Bastian, S. Laham, and S. Loughnan. 2012. Humanness, dehumanization, and moral psychology.
54. Sriram, N., and A.G. Greenwald. 2009. The brief implicit association test. *Experimental Psychology (formerly Zeitschrift für Experimentelle Psychologie)* 56 (4): 283–294.
55. Ishiguro, H. 2006. Android science: Conscious and subconscious recognition. *Connection Science* 18 (4): 319–332.
56. Rosenthal-von der Pütten, A.M., and N.C. Krämer. 2014. How design characteristics of robots determine evaluation and uncanny valley related responses. *Computers in Human Behavior* 36: 422–439.
57. Bartneck, C., D. Kulic, E. Croft, and S. Zoghbi. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics* 1 (1): 71–81.
58. Haslam, N., S. Loughnan, Y. Kashima, and P. Bain. 2009. Attributing and denying humanness to others. *European Review of Social Psychology* 19 (1): 55–85.
59. Nunnally, J. 1978. *Psychometric methods*. New York: McGraw.
60. Wheeler, B. 2010. *lmPerm: Permutation tests for linear models*. R package version 1.1-2.
61. R Core Team. 2014. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.