The Robots are Coming! The Robots are Coming!

Fear and Empathy for Human-like Entities

Shelby Ceh

Eric J. Vanman

University of Queensland

Abstract

Social robots will soon be a part of our everyday work and home lives. Roboticians have assumed that a more human-like appearance in robots will ensure pro-sociality and cooperation between robots and humans. Indeed, people more readily attribute human-like capacities and a perception of mind to such robots. In turn, people show greater empathy for human-like robots. However, a human-like design has also been linked to beliefs about the potential of robots to be evil and cause damage to humanity. The aim of this research was to further investigate these ambivalent emotional reactions. In two studies, participants viewed images of robots of varying human likeness that were also depicted as the main character in a sad situation. Participants rated their sympathy for the robot but also their perceptions of threat. Facial electromyography (EMG) was added in the second study to measure empathic facial responses. We predicted that human likeness would have effects on both empathy and threat responses. We also tested whether these effects would be moderated by social categorization and threat information. Study 1 ($n = 201$) manipulated the categorization of the robot through changes in the phrasing of the scenarios, whereas Study 2 ($n = 120$) manipulated threat information about robots by first presenting positive or negative news stories about robots. Consistent with previous research, more human-like robots elicited greater empathic and threat responses, both on self-report and EMG measures. Manipulations of social categorization and threat failed to moderate the effects of human likeness on the measures. These results have considerable implications for the existing human-robot interaction literature and could help inform social robot designs for the future.

The Robots are Coming! The Robots are Coming!

Fear and Empathy for Human-like Entities

The futuristic notion of robots operating among humans in our everyday work and private lives is fast becoming a reality. Modern robots are no longer primarily designed for operation within the confines of factories, but rather for implementation into face-to-face human contact jobs, including care for the elderly and people with disabilities (de Graaf, Allouch, & Klamer, 2015). In fact, demand for social robots is expected to grow seven times faster than the demand for industrial robots in the next few years (Ambasna-Jones, 2016). The growing emphasis on social robot development has heightened concerns over how they will be received and accepted by potential users and co-workers. As with any new technology, there is likely to be some resistance to change when it comes to the initial introduction of social robots (Strait, Aguilon, Contreras, & Garcia, 2017). However, unlike most technologies, roboticists also have to contend with people's preconceived ideas about robots, developed over many years through books, TV shows, and movies (Szollosy, 2017; Złotowski, Proudfoot, Yogeeswaran, & Bartneck, 2015). This kind of destructive reputation could put successful robot introduction into jeopardy.

One of the most powerful design tools afforded to roboticists to facilitate social integration is the ability to make the robots look more human (Duffy, 2003; Złotowski et al., 2015). Modern day roboticists have the capacity to produce increasingly lifelike entities, typically termed androids, which possess uncanny human-like features such as hair and skin (Choi, Lee, Lee, Ahn, & Lee, 2011). These design decisions aim to increase identification between humans and robots, and have reliably been found to affect perceptions of anthropomorphism (Krach et al., 2008; Kuchenbrandt, Eyssel, Bobinger, & Neufeld, 2013).

By increasing anthropomorphism (Tondou, 2012), a more human-like appearance has been found to increase empathy (Riek, Rabinowitch, Chakrabarti, & Robinson, 2009) but also simultaneously a fear of robots (Ferrari, Paladino, & Jetten, 2016). These two separate, apparently contradictory, responses to robots have not been examined simultaneously. Therefore, the research we report here examined potential moderators of the effect of human likeness on these co-occurring positive and negative reactions.

Human Likeness and Empathy

People tend to see nonhuman entities as other social beings, whether the entities are animals, toys, cartoon characters, strange beasts in a movie, or robots in a store, and therefore respond to them as such (Bartneck & Hu, 2008; Reeves & Nass, 1996). For example, when people viewed a video of a robot dinosaur being tortured, they reported greater feelings of discomfort and negative affect, as well as displayed greater physiological response of distress (Rosenthal-von der Pütten, Krämer, Hoffmann, Sobieraj, & Eimler, 2013). A later study found that participants watching the same dinosaur in a friendly interaction or being tortured showed facial expression activity that reflected positive and negative emotions, respectively (Menne & Schwab, 2018). In another study (Suzuki, Galli, Ikeda, Itakura, & Kitazaki, 2015), participants viewed images of robot or human hands in painful or non-painful situations (i.e., about to be cut by scissors or posed a safe distance from scissors). Electroencephalography (ERP) readings were similar for robot and human hands, suggesting that participants were able to empathize with the robots' pain in a similar manner as they did with other humans. Studies like these provide strong evidence for humans' capacity to empathize and treat robots as social beings.

Although humans tend to be social with nonhuman entities, if the entity has a human-like appearance, people are even more likely to empathize with it. When participants in a study by Riek et al. (2009) viewed videos of various robots being treated cruelly, they rated feeling empathetic and willing to provide assistance to a greater extent when the robot was human-like than when the robot was mechanical in appearance. Why is this so? An attribution of anthropomorphic characteristics and perception of mind could be responsible for the greater degree of empathy expressed towards human-like robots. For example, Krach et al. (2008) found using functional magnetic resonance imaging (fMRI) that, when engaging in an interactive game with a human-like robot partner, activation in participants' brain areas associated with Theory of Mind perception was greater than when interacting with mechanical robots. In another study (Hegel, Krach, Kircher, Wrede, & Sagerer, 2008), after competing with a robot on a game, participants rated human-like android robots as more intelligent and more likely to possess human-like mind capacities than functional robots or computers. It follows from these findings that the effect of human likeness on empathy found by Riek et al. (2009) could be explained by an increased perception of robot's capacity to experience emotion.

Human Likeness and Perception of Threat

Despite the potential benefit afforded by more human-like designs of robots, there are also unintended negative consequences. The prospect of introducing robots into everyday human lives is frightening for many and is more threatening when the robot looks like a human (Ferrari et al., 2016; Strait et al., 2017). Traditionally, research in this domain has come under the umbrella of the Uncanny Valley theory (Mori, MacDorman, & Kageki, 2012). The Uncanny Valley describes that at a certain degree of human likeness, artificial entities

begin to cause unease and discomfort. Despite mixed support in the literature (Kätsyri, Förger, Mäkäräinen, & Takala, 2015; Wang, Lilienfeld, & Rochat, 2015), the Uncanny Valley has been affirmed by some researchers as one of the greatest obstacles to the successful integration of androids into society (Gee, Browne, & Kawamura, 2005; Mathur & Reichling, 2016).

Application of the Uncanny Valley in human-robot interaction research has received some support in studies that used both explicit and implicit measures. Self-report studies have shown that, at a particular degree of human likeness, participants dramatically increase ratings of eeriness and discomfort (MacDorman & Ishiguro, 2006; Seyama & Nagayama, 2007). Such feelings of discomfort can manifest themselves in how people behave towards robots. For example, Mathur and Reichling (2016) used an investment game task to measure how human likeness affected participants' degree of trust in hypothetical robot partners. They found that Uncanny Valley effects penetrated participants' implicit social decisions, such that they were less likely to trust human-like robots to act pro-socially to both parties' advantage. Human-like robots can be perceived as eerie or creepy, and therefore we trust them less.

Most explanations of Uncanny Valley effects have pertained to aspects of physical appearance. For instance, some have argued that the unnerving nature of lifelike robots stems from unnatural or slightly abnormal looking facial features, which do not quite resemble that of a real human (Seyama & Nagayama, 2007). Others have argued instead that any subtle flaws in the design of human-like robots are unnerving because they remind us of the inevitability of death (MacDorman & Ishiguro, 2006). More recently, explanations of discomfort towards human-like robots have focused less on physical features and more on inferences made from appearance. Consistent with the robot empathy literature, the unnerving

nature of human-like robots has been attributed by some researchers to be due to an attribution of mind to the robot. Like Krach et al. (2008), Gray and Wegner (2012) found that participants tended to perceive the same robot as more likely to possess human-like qualities and sentience when it appeared to be more human-like, compared to when it appeared to be mechanical. However, Gray and Wegner also found that simply describing the robot as having some sort of mind (e.g., capable of feeling hunger, fear or sadness) caused participants to feel unnerved, regardless of whether it was human-like or mechanical in appearance. This finding is evidence that anthropomorphism and attribution of sentience, not appearance itself, could be driving the Uncanny Valley effects seen in human-robot interaction studies.

More recently, Ferrari et al. (2016) proposed that human-like robots cause unease due to their perceived threat to the distinctiveness and uniqueness of humankind. That is, when robots are perceived to be too physically similar to humans, they blur the boundaries between human and machine, and are therefore considered more threatening. In this study, participants rated images of android, humanoid and mechanical robots on their capacity to cause damage to human identity, their likelihood of having mind capacities, and the extent to which they undermine the distinctiveness between man and machine. Contrary to Gray and Wegner (2012), Ferrari et al. found that a perception of mind was not a significant predictor of robots' perceived damage to humans, after accounting for human-like appearance. This finding suggests that robot appearance influences perception of threat, rather than attribution of sentience as found by Gray and Wegner.

Minimizing the perception of threat and danger from robots is likely to be important for successful human-robot interaction. Through this research, we aimed to contribute to the ongoing discussion surrounding human-robot interaction and robot acceptance by exploring

how human likeness affects individuals' abilities to empathize with robots, and how threatening they perceive robots to be. We further aimed to test two potential moderators of the relationship between human likeness and empathy and threat responses. Both Study 1 and Study 2 involved the presentation of a series of robot images and accompanying fictitious sad scenarios. The images varied in human likeness such that they fit into three categories: accurate depictions of humans (*android*), basic human figures (*humanoid*) and not resembling humans (*mechanical*). The hypotheses for the two studies were that (a) participants would have more empathy for androids than humanoid or mechanical robots and (b) participants would rate greater perception of threat and blurring of man-machine boundaries for android than humanoid or mechanical robots. Additional hypotheses pertained to potential moderators of the effect of human likeness on empathy and threat measures. Study 1 investigated the role of social categorization on this relationship. Study 2 explored the effects of manipulating threat information about robots. Specific hypotheses are discussed below. For both experiments, we report all conditions, measures, data exclusions and how we determined the sample size. The complete data files and materials used for both studies are permanently archived at https://osf.io/eu9p7/?view_only=07f3bab6cc9049dfb9c86879f1b81ee0

Study 1 – The Effect of Social Categorization

Initial social categorizations allow us to gain information about an individual with relatively little cognitive effort (Fiske, Neuberg, Beattie, & Milberg, 1987). These categorizations can greatly influence how we perceive and interact with others. For example, people who are categorized as outgroup members are often perceived less favorably (Castelli, Tomelleri, & Zogmaister, 2008) and less empathetically than ingroup members (Montalan, Lelard, Godefroy, & Mouras, 2012; Tarrant, Dazeley, & Cottom, 2009). Interestingly, social

categorization effects have been found in impression formation of robots. When people were presented with robots described to be part of their ethnic ingroup, they were more positive towards the robot, and rated it as more anthropomorphic than robots labelled as part of a different ethnic outgroup (Eyssel & Kuchenbrandt, 2012). In another study, arbitrary assignment of participants to groups associated with a particular robot made them feel more positive towards the robot, and more likely to attribute anthropomorphic characteristics (Kuchenbrandt et al., 2013).

The findings above demonstrate that group membership influences both likeability and anthropomorphism ratings, which suggests that social categorization may be a potential moderator of the effects of human likeness. If robots are given individual identities, they may be more readily attributed human-like capabilities and sentience. In turn, these robots could elicit greater empathy and threat responses. However, if outgroup robots receive lower attribution of anthropomorphism, they may be more likely to be treated categorically as robots, rather than as social beings. In effect, this could minimize the effect of human-like appearance on both empathy and threat responses.

In this study, we investigated whether making the outgroup status of the robots salient (i.e., calling them “a robot”), or framing them as individuals (i.e., attributing names and genders), would moderate the relationship between human-like appearance and empathy and threat responses. We hypothesized that, across conditions, participants would show greater ratings of both empathy and threat measures for android over humanoid and mechanical robots. Further, we predicted that the robots framed as individuals, with names and genders, would elicit greater empathy and higher threat perception across human likeness conditions, than those who are presented as simply “robots”. Finally, we hypothesized that the effects of

social categorization would interact with human likeness, such that the categorization effects would be greatest for participants in the condition where the robots were androids compared to when they were mechanical, for instance.

Method

Participants. Using a 3 x 2 mixed factorial design, we planned to recruit a minimum of $n = 60$ per group was aimed to ensure sufficient power (.80), assuming medium effect sizes (Faul, Erdfelder, Lang, & Buchner, 2007). There were 201 participants in this study, including 99 males and 101 females. Participants ranged in age from 18-66 with a mean age of 32.65 ($SD = 10.56$). Participants were primarily located in Europe ($n = 95$) and the United States ($n = 96$), as well as Oceania ($n = 4$) and Canada ($n = 6$). English was the primary language for 97.5% of participants. Participants were recruited through the online data collection site Prolific (www.prolific.ac) and were compensated £1.50 for their time.

Materials

Robot images and sad scenarios. The android, humanoid and mechanical robot images used for this study were accessed from online resources. The android robot images included Kokoro and Ishiguro's Geminoid DK and Geminoid F as well as Hanson Robotics' Jules, Philip and Philip DK. Humanoid robot images were of KOBIAN of Tmsuk and Waseda University; Hanson Robotics' Diego, University of Tokyo's Kojiro; Carnegie Mellon University's Snackbot and University of Kawada's HRP 4. Mechanical robots included Toshiba's Four-Legged Explorer, Modular Snake from CMU Biorobotics, Modular Robotic's MOSS and BioRob's Cheetah-Cub. The majority of these images were sourced from Ferrari et al. (2016) and the rest were obtained from multiple search engines. Original, sad scenarios were created for this research. Examples of the scenarios include robots that have lost a loved

one, been overworked, ignored and abused. All images and scenarios can be found at the study website.

A pilot study ($n = 80$) was conducted online through Prolific to check the stimuli. In the first part of the short study, participants rated how human-like each of the robot images were. Importantly, all the android robots were rated significantly more human-looking than humanoids or mechanical robots. Robot images that were more ambiguous were eliminated for the subsequent studies. In part two of the study, participants rated how sad 18 scenarios were. The three scenarios rated the least sad were deleted.

Response items. Two response items were designed to establish empathy. The items were “I think this scenario is sad” and “I would have sympathy for someone in this situation.” Two items were also included to assess feelings of threat towards robots, as well as specific threat to identity. These items were “I think this robot could be used for evil (to trick or fool, etc.)” and “This kind of robot blurs the boundaries between human beings and machines.” Threat items were obtained from scales used by Ferrari et al. (2016). The latter item was included in order to assess Ferrari et al.’s assertion that threat responses are caused by the blurring of human distinctiveness. All items had 6-point response scale, from 1 (*strongly disagree*) to 6 (*strongly agree*).

Procedure. Phrasing of the scenarios was varied between-groups, where one group of participants received scenarios containing individual information about the robot (i.e., name and gender) and the other group received no individual information (i.e., group category condition). For example, one individual group scenario was “Philip has worked as a carer for an elderly woman for the last three years, and he has just received word that she has passed away.” In contrast, the category-phrased group received a description such as, “This robot has

worked as a carer for an elderly woman for the last three years and has just received word that she has passed away.”

The study was conducted entirely online. After reading the information and consent forms, participants provided demographic details. All participants then viewed robot images (android, humanoid and mechanical randomized) along with sad scenarios, and responded to the two empathy items and the two threat items. Participants were randomly allocated to one of the two categorization conditions, such that approximately half were given scenarios with names and genders, and the other half were presented with the category (i.e., “this robot”) scenarios.

Results

A two-way mixed ANOVA was conducted for the two empathy and threat measures separately. Mauchly’s Test of Sphericity indicated that the assumption of sphericity had been violated for the effect of human likeness on sadness, $\chi^2(2) = 14.97$ $p = .001$, sympathy $\chi^2(2) = 6.72$, $p = .03$, propensity for evil, $\chi^2(2) = 144.16$, $p < .001$ and blurring of boundaries items, $\chi^2(2) = 116.88$, $p < .001$. Violation of sphericity can lead to overly liberal results, therefore Greenhouse-Geisser corrections were used for interpretation.

Empathy Measures.

Sadness. Figure 1 depicts the means and standard error for both sadness and sympathy ratings. A main effect of human likeness was observed for sadness, $F(2, 199) = 89.80$, $p < .001$, $\eta_p^2 = .311$. Follow-up pairwise comparisons showed that scenarios were rated to be significantly sadder when presented with android ($M = 3.51$, $SD = 1.37$) than humanoid robots ($M = 2.99$, $SD = 1.34$), and mechanical robots ($M = 2.79$, $SD = 1.36$). Scenarios were also rated higher on the sadness item when presented with humanoid over mechanical robots.

Mean sadness ratings were not different between individual ($M = 3.13$, $SD = 1.46$) and category phrasing conditions ($M = 3.06$, $SD = 1.25$), $F(1, 199) = 0.16$, $p = .69$, $\eta_p^2 < .001$. No interaction was found between human likeness and condition, $F(2, 199) = 0.99$, $p = .37$, $\eta_p^2 = .005$.

Sympathy. As can be seen in Figure 1, human likeness was also found to have an effect on ratings of sympathy, $F(2, 199) = 80.79$, $p < .001$, $\eta_p^2 = .289$. Pairwise comparisons revealed that scenarios were rated as more likely to evoke sympathy when presented with android ($M = 4.19$, $SD = 1.36$) over humanoid ($M = 3.83$, $SD = 1.46$), or mechanical robots ($M = 3.61$, $SD = 1.54$). Humanoid robot images also had higher sympathy ratings than mechanical robots. Individual ($M = 3.87$, $SD = 1.53$) and category conditions ($M = 3.88$, $SD = 1.37$) were not found to have an effect on ratings of sympathy, $F(1, 199) = 0.001$, $p = .97$, $\eta_p^2 = .001$. There was no interaction found between human likeness and scenario condition for sympathy, $F(2, 199) = 1.77$, $p = .17$, $\eta_p^2 = .009$.

Figure 1. Mean ratings, from 1 (*strongly disagree*) to 6 (*strongly agree*), across human likeness and condition for sadness (left) and sympathy (right) items. Asterisks indicate

significant differences across human likeness levels ($p < .05$). Standard errors are represented by error bars attached to each column.

Threat Measures.

Propensity for evil. Figure 2 depicts a main effect of human likeness was found for perceptions of robots' propensity for evil, $F(2, 199) = 145.72, p < .001, \eta_p^2 = .423$. Follow up pairwise comparisons revealed greater propensity for evil ratings for androids ($M = 3.71, SD = 1.45$) over humanoid ($M = 2.76, SD = 1.14$), or mechanical robots ($M = 2.41, SD = 1.13$), and humanoid over mechanical robots. Individualized robots ($M = 3.08, SD = 1.35$) were not rated differently on the propensity for evil item compared to those with the category phrasing ($M = 2.84, SD = 1.11$), $F(1, 199) = 2.65, p = .11, \eta_p^2 = .013$. There was no interaction between human likeness and scenario condition, $F(2, 199) = 0.32, p = .64, \eta_p^2 = .002$.

Figure 2. Mean ratings, from 1 (*strongly disagree*) to 6 (*strongly agree*), across human likeness and condition for propensity for evil (left) and blurring of boundaries (right) items. Asterisks indicate significant differences across human likeness levels ($p < .05$). Standard errors are represented by error bars attached to each column.

Blurring of Boundaries. Human likeness was also found to affect how participants perceived blurring of boundaries between man and machine, $F(2,199) = 578.50, p < .001, \eta^2 = .744$, as can be seen in Figure 2. Pairwise comparisons revealed differences between android ($M = 4.49, SD = 1.30$) and humanoid robots ($M = 2.60, SD = 1.05$), as well as between android and mechanical robots ($M = 1.85, SD = 0.96$). Differences were also found between humanoid and mechanical robots. No difference was found between individual ($M = 3.04, SD = 1.14$) and outgroup category ($M = 2.91, SD = 1.06$) conditions, $F(1, 199) = 0.96, p = .33, \eta^2 = .005$. Further, no interaction between human likeness and scenario condition was found for the blurring of boundaries item, $F(2, 199) = 1.37, p = .25, \eta^2 = .007$.

Discussion

Study 1 investigated whether social categorization of robots influenced the effects of human likeness on both threat and empathy responses. First, we hypothesized that participants would rate android robots higher on sadness and sympathy items than humanoid or mechanical robots. This hypothesis was supported. The findings revealed that although participants were asked to indicate their sadness and sympathy towards the scenarios, the human likeness of the robot in the simultaneously presented image affected their ratings. Consistent with the results of Riek et al. (2009), scenarios were rated as sadder and more worthy of sympathy when presented alongside android over humanoid or mechanical robots.

We also predicted that human likeness would affect ratings of propensity for evil and blurring of boundaries, such that participants would find androids more threatening than humanoid or mechanical robots. This hypothesis was also supported. Androids were rated as more capable of evil and more threatening to the blurring of man-machine boundaries than humanoid robots, and both were higher than mechanical robots. This result is similar to those

of Ferrari et al. (2016). None of the sad scenarios implied that the robots depicted were more or less intelligent than one another. Indeed, in reality, some of the mechanical-looking robots are more advanced than the android robots. Similarly, none were described to be designed for weaponry, warfare, or any other field that may affect how threatening they appear. Thus, any perceptual differences in how likely the robots were to participate in acts of evil were made based on the appearance of the robots alone.

We also tested the hypothesis that participants presented with scenarios providing individual identities to the robots would display greater empathy and perception of threat responses, due to greater anthropomorphism. This hypothesis was not supported. Social categorization of robots was not found to affect any of the provided measures. We also predicted that categorization would moderate the relationship between human likeness and empathy or threat perception, such that the effect of human likeness would be greatest for participants in the individual framing condition. This moderation effect was not supported in the findings. It is possible that the strong human likeness effect overcame any differences in the way the scenarios were phrased. It is equally plausible that the manipulation was not strong enough to incite meaningful differences in perception of categorization and anthropomorphism. Thus, social categorization in Study 1 did not affect the influence of human likeness on either empathy or threat. A different potential moderator, threat information, was chosen for investigation in Study 2.

Study 2. Message Framing

General positive and negative information towards robots may play a role in the way in which individuals perceive and react towards robots in their environment. Using the Negative Attitudes Towards Robots scale (NARS; Nomura, Suzuki, Kanda, & Kato, 2006), Syrdal,

Dautenhahn, Koay, and Walters (2009) found that negative attitudes reliably and accurately predicted how participants behaved when presented with robots. In accordance with this finding, Nomura et al. (2006) found that participants with more negative attitudes (i.e., greater NARS scores) were more likely to display avoidance behaviors when interacting with robots, than participants with neutral or positive attitudes. These studies provide evidence that attitudes may be accurate predictors of anti-social reactions towards robots in real life.

Thus far, the effect of attitudes on empathy for robots has yet to be explored. However, there is some evidence in the social psychology literature that negative attitudes towards individuals or outgroups predicts lower displays of empathy (Castiglione, Licciardello, Rampullo, & Campione, 2013; Gutsell & Inzlicht, 2010). Therefore, it is reasonable to suggest that this effect extends to robots as well. That is, it is possible that individuals with more positive attitudes towards robots will more readily empathize with them, than those with negative attitudes.

Interestingly, how threatening individuals find artificial entities has been shown to affect attribution of mind and anthropomorphism. Hackel et al. (2014) presented participants with a series of morphed images between humans and human-like dolls, described to be ingroup or outgroup members. They found that when an outgroup was perceived to be a potential threat, the threshold of human likeness for perceiving mind was significantly lower than when the outgroup was not perceived to be a threat. That is, participants attributed mind to images with lower degrees of human likeness when they were potentially threatening. This finding was attributed to evolutionary mechanisms, inferring that we may attribute mind to threatening outgroup members in an attempt to understand them and predict their next move. These results suggest that negative attitudes towards robots, which are likely to encompass

feelings of threat and anxiety, could influence how humans anthropomorphize and attribute mind to robots.

Physiological Measures

Facial electromyography (EMG) is a non-invasive tool for measuring affective responses (Tassinari, Cacioppo, & Vanman, 2016). Facial EMG records particular facial muscle regions, including the *corrugator supercilii*, which pulls the brow downwards to form a frown, and the *zygomaticus major*, which elevates the edges of the mouth to produce a smile. Through the activation of these muscles, it is possible to discriminate both the valence and intensity of an individual's emotions. Importantly, corrugator activity has been associated with not only negative affect such as sadness, but also with empathic concern for another individual (Sun, Wang, Wang, & Luo, 2015). It has been proposed that the corrugator muscle is activated when witnessing another's pain or sadness due to a process of facial mimicry, which facilitates greater understanding of another's emotional state (Gallese, 2005; Goldman & Sripada, 2005). "Mimicry" here might involve a one-to-one correspondence between the observer's and target's facial expressions, or it could more simply refer to the viewer making a facial expression that is concordant with the target's perceived emotion. For example, one might hear the story of about a friend's sad event and thus display a sad facial expression, even though the friend is not present to mimic directly. The one-to-one mimicry, however, is believed to be essential for the recognition of emotions and the empathy response. As such, interference with mimicking behavior has been found to inhibit recognition of emotion and empathic response (Oberman, Winkielman, & Ramachandran, 2007).

Do humans mimic the expressions of robots in the same way that they do other humans? Only one study to date has investigated that question (Hofree, Ruvolo, Bartlett, &

Winkielman, 2014). In that study, when participants interacted with a human-like robot they spontaneously mimicked the robot's expression. It is important to note that the robot used in that study was capable of making facial expressions, so mimicry was relatively straightforward to measure. This leaves open the question of what happens when there are no such facial expressions, either because the robot has no ability to make them (and perhaps does not even have a face), or that still photos of robots are used that do not show changes in facial expressions. This study is the first to compare facial mimicry across varying degrees of human likeness. More recently, however, Menne and Schwab (2018) analyzed facial actions associated with emotion and found that when participants viewed a dinosaur robot being tortured, they expressed more brow (corrugator) activity. That study did not demonstrate facial mimicry—the toy dinosaur was incapable of making facial expressions—but it nevertheless suggests that the activity of the corrugator muscle can serve as a marker of an empathic response when people interact with nonhuman entities (Rae Westbury & Neumann, 2008) like robots.

Aims and Hypotheses

Here we tested whether the effects of human-like appearance on empathy and threat perception found in Study 1 would be moderated by a positive or negative message frame about robots. These messages were 5-min audio clips, which either advocated the benefits of robots in society, or the impending threat of robots to mankind. We also measured attitudes about robots using the NARS as a potential covariate in exploratory analyses. As in Study 1, we predicted that participants would display more empathy and greater perception of threat for androids, over humanoids and mechanical robots. Accordingly, we predicted that greater corrugator activity associated with empathy would occur for androids over humanoid and

mechanical robots. We also predicted that participants in the threatening message condition would show a greater threat response than those in the positive attitude condition, and that participants in the positive attitude condition would have greater empathy response, on both self-report and the EMG measures. Finally, we hypothesized that the effect of human likeness on both empathy and perception of threat would be more pronounced or amplified in the positive attitude condition, than in the negative attitude condition.

Method

Participants. Using a 2 (Message frame: positive, negative) x 3 (Human likeness: android, humanoid, mechanical) mixed factorial design, with the first factor manipulated between subjects, we aimed for $n > 60$ for each of the two message frame levels to allow ample statistical power (.80) assuming a medium effect size (Faul et al., 2007). There were 120 participants in this study, including 83 females and 37 males. Ages ranged from 17 to 42, with a mean of 20.54 ($SD = 4.46$). Participants were students of an undergraduate psychology course at the University of Queensland. Participants were awarded credit towards their first-year course grade for taking part in the study. Ethnicity was predominantly Caucasian (51.7%) and Asian (36.7%), as well as Indian (5.8%). Other ethnicities, including African, Hispanic, Maori and Samoan made up a further 5.8% of participants.

Materials.

Message Framing Manipulation. Participants were presented with a 5-min audio clip, which consisted of news stories that were either pro-robots or anti-robots (depending on condition). These audios were recorded originally for this study. Scripts are available at the project's OSF site.

Physiological and Self-Report Measures. This study involved the use of facial electromyography (EMG) electrodes, which recorded activity in the zygomaticus and corrugator muscles on the face. Skin conductance was also recorded as an exploratory measure of threat response. However, this measure was excluded from analysis following testing due to low SCR responding from most participants.

DirectRT computer software was used to run the task, which included the same 15 images and scenarios used in Study 1. Scenarios were used in their individual form (i.e., with robot names and genders retained,) in order to maximize the likelihood of anthropomorphism from participants. Items of empathy and perception of threat were the same as those used in Study 1.

Post Study Questionnaire and Forms. The Negative Attitudes Towards Robots scale (NARS) was also used to assess longer term attitudes and served as a covariate. This 16-item scale assesses feelings of threat and nervousness, such that higher scores indicate more negative views (Nomura, Kanda, & Suzuki, 2006). Participants were also asked three brief items pertaining to their experience with robots in the past, in order to assess their familiarity with robot technology. Debrief forms giving further information about the purpose and hypotheses of the study were also developed for distribution following completion of the study.

Design and Procedure. Participants sat approximately 1 m from a 17-in computer monitor in a small, quiet room. They were first given the informed consent form to read, followed by the short demographics form. The experimenter orally informed participants that the study aimed to assess how individuals feel towards robots and instructed that the sensors would be measuring signals emanating from the head.

Following this, the experimenter prepared the participant's face for electrode attachment following recommended guidelines (Fridlund & Cacioppo, 1986). Participants removed facial hair or makeup in the areas required for electrode attachment prior to attending the study. To further maximize sensitivity for the electrode readings, the experimenter cleansed the electrode sites on the participant's face using sensitive soap and alcohol wipes. Two pairs of 2mm Ag/AgCl EMG electrodes were attached to the left cheek (zygomaticus muscle) and inner left eyebrow (corrugator muscle). Skin impedance at these sites was below 10 kohms. A large 8mm Ag/AgCl ground electrode was also attached to the forehead on the right side of the face. Two large 8mm Ag/AgCl electrodes were also placed, following cleaning, on the inside, middle third of two fingers on the participant's non-dominant hand. These were used to record skin conductance data.

The participants wore headphones throughout the testing session. They first listened to the audio news story about robots. They then viewed a fixation cross in the center of the screen for a random period of time between 1.5 and 2 s, before the robot image and accompanying sad scenario were presented for 12 s. The order of presentation of robots was randomized, and scenarios were also randomly allocated and varied between two sets between participants. The image then disappeared and participants were prompted to respond to the same items used in Study 1 (sadness, sympathy, propensity for evil and blurring of boundaries items) on a scale from 1 (*strongly disagree*) to 6 (*strongly agree*).

During the task, facial EMG was recorded continuously. The experimenter, who was in the room next door, had a camera view of the participant throughout the experiment, allowing her to make note of any facial movements unrelated to the task (e.g., if the participant yawned or sneezed). At the conclusion of the task, the experimenter prompted

participants to make exaggerated smiles and frowns, in order to ensure the equipment was recording properly. Following this, the experimenter entered the room, removed the participant's headphones and presented them with the NARS and the experience with robots questionnaire. When the participant had completed this questionnaire, the experimenter removed the electrodes and debriefed the participant.

Results

Data Reduction and Clean Up. Facial EMG data were reduced according to recommended guidelines (Fridlund & Cacioppo, 1986). Raw EMG signals were amplified by a factor of 5000 using a Biopac MP-150 amplifier (Biopac Systems, Inc., Santa Barbara, CA). Acknowledge software was used to record data at 1000 Hz, which was later filtered using 20 to 500 Hz filters, rectifying and integrating at a 50-ms time constant. Signals were segmented into 12000-ms long intervals post onset of stimuli and averaged into three time periods (1000-4000 ms, 5000-8000 ms and 8000-12000 ms).

Offline, an examination of the EMG data identified participants with poor readings according to testing exaggerated expressions immediately following the experiment, or outlier readings caused by unrelated facial movements (e.g., twitches, chewing, consistent coughing, etc.). These data were excluded from the analysis. Self-report data were also checked for outliers and missing data but were all found to be usable for analysis. Mauchly's Test of Sphericity indicated that the assumption of sphericity had been violated for the effect of human likeness on sadness, $\chi^2(2) = 12.89, p = .002$, sympathy $\chi^2(2) = 15.75, p < .001$, propensity for evil, $\chi^2(2) = 44.86, p < .001$ and blurring of boundaries items, $\chi^2(2) = 39.33, p < .001$. Sphericity was also violated for zygomaticus, $\chi^2(2) = 11.10, p = .004$ and corrugator

readings, $\chi^2(2) = 12.91, p = .002$. Thus, as in Study 1, Greenhouse-Geisser corrections were used.

Self-report Data. A two-way mixed factorial ANOVA was conducted for each dependent variable. Both NARS and previous experience with robots scores were assessed as covariates in separate exploratory analyses for each item.

Figure 3. Mean ratings, from 1 (*strongly disagree*) to 6 (*strongly agree*), across human likeness and audio condition for sadness (left) and sympathy (right) items. Asterisks indicate significant differences across human likeness levels ($p < .05$). Standard errors are represented by error bars attached to each column.

Sadness. A mixed factorial ANOVA revealed a significant main effect of human likeness for the sadness measure, $F(2, 118) = 74.87, p < .001, \eta_p^2 = .388$, as depicted in Figure 3. This was followed up with pairwise comparisons, which indicated that participants rated scenarios as sadder when presented with accompanying images of androids ($M = 3.88, SD = 1.28$) over humanoids ($M = 3.22, SD = 1.33$), or mechanical robots ($M = 3.07, SD = 1.40$). Sadness ratings were also significantly higher for humanoid than mechanical robots.

Participants in the positive message condition ($M = 3.35$, $SD = 1.35$) did not rate differently in sadness than those exposed to the negative message ($M = 3.42$, $SD = 1.28$), $F(1, 118) = 0.10$, $p = .76$, $\eta^2 = .001$. There was no interaction between human likeness and message frame for sadness, $F(2, 118) = 1.75$, $p = .18$, $\eta^2 = .014$. Analyses with NARS and experience scores as covariates revealed that neither score was found to affect sadness ratings.

Sympathy. Additionally, as can be seen in Figure 3, human likeness was found to have an effect on sympathy, $F(2, 118) = 45.29$, $p < .001$, $\eta^2 = .277$. Pairwise comparisons confirmed that participants considered that they would be more likely to have sympathy for scenarios presented with android ($M = 4.52$, $SD = 1.27$) over humanoid ($M = 4.10$, $SD = 1.44$), or mechanical robot images ($M = 3.94$, $SD = 1.50$). Sympathy ratings were also higher for humanoid than mechanical robots. Message frame also appeared to have no effect on sympathy ratings, $F(1, 118) = 0.03$, $p = .86$, $\eta^2 < .001$, ($M = 4.16$, $SD = 1.43$ and $M = 4.20$, $SD = 1.28$ for good and evil conditions respectively). There was no interaction between human likeness and message frame for sympathy, $F(2, 118) = 0.56$, $p = .55$, $\eta^2 = .005$. NARS and experience with robots scores were analyzed as covariates but, as with the sadness item, were not found to affect sympathy.

Propensity for evil. Mixed factorial ANOVA analyses revealed a main effect of human likeness for perception of propensity for evil, $F(2, 118) = 7.06$, $p = .003$, $\eta^2 = .057$, as depicted in Figure 3. Follow up pairwise comparisons found that androids ($M = 3.79$, $SD = 1.22$) were rated significantly higher than humanoid ($M = 3.14$, $SD = 1.10$), or mechanical robots ($M = 2.45$, $SD = 1.02$) on the propensity for evil item. Humanoids were also rated more capable of evil than mechanical robots. Additionally, a significant main effect of message frame was found for the evil item, $F(1, 118) = 6.80$, $p = .01$, $\eta^2 = .055$, such that individuals

in the negative message condition rated robots as more capable of being evil ($M = 3.38$, $SD = 0.90$) than those in the positive message condition ($M = 2.86$, $SD = 0.95$). There was no interaction between human likeness and message frame found for the propensity for evil item, $F(2,118) = 1.79$, $p = .18$, $\eta_p^2 = .015$.

Analyses with the NARS as a covariate revealed it was related to the propensity for evil item, $F(1, 118) = 23.95$, $p < .001$, $\eta_p^2 = .170$. This highlights the validity of this measure. The estimated effect size of human likeness on perception of propensity for evil was reduced by entering NARS as a covariate (from $\eta_p^2 = .472$ to $\eta_p^2 = .057$), but not eliminated. Experience scores were not found to affect propensity for evil ratings.

Blurring of Boundaries. Further, human likeness was found to have an effect on participants' ratings of the extent to which robots blurred the boundaries between man and machine, $F(2,118) = 14.74$, $p < .001$, $\eta_p^2 = .112$. This effect is depicted in Figure 4. Pairwise comparisons found that androids ($M = 4.76$, $SD = 1.10$) were rated higher on the blurring of boundaries items than humanoids ($M = 3.05$, $SD = 0.99$), or mechanical robots ($M = 2.08$, $SD = 1.01$). Humanoids were also rated significantly higher on this measure than mechanical robots. There was no main effect of message frame for the blurring of boundaries item, $F(1, 118) = 1.15$, $p = .28$, $\eta_p^2 = .010$ ($M = 3.40$, $SD = 0.85$ and $M = 3.20$, $SD = 0.70$ for positive and negative message conditions, respectively). There was also no interaction between human likeness and message frame found for the blurring of boundaries item, $F(2,118) = 0.30$, $p = .68$, $\eta_p^2 = .003$.

NARS results were significant for blurring of boundaries, $F(1, 118) = 4.31$, $p = .04$, $\eta_p^2 = .035$. The estimated effect size of human likeness on this item was reduced by entering

NARS as a covariate (from $\eta_p^2 = .732$ to $\eta_p^2 = .112$). Experience with robots did not affect blurring of boundaries scores.

Figure 4. Mean ratings, from 1 (*strongly disagree*) to 6 (*strongly agree*), across human likeness and audio condition for propensity for evil (left) and blurring of boundaries (right) items. Asterisks indicate significant differences across human likeness or conditions ($p < .05$). Standard errors are represented by error bars attached to each column.

Facial EMG Data. A mixed factorial ANOVA was conducted with human likeness (android, humanoid and mechanical) and time (first four seconds, second four seconds and third four seconds of recording) as within-group variables and message frame as the between-groups variable. Zygomaticus responses are reported separately from corrugator results. NARS and experience scores were also analyzed as covariates in separate analyses.

Zygomaticus. No difference in zygomaticus activity was found between android, ($M = 0.055$, $SD = 0.053$), humanoid ($M = 0.055$, $SD = 0.051$) or mechanical robots ($M = 0.055$, SD

$= 0.052$), $F(2, 100) = 0.26$, $p = .75$, $\eta^2 = .003$. Additionally, zygomaticus activity was not different between participants in the positive ($M = 0.055$, $SD = 0.047$) or negative message conditions ($M = 0.056$, $SD = 0.057$), $F(1, 100) = 0.01$, $p = .93$, $\eta^2 = .001$. No interaction between human likeness and message frame was found, $F(2, 100) = 0.78$, $p = .45$, $\eta^2 = .008$. There was no significant main effect of time, $F(2, 100) = 2.93$, $p = 0.06$, $\eta^2 = .028$, nor interactions between time and human likeness, $F(4, 100) = 0.73$, $p = .55$, $\eta^2 = .002$, or time and message frame, $F(2, 100) = 0.16$, $p = .84$, $\eta^2 = .007$. No three way interaction between time, human likeness and message frame was found, $F(2, 100) = 0.65$, $p = .60$, $\eta^2 = .002$. Separate analyses including NARS and experience scores as covariates were not found to affect zygomaticus activity.

Figure 5. Corrugator activity across human likeness. Standard errors are represented by error bars attached to each data point.

Corrugator. A mixed ANOVA with human likeness, time and condition revealed a main effect of human likeness for corrugator activity, $F(2, 106) = 3.46$, $p = .04$, $\eta^2 = .031$,

such that participants showed more corrugator activity toward androids than they did to humanoid and mechanical robots. This effect is depicted in Figure 5. No other main effects or interactions were significant. Analyses were conducted with NARS and experience scores as covariates, but neither were found to affect corrugator activity, or add to the relationships found.

Discussion

Here we explored whether the effects of human likeness found in Study 1 could be influenced by threatening information about robots and individual attitudes towards robots. As hypothesized, participants rated scenarios as sadder and more likely to evoke sympathy when presented with images of androids, over humanoid and mechanical robots. Facial EMG recordings provided physiological evidence of these results, as participants were found to display greater corrugator activity towards androids over humanoid robots. This finding is consistent with those of Riek et al. (2009), if one considers the corrugator activity to be measure of empathy.

And, like Study 1, human likeness was also predicted to affect ratings on threat items, such that greater human likeness would result in higher ratings on both propensity for evil and blurring of boundaries items. This hypothesis was supported. In alignment with previous research (Ferrari et al., 2016; Gray & Wegner, 2012), participants rated androids higher for self-report threat measures than humanoid or mechanical robots. These findings are in concurrence with those of Study 1 and provide convincing support for the existing body of literature.

We further hypothesized that participants presented with a threatening audio message about the dangers of modern robots would have a greater threat response than those in the

positive audio condition. This hypothesis was partially supported. This threat manipulation was found to significantly increased perceptions of evil but failed to affect the blurring of boundaries item. In effect, this could indicate that the two ideas are separable, rather than blurring of boundaries causing perception of evil as suggested by Ferrari et al. (2016). NARS scores were positively associated with both threat items. This highlights the validity of the measure. Although accounting for some variance in threat scores, NARS did not account for all variance encapsulated by human likeness.

In addition, we also expected that participants exposed to a positive message about robots would have a greater empathic response. Contrary to this hypothesis, neither the message manipulation nor the NARS score was found to predict empathy measures or corrugator activity. This is in contradiction to previous social psychology literature with humans (Castiglione et al., 2013; Gutsell & Inzlicht, 2010). Also contrary to predictions, attitudes did not interact with human likeness for any measures (i.e., neither empathy nor threat).

General Discussion

In two studies, using different populations with different methodologies, we investigated the ambivalent feelings humans hold towards human-like robots. The effects of the degree of human-likeness of a robot on empathy (for a sad situation) and perceptions of threat were assessed simultaneously. This work builds on those of other studies, but which addressed these variables separately. Further, we aimed to identify potential moderators of the effects of human likeness, on both empathy and threat reactions. Specifically, we investigated whether manipulating social categorization (Study 1) or threatening information (Study 2), affected the effects of human likeness and empathy and perception of threat.

Consistent with the findings of Riek et al. (2009), we predicted that participants would display greater empathy for androids over humanoid and mechanical robots. This hypothesis was supported in both Study 1 and Study 2. Greater corrugator activity found in Study 2 further provided physiological evidence of this empathic response. Although previous studies (e.g., Hofree et al., 2014) have found empathic mimicry for robots, this is the first study to find differences in such implicit reactions that varied with robot appearance.

In accordance with Gray and Wegner (2012) and Ferrari et al. (2016), it was further hypothesized that androids would be perceived as more capable of evil and more likely to blur the boundaries between man and machine, than less human-like models. This hypothesis was supported in both studies. Study 2 revealed that human likeness was still a significant predictor of perception of propensity for threat and blurring of boundaries, even after accounting for existing negative attitudes towards robots. This finding provides further support for the assertion that human-like robotic design can have negative effects on their acceptance into human lifestyles.

We also tested potential moderators (social categorization, threatening information) that we predicted affect the strength of the effects of human likeness on empathy and perceptions of threat. Contrary to these predictions, neither social categorization (Study 1) nor message framing (Study 2) were found to moderate the relationship between human likeness and empathy or perceptions of threat. We next discuss this lack of moderation.

Why Did We Find No Moderators?

It is possible that the manipulations were not strong enough to change participants' pre-conceived attitudes towards robots. For Study 2, it is likely that participants would have entered the experiment with their own ideas and conceptions about robots, and that a 5-min

audio message may not be strong enough to sway these attitudes. A more powerful stimulus, such as a video of robots harming or saving people, may have had a greater effect. However, we did find a main effect of the audio manipulation on the ratings of robots' propensity for evil. Participants primed with the threatening message rated robots as more likely to be used for evil than those presented with the positive audio. This finding suggests that the manipulation was strong enough to influence participants' perceptions of the robots to some degree.

It is also plausible that the effect of human likeness on empathy and threat responses was just more powerful than the other manipulated variables. Thus, the lack of moderation effects found in this study could be interpreted as highlighting the pervasiveness of the effects of human-like appearance on these ambivalent emotions, such that they were able to override contextual manipulations and prejudices. When attempting to understand the driving forces behind the increase in threat and emotive response as a result of human-like appearance, anthropomorphism is a strong factor. Androids have consistently been found to elicit a greater unconscious attribution of human capacities (Hegel et al., 2008; Krach et al., 2008; Kuchenbrandt et al., 2013). Attribution of human-like qualities to inanimate objects has clear implications for empathy. It may be that a human-like appearance triggers the attribution of human-like ability to feel pain and sadness (i.e., sentience), thereby evoking empathy in the onlooker. Whereas mechanical robots are unlikely to be attributed this human-like capacity for feeling sadness, they do not evoke the same degree of empathy.

At the same time, the process of anthropomorphism could be underlying the greater perceptions of threat for android robots over less human-like models. The findings of the current study suggest that human-like appearance alone is enough to cause people to believe

that a robot is more capable of evil acts, without any supporting evidence of the robot's abilities, purpose, or its level of sophistication. Gray and Wegner (2012) suggested that it is the attribution of capacity to sense and feel emotion that causes people to feel uneasy and threatened by human-like robots. According to Gray and Wegner, sentience is considered one of the most fundamentally human capacities, and when this ability is ascribed to robots, it brings about discomfort and feelings of threat in humans. In this way, it could be precisely because we feel empathy for human-like robots, that we also feel threatened by them. This explanation is also consistent with Hackel et al. (2014), who found that the threshold for perceiving another as having a mind is lower when the entity is considered a threatening outgroup member (i.e., a robot). Note too that all instances of robot presentations were accompanied by sad stories, and yet the effects on perceived evil remained.

However, we can also consider an alternative explanation. The Threat to Distinctiveness Hypothesis proposes that too much perceived similarity between robots and humans triggers concerns because similarity blurs the boundaries between humans and machines (Ferrari et al., 2016). Ferrari et al. suggested that people feel threatened by human-like robots because their appearance increases perceptions of their threat to the distinctiveness of humans, rather than as a function of perceived sentience proposed by Gray and Wegner (2012). To further investigate the Threat to Distinctiveness Hypothesis, the blurring of boundaries item was included in the current study. Studies 1 and 2 found that the more human-like the robot, the higher the ratings of perception of propensity for evil, and the higher perception of robots blurring boundaries. These findings provide some support for the Threat to Distinctiveness Hypothesis.

However, Study 2 found listening to a negative message about robots affected how participants rated the robots' propensity for evil but had no effect on their ratings of distinctiveness threat. If the assertion that human-like robots are perceived to be evil because of the threat they pose to human distinctiveness were true, it would be expected that the blurring of boundaries item would have been affected in the same way as propensity for evil ratings by the message frame. This finding suggests that human distinctiveness and perception of evil propensity may be separate constructs. Ferrari et al. (2016), on the other hand, conceptualized threat as a follow-on effect of the blurring of human distinctiveness. Clearly, future research could address these issues. Although the findings of the current study support the suggestion that humans feel more threatened by human-like robots, whether this reaction is due to anthropomorphism and attributions of sentience (Gray & Wegner, 2012), or due to the blurring of robot/human likeness (Ferrari et al., 2016), remains inconclusive.

Limitations and Future Directions

Corrugator activity (i.e., frowning), measured through facial EMG, has been found to differentiate between negative and positive emotions reliably (Tan et al., 2012; Tassinari et al., 2016). However, corrugator activity can occur for a variety of reasons, not solely empathic response. Therefore, it is important to assess whether the corrugator activity found for android images is the result of empathy, threat, or some other response (e.g., surprise, discomfort, etc.). An attempt was made to account for this limitation in the current study, by analyzing corrugator recordings in time intervals. Analyses indicated that the corrugator activity response time and duration were more consistent with empathic concern, than other explanations such as discomfort, unease, threat or surprise. Specifically, activity typically occurred between 5 s and 8 s after the presentation of the image and scenario. As this is likely

to be the timeframe where participants would be finishing reading the sad scenarios, it is suggested that this is evidence of empathic response. If participants frowned more as a result of threat or surprise, it would be predicted that the activity would occur immediately after presentation of the image. Other explanations, such as discomfort or unease, might be expected to persist throughout the recording period, which they did not. Therefore, an effort was made to address alternative explanations for corrugator activity, but empathic concern was the most viable explanation.

The use of facial EMG in the present study is a valuable addition to the findings and field of literature. However, if researchers were to use this measure for similar study in the future, a few improvements would be recommended. First, we would suggest that participants are presented with neutral, as well as sad scenarios. This design feature would allow researchers to compare corrugator activity to each category of robot across neutral and sad conditions. If greater activity is observed for sad scenarios, this would be considerable evidence that the response is due to empathy, not an alternative explanation. Further, we would recommend that the scenarios are presented alone for 12 s, to allow the participants to read the scenario before the image appears alongside. This would allow for an isolated view of corrugator activity when seeing the robot, not when reading the scenario.

Additionally, future research may aim to further explore the theories of Ferrari et al. (2016) and Gray and Wegner (2012). Contradictory findings in regards to whether anthropomorphism and attribution of sentience or intergroup distinctiveness threat underlies perception of evil in robots requires further research. Although the present study provides some support for Gray and Wegner's assertions that ascription of sentience increases perception of threat, Ferrari et al. found that sentience did not affect perceived propensity for

evil in their study. One way to address this discrepancy would be to design an experiment where participants are given several images of robots of varying human likeness, and short scenarios describing something neutral about them (e.g., where they work, what they are programmed to do) or scenarios describing them as having sentience (e.g., why they are feeling sad, upset, joyful or other human emotions). An examination of human perception of threat between these groups may elucidate the influence of human likeness alone, versus sentience.

Conclusion

At present, human-robot interaction research is a field of pressing importance and constant development. The current study found evidence for the effects of human likeness on both empathy and threat responses. Humans possess ambivalent and contradictory feelings towards human-like robots, which is likely determined by the same anthropomorphism mechanism. Although competing theories exist, there seems to be a pervasive understanding in the current literature that humans feel threatened by robots which undermine aspects of our identity, whether in terms of physical appearance or sentience. This understanding may have potential for informing the design decisions of future roboticists and the social psychology of humans and robots.

Appendix A

Robot Images and Results from Pilot Test – “I think this robot looks like a human” from 1 (*strongly disagree*) to 6 (*strongly agree*). Androids (top row), Humanoids (second row), Mechanical robots (third row).

Geminoid DK	Jules	Geminoid F	Philip DK	Philip

$M = 5.68$ $SD = 0.61$	$M = 5.21$, $SD = 1.93$	$M = 5.30$ $SD = 1.02$	$M = 5.55$ $SD = 0.81$	$M = 5.28$ $SD = 0.94$
HRP4 Kawada	Kojiro	Snackbot	Kobian	Diego San Hanson

$M = 2.39$ $SD = 1.15$	$M = 2.84$ $SD = 1.35$	$M = 2.09$ $SD = 1.03$	$M = 2.60$ $SD = 1.28$	$M = 3.29$ $SD = 1.19$
Four Legged Explorer a2	Modular Snake	MOSS	Four Legged Explorer	EPFL Biorobotics

$M = 1.04$, $SD = 0.19$	$M = 1.01$ $SD = 0.11$	$M = 1.10$ $SD = 0.49$	$M = 1.18$ $SD = 0.74$	$M = 1.14$ $SD = 0.49$

References

- Ambasna-Jones, M. (2016, May 9). How social robots are dispelling myths and caring for humans. *The Guardian*. Retrieved from <https://www.theguardian.com/media-network/2016/may/09/robots-social-health-care-elderly-children>
- Bartneck, C., & Hu, J. (2008). Exploring the abuse of robots. *Interaction Studies*, 9(3), 415–433. <https://doi.org/10.1075/is.9.3.04bar>
- Castelli, L., Tomelleri, S., & Zogmaister, C. (2008). Implicit ingroup metafavoritism: Subtle preference for ingroup members displaying ingroup bias. *Personality and Social Psychology Bulletin*, 34(6), 807–818. <https://doi.org/10.1177/0146167208315210>
- Castiglione, C., Licciardello, O., Rampullo, A., & Campione, C. (2013). Intergroup anxiety, empathy and cross-group Friendship: Effects on attitudes towards gay Men. *Procedia - Social and Behavioral Sciences*, 93, 969–973. <https://doi.org/10.1016/J.SBSPRO.2013.09.312>
- Choi, D., Lee, D.-W., Lee, D. Y., Ahn, H. S., & Lee, H. (2011). Design of an android robot head for stage performances. *Artificial Life and Robotics*, 16(3), 315–317. <https://doi.org/10.1007/s10015-011-0938-4>
- de Graaf, M. M. A., Allouch, S. Ben, & Klamer, T. (2015). Sharing a life with Harvey: Exploring the acceptance of and relationship-building with a social robot. *Computers in Human Behavior*, 43, 1–14. <https://doi.org/10.1016/J.CHB.2014.10.030>
- Duffy, B. R. (2003). Anthropomorphism and the social robot. *Robotics and Autonomous Systems*, 42(3–4), 177–190. [https://doi.org/10.1016/S0921-8890\(02\)00374-3](https://doi.org/10.1016/S0921-8890(02)00374-3)
- Eyssel, F., & Kuchenbrandt, D. (2012). Social categorization of social robots: Anthropomorphism as a function of robot group membership. *British Journal of*

- Social Psychology*, 51(4), 724–731. <https://doi.org/10.1111/j.2044-8309.2011.02082.x>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <https://doi.org/10.3758/BF03193146>
- Ferrari, F., Paladino, M. P., & Jetten, J. (2016). Blurring human–machine distinctions: Anthropomorphic appearance in social robots as a threat to human distinctiveness. *International Journal of Social Robotics*, 8(2), 287–302. <https://doi.org/10.1007/s12369-016-0338-y>
- Fiske, S. T., Neuberg, S. L., Beattie, A. E., & Milberg, S. J. (1987). Category-based and attribute-based reactions to others: Some informational conditions of stereotyping and individuating processes. *Journal of Experimental Social Psychology*, 23(5), 399–427. [https://doi.org/10.1016/0022-1031\(87\)90038-2](https://doi.org/10.1016/0022-1031(87)90038-2)
- Fridlund, A. J., & Cacioppo, J. T. (1986). Guidelines for human electromyographic research. *Psychophysiology*, 23(5), 567–589. <https://doi.org/10.1111/j.1469-8986.1986.tb00676.x>
- Gallese, V. (2005). Embodied simulation: From neurons to phenomenal experience. *Phenomenology and the Cognitive Sciences*, 4(1), 23–48. <https://doi.org/10.1007/s11097-005-4737-z>
- Gee, F. C., Browne, W. N., & Kawamura, K. (2005). Uncanny valley revisited. In *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication* (Vol. 2005, pp. 151–157). IEEE. <https://doi.org/10.1109/ROMAN.2005.1513772>
- Goldman, A. I., & Sripada, C. S. (2005). Simulationist models of face-based emotion recognition. *Cognition*, 94(3), 193–213.

<https://doi.org/10.1016/J.COGNITION.2004.01.005>

Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition*, 125(1), 125–130.

<https://doi.org/10.1016/J.COGNITION.2012.06.007>

Gutsell, J. N., & Inzlicht, M. (2010). Empathy constrained: Prejudice predicts reduced mental simulation of actions during observation of outgroups. *Journal of Experimental Social Psychology*, 46(5), 841–845.

<https://doi.org/10.1016/J.JESP.2010.03.011>

Hackel, L. M., Looser, C. E., & Van Bavel, J. J. (2014). Group membership alters the threshold for mind perception: The role of social identity, collective identification, and intergroup threat. *Journal of Experimental Social Psychology*, 52, 15–23.

<https://doi.org/10.1016/J.JESP.2013.12.001>

Hegel, F., Krach, S., Kircher, T., Wrede, B., & Sagerer, G. (2008). Theory of mind (ToM) on robots. In *Proceedings of the 3rd international conference on Human robot interaction - HRI '08* (p. 335). New York, New York, USA: ACM Press.

<https://doi.org/10.1145/1349822.1349866>

Hofree, G., Ruvolo, P., Bartlett, M. S., & Winkielman, P. (2014). Bridging the mechanical and the human mind: Spontaneous mimicry of a physically present android. *PLoS ONE*, 9(7), e99934. <https://doi.org/10.1371/journal.pone.0099934>

Kätsyri, J., Förger, K., Mäkräinen, 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(MAR), 390.

<https://doi.org/10.3389/fpsyg.2015.00390>

Krach, S., Hegel, F., Wrede, B., Sagerer, G., Binkofski, F., & Kircher, T. (2008). Can machines think? Interaction and perspective taking with robots investigated via fMRI.

- PLoS ONE*, 3(7), e2597. <https://doi.org/10.1371/journal.pone.0002597>
- Kuchenbrandt, D., Eyssel, F., Bobinger, S., & Neufeld, M. (2013). When a robot's group membership matters. *International Journal of Social Robotics*, 5(3), 409–417. <https://doi.org/10.1007/s12369-013-0197-8>
- MacDorman, K. F., & Ishiguro, H. (2006). The uncanny advantage of using androids in cognitive and social science research. *Interaction Studies*, 7(3), 297–337. <https://doi.org/10.1075/is.7.3.03mac>
- 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>
- Menne, I. M., & Schwab, F. (2018). Faces of emotion: investigating emotional facial expressions towards a robot. *International Journal of Social Robotics*, 10(2), 199–209. <https://doi.org/10.1007/s12369-017-0447-2>
- Montalan, B., Lelard, T., Godefroy, O., & Mouras, H. (2012). Behavioral investigation of the influence of social categorization on empathy for pain: A minimal group paradigm study. *Frontiers in Psychology*, 3, 389. <https://doi.org/10.3389/fpsyg.2012.00389>
- Mori, M., MacDorman, K., & Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine*, 19(2), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>
- Nomura, T., Kanda, T., & Suzuki, T. (2006). Experimental investigation into influence of negative attitudes toward robots on human–robot interaction. *AI & SOCIETY*, 20(2), 138–150. <https://doi.org/10.1007/s00146-005-0012-7>
- Nomura, T., Suzuki, T., Kanda, T., & Kato, K. (2006). Measurement of negative attitudes toward robots. *Interaction Studies*, 7(3), 437–454.

<https://doi.org/10.1075/is.7.3.14nom>

Oberman, L. M., Winkielman, P., & Ramachandran, V. S. (2007). Face to face: Blocking facial mimicry can selectively impair recognition of emotional expressions. *Social Neuroscience*, 2(3–4), 167–178. <https://doi.org/10.1080/17470910701391943>

Rae Westbury, H., & Neumann, D. L. (2008). Empathy-related responses to moving film stimuli depicting human and non-human animal targets in negative circumstances. *Biological Psychology*, 78(1), 66–74.

<https://doi.org/10.1016/J.BIOPSYCHO.2007.12.009>

Reeves, B., & Nass, C. I. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. CSLI Publications. Retrieved from <http://press.uchicago.edu/ucp/books/book/distributed/M/bo3618528.html>

Riek, L. D., Rabinowitch, T.-C., Chakrabarti, B., & Robinson, P. (2009). How anthropomorphism affects empathy toward robots. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction - HRI '09* (p. 245). New York, New York, USA: ACM Press. <https://doi.org/10.1145/1514095.1514158>

Rosenthal-von der Pütten, A. M., Krämer, N. C., Hoffmann, L., Sobieraj, S., & Eimler, S. C. (2013). An experimental study on emotional reactions towards a robot. *International Journal of Social Robotics*, 5(1), 17–34. <https://doi.org/10.1007/s12369-012-0173-8>

Seyama, J., & Nagayama, R. S. (2007). The uncanny valley: Effect of realism on the impression of artificial human faces. *Presence: Teleoperators and Virtual Environments*, 16(4), 337–351. <https://doi.org/10.1162/pres.16.4.337>

Strait, M. K., Aguillon, C., Contreras, V., & Garcia, N. (2017). The public's perception of humanlike robots: Online social commentary reflects an appearance-based uncanny valley, a general fear of a “Technology Takeover”, and the unabashed sexualization

- of female-gendered robots. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (pp. 1418–1423). IEEE.
<https://doi.org/10.1109/ROMAN.2017.8172490>
- Sun, Y.-B., Wang, Y.-Z., Wang, J.-Y., & Luo, F. (2015). Emotional mimicry signals pain empathy as evidenced by facial electromyography. *Scientific Reports*, 5(1), 16988.
<https://doi.org/10.1038/srep16988>
- Suzuki, Y., Galli, L., Ikeda, A., Itakura, S., & Kitazaki, M. (2015). Measuring empathy for human and robot hand pain using electroencephalography. *Scientific Reports*, 5(1), 15924. <https://doi.org/10.1038/srep15924>
- Syrdal, D. S., Dautenhahn, K., Koay, K. L., & Walters, M. L. (2009). The Negative Attitudes Towards Robots Scale and reactions to robot behaviour in a live human-robot interaction study. In *Proceedings of the 23rd Convention of the Society for the Study of Artificial Intelligence and Simulation of Behaviour, AISB 2009* (pp. 109–115). 23rd Convention of the Society for the Study of Artificial Intelligence and Simulation of Behaviour, AISB 2009; Edinburgh; United Kingdom; 6 April 2009 through 9 April 2009; Code 89166.
- Szollosy, M. (2017). Freud, Frankenstein and our fear of robots: projection in our cultural perception of technology. *AI & SOCIETY*, 32(3), 433–439.
<https://doi.org/10.1007/s00146-016-0654-7>
- Tan, J.-W., Walter, S., Scheck, A., Hrabal, D., Hoffmann, H., Kessler, H., & Traue, H. C. (2012). Repeatability of facial electromyography (EMG) activity over corrugator supercilii and zygomaticus major on differentiating various emotions. *Journal of Ambient Intelligence and Humanized Computing*, 3(1), 3–10.
<https://doi.org/10.1007/s12652-011-0084-9>
- Tarrant, M., Dazeley, S., & Cottom, T. (2009). Social categorization and empathy for

outgroup members. *British Journal of Social Psychology*, 48(3), 427–446.

<https://doi.org/10.1348/014466608X373589>

Tassinary, L. G., Cacioppo, J. T., & Vanman, E. J. (2016). The somatic system. In J. T.

Cacioppo, L. G. Tassinary, & G. G. Berntson (Eds.), *Handbook of Psychophysiology* (4th Ed., pp. 151–182). Cambridge: Cambridge University Press.

Tondu, B. (2012). Anthropomorphism and service humanoid robots: an ambiguous relationship. *Industrial Robot: An International Journal*, 39(6), 609–618.

<https://doi.org/10.1108/01439911211268840>

Wang, S., Lilienfeld, S. O., & Rochat, P. (2015). The uncanny valley: Existence and explanations. *Review of General Psychology*, 19(4), 393–407.

<https://doi.org/10.1037/gpr0000056>

Złotowski, J., Proudfoot, D., Yogeeswaran, K., & Bartneck, C. (2015).

Anthropomorphism: opportunities and challenges in human–robot interaction. *International Journal of Social Robotics*, 7(3), 347–360.

<https://doi.org/10.1007/s12369-014-0267-6>