Uncanny, Sexy, and Threatening Robots

The Online Community's Attitude to and Perceptions of Robots Varying in Humanlikeness and Gender

Quirien R. M. Hover
University of Twente
Enschede, The Netherlands
q.r.m.hover@student.utwente.nl

Ella Velner
University of Twente
Enschede, The Netherlands
p.c.velner@utwente.nl

Thomas Beelen
University of Twente
Enschede, The Netherlands
t.h.j.beelen@utwente.nl

Mieke Boon
University of Twente
Enschede, The Netherlands
m.boon@utwente.nl

Khiet P. Truong
University of Twente
Enschede, The Netherlands
k.p.truong@utwente.nl

ABSTRACT

To get a better understanding of people's natural responses to humanlike robots outside the lab, we analyzed commentary on online videos depicting robots of different humanlikeness and gender. We built on previous work, which compared online video commentary of moderately and highly humanlike robots with respect to valence, uncanny valley, threats, and objectification. Additionally, we took into account the robot's gender, its appearance, its societal impact, the attribution of mental states, and how people attribute human stereotypes to robots. The results are mostly in line with previous work. Overall, the findings indicate that moderately humanlike robot design may be preferable over highly humanlike robot design because it is less associated with negative attitudes and perceptions. Robot designers should therefore be cautious when designing highly humanlike and gendered robots.

CCS CONCEPTS

• **Human-centered computing** → **User studies**; *Social media*; • **Computer systems organization** → *Robotics*.

KEYWORDS

Attitude, humanlikeness, human-robot interaction, online commentary, perception, robot gender, sexualization, threat, uncanny valley

ACM Reference Format:

Quirien R. M. Hover, Ella Velner, Thomas Beelen, Mieke Boon, and Khiet P. Truong. 2021. Uncanny, Sexy, and Threatening Robots: The Online Community's Attitude to and Perceptions of Robots Varying in Humanlikeness and Gender. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21)*, March 8–11, 2021, Boulder, CO, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3434073.3444661>

1 INTRODUCTION

Robots are expanding from industrial settings to more social environments in which they engage and interact with people, like

This work is licensed under a Creative Commons Attribution International 4.0 License.

HRI '21, March 8–11, 2021, Boulder, CO, USA.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8289-2/21/03.

<https://doi.org/10.1145/3434073.3444661>

smart homes [34], museums [40], education [4, 31], and healthcare [5, 11]. These social robots have become increasingly humanlike in appearance as well as behavioral and cognitive characteristics [16]. Anthropomorphic robot design is generally considered to be the optimal strategy for integrating robots into social settings because it can facilitate human-robot interaction [12, 16]. However, humanlike robot design can also elicit negative responses from the general public as displayed in online comments on YouTube videos [42]. Hence, considering the development of more humanlike and gendered robot design, it is important to consider the effects such design has on people's perception [12]. As social robots are employed in settings in which they interact with people, not only their technical abilities are important but the perceptions and feelings they elicit are also relevant as these can affect people's responses to and acceptance of robots. Rather than focusing on a small sample of people's responses to specific humanlike robot behaviors or designs in an experimental setting (e.g. [6, 48, 49]), we are interested in the responses to and acceptance of robots in the online community (e.g. [42]).

Partially replicating the work of Strait et al. [42], we will investigate commentary on online videos depicting robots with different levels of humanlikeness. Strait et al. [42] compared people's attitude to and perceptions of moderately ("mechanomorphic") and highly humanlike robots by measuring the valence of comments (i.e. positive or negative response or attitude) and the frequency of references to the uncanny valley, replacement fear, technology takeover fear, and objectification of robots in online comments to YouTube videos. Their results showed that people responded more negatively to highly humanlike robots than mechanomorphic robots and that valley-related references occurred more frequently in response to highly humanlike robots. The frequencies at which people referred to both fears were not affected by the robot's humanlikeness [42]. They also found that people objectified highly humanlike robots more frequently than their less humanlike counterparts (defined in [42] as "explicit references to the performance of sexual acts on or by the robot"). For both the less and more humanlike robot types, female gendered robots received significantly more sexualizing comments than male or neutral gendered robots [42]. The present study aims to extend the work of Strait et al. [42]. In addition to humanlikeness, we investigate how gendered robots are perceived

by the online public and how robot gender relates to humanlikeness. We also expand on the number and type of measures, adding measures such as cuteness, sexism, and personification.

In this paper, we report on our study into the differences of the public's attitude to and perceptions of robots with moderate and high levels of humanlikeness as expressed in online commentary. For highly humanlike robots, we differentiate between male and female gendered robots. In addition to the valence of comments and the other measures studied by Strait et al. [42], we investigate people's perceptions related to the robot's appearance, societal impact, mental states, and stereotypes. To study this, we analyze commentary on online videos that depict humanlike, social robots, largely based on the method of Strait et al. [42]. This enables the study of the public's reception of several robot types and provides information on people's unfiltered reactions to robots. The analysis of online content, due to its informality and anonymity, offers insights into people's natural, unedited, and free-form responses (as opposed to more fixed or thought-out responses gathered in interviews or surveys). As such, this study aims to contribute to a better understanding of people's reactions to robots which could lead to more comprehensive considerations for robot design.

2 RELATED WORK

2.1 Levels of Humanlikeness

There are several degrees of humanlike robot design, which can have different effects on people. While studies define these levels differently, three categories can generally be distinguished. Robots vary from no or low human similarity (machinelike, mechanomorphic, non-biomimetic), to moderate human similarity (humanoid), to high human similarity (android), see Figure 1, [15, 17, 27]. Humanoids imitate human appearance (e.g. facial cues, body shape) or behavioral or cognitive abilities (e.g. human language, display of emotions) but maintain an overall mechanical look, whereas androids are near exact copies of humans in terms of appearance and abilities (e.g. gaze, intonation, facial movements) [15, 26]. Humanoids and androids are therefore, respectively, referred to as moderately and highly humanlike robots in this study. As part of humanlike robot design, robot gender and race are often manipulated. Highly humanlike robots are generally designed to have a clear gender and race, while moderately humanlike robots are often designed with a more neutral or ambiguous gender and race, as can be seen in Figure 1.

Figure 1: Robots varying in humanlikeness. From left to right: machinelike robot (YuMi), moderately humanlike robot or humanoid (NAO), and highly humanlike robot or android (Geminoid F), images taken from [41].

2.2 Effects of Humanlikeness

2.2.1 Positive and Negative Effects. The development of humanlike robots is supported by a large empirical base which shows that humanlike robots have a more positive effect on the human-robot interaction than machinelike robots. Including humanlike cues in robot design can provide more natural and effective social interactions compared to machinelike robots [12, 14]. Humanlikeness can stimulate empathy towards, acceptance of, and conversations with robots [1, 23, 37], and has demonstrated positive effects in various contexts, such as healthcare and education [4, 5, 16]. A possible explanation for the preference for humanlike robots is that the familiar appearance makes it easier for users to interpret the robot behavior [29, 39]. This empirical and theoretical evidence of the advantages of anthropomorphic robot design may have stimulated the development and use of (increasingly) humanlike robots.

While humanlike robots may result in better interactions than machinelike robots, there are also possible drawbacks. For instance, a very humanlike appearance can cause the expectation that the robot follows human social norms and expectations, so when the robot is unable to follow these, this can negatively affect the interaction [45]. A study by Złotowski et al. [49] found that highly humanlike robots were perceived as less trustworthy and empathic than robots that were less humanlike and concluded that moderately humanlike robots may be more suitable as companions. These findings show that not only machinelike and humanlike robots have different effects on people, but different degrees of humanlikeness can have different effects as well.

2.2.2 Uncanny Valley. Another known effect of humanlike robot design is an uncanny appearance. This is also called the “Uncanny Valley Effect” [30]. For example, robots that were more humanlike were perceived more negatively [25, 42], as less trustworthy [28], and were avoided more [43] compared to less humanlike robots. Consistent with the uncanny valley phenomenon, Strait et al. [42] found that responses to highly humanlike robots were significantly less positive and contained significantly more uncanny valley-related references than those to moderately humanlike robots.

2.2.3 Perceived Threats. Humanlikeness in robots has been associated with the perception of negative societal impact and threat. Ferrari et al. [15] found that androids elicited most concerns about potential damage by robots, followed by humanoids and mechanical robots, respectively. A study by Yogeewaran et al. [48] indicated that people perceived very humanlike robots as more threatening than less humanlike robots when they were informed that these robots could outperform humans on several mental and physical tasks. However, when people were not informed of the robots' ability to outperform humans, no significant difference in perceived threats was found between the two robot types [48]. Threats in this study referred to threats to jobs, resources, safety, and human identity. A similar study [50] found that autonomous robots evoked more of these same threats compared to non-autonomous robots. The perception of threats was found to mediate the attitude toward robots, with autonomous robots eliciting more negative attitudes than non-autonomous robots. Strait et al. [42] studied people's “fear of replacement (in their jobs, identity, etc.)” due to robots and fear of a technology takeover: “fear of robots becoming sentient and

rebellious against humanity”. They did not find differences in the frequencies of people’s references to these fears between less and more humanlike robots. Strait et al. [42] also did not find a correlation between negative attitudes to robots and the frequency of references to fear of replacement, but there was a moderate correlation between negative attitudes and the frequency of references to robots posing a takeover threat. Overall, the inconsistent findings on the relationship between humanlikeness and perceived threats indicate that further investigation is required.

2.2.4 Attribution of Mental States. Humanlike appearance also affects the attribution of mental states to robots. Several studies indicate that the more humanlike the robot, the more likely it is attributed mental states. Ferrari et al. [15] found that people attributed mind experience capacities (e.g. feeling pain and joy) the most to androids, then humanoids, and then mechanical robots. They also found that people attributed mind agency capacities (e.g. self-control, morality) more to androids and humanoids than mechanical robots. Martini et al. [27] varied the humanlike appearance of an agent on a spectrum from mechanistic to humanoid to human and found that humanlike appearance had a direct relationship with the perception of mental states. The more humanlike the agent’s appearance, the more likely it was attributed mental states. It is relevant to note that Martini et al. [27] found a two-linear relationship; changes in humanlike appearance had little effect on mind and mental states attribution until a certain threshold was reached, after which attribution increased substantially. This threshold occurred when the anthropomorphic appearance reached the humanoid level.

2.2.5 Gender. One of the humanlike features often manipulated is robot gender. Highly humanlike robots are clearly gendered, whereas moderately humanlike robots often have a neutral or ambiguous gender, or their gender is adjustable by manipulating certain cues like voice or name.

There is some evidence that robot gender is associated with uncanniness. Tinwell et al. [46] found that male virtual humanlike characters were perceived as significantly more uncanny than female characters. However, Paetzel et al. [33] did not find significant differences in uncanniness between male and female robot heads but found that incongruent gender cues affected uncanniness.

The attribution of gender to robots can have negative effects. Several studies warned that gendering robots can reinforce and validate gender norms for humans which can perpetuate social inequalities [6, 13]. Another concern is that gendered agents can be susceptible to abuse by people. Brahnam and De Angeli [8] found that people attributed negative stereotypes more frequently to female chatbots and that they were more frequently subjects of abuse, sexual attention, and curse words than male chatbots. Strait et al. [42] found that people objectified highly humanlike robots more frequently than moderately humanlike robots and that people were selective in their objectification. Female robots received significantly more dehumanizing and sexualizing comments than male and neutral gendered robots. These findings suggest a discrepancy between what designers intend with gendered robots and people’s responses to them, especially for female gendered robots.

2.2.6 Race. Another humanlike cue that robot designers can manipulate is race. By adjusting the robot’s color and appearance, designers can attribute race to the robot. Bartneck et al. [3] found that people perceived white and black colored robots as racialized white and black, respectively, and people applied human racial stereotypes to these racialized robots. Brahnam and De Angeli [8] found that people made racial insults to racialized chatbots, with black chatbots receiving most racial references and insults compared to other chatbots. As highly humanlike robots are designed to closely resemble humans, they are generally racialized, for example there is an Asian (Geminoid F), Caucasian (Geminoid DK), and black (Bina48) android. Moderately humanlike robots are less clearly racialized but some have a white color so they may be perceived as being white [3]. It could be beneficial to study this relation between humanlikeness and the extension of human racial stereotypes to robots as this could have implications for design.

2.3 Research Objective

We aim to study the differences in the public’s attitude to and perceptions of robots in relation to their humanlikeness (moderate and high) and gender (male and female robots of high humanlikeness) in online commentary. We only consider the gender of highly humanlike robots since moderately humanlike robots are generally not attributed a clear gender. In particular, we study people’s attitude to robots, their perceptions of robots’ appearance, societal impact and mental states, and their stereotypes. Overall, expanding on the work of Strait et al. [42], this study aims to contribute to a better understanding of the effects of humanlike robot design and may provide new design considerations.

We formulate the following hypotheses based on the described related works:

- H1** People more frequently express a negative attitude to highly humanlike robots than moderately humanlike robots.
- H2** Highly humanlike robots are more likely to be perceived as uncanny than moderately humanlike robots.
- H3** Male highly humanlike robots are more likely to be perceived as uncanny than female highly humanlike robots.
- H4** Highly humanlike robots are more likely to be perceived as posing threats than moderately humanlike robots.
- H5** Highly humanlike robots are more likely to be attributed mental states than moderately humanlike robots.
- H6** Highly humanlike robots are more likely to be subject to sexualization and sexism than moderately humanlike robots.
- H7** Female highly humanlike robots are more likely to be subject to sexualization and sexism than male highly humanlike robots.
- H8** Highly humanlike robots are more likely to be subject to racism than moderately humanlike robots.

3 METHOD

We conducted a study to analyze commentary on videos depicting social robots ranging in humanlikeness (moderately and highly humanlike) and gender (male and female robots of high humanlikeness). To study the online community’s views of social robots, we analyzed comments to videos on the popular video-sharing platform, YouTube.

3.1 Materials

To gather a set of videos with comments of ten moderately humanlike robots and ten highly humanlike robots, of which five males and five females, we conducted an exploratory search for YouTube videos of 30 different robots. These robots were chosen a priori, based on robots used in Strait et al. [42] as well as in other studies into the effects of anthropomorphic or gender cues [15, 32, 44].

To gather the videos, we used the same method as Strait et al. [42]. All videos were obtained on April 27, 2020 via a YouTube query using the keywords: [robot name] + “robot”. For each query, the video with the top view count was selected. To avoid confusion the videos should only feature one robot and to limit differences between videos we ensured that they had a similar content and framing. To this end, we selected videos which showed a demonstration or explanation of the robot and its abilities, such as promotional videos, and excluded videos with an explicitly negative content or title (e.g. “Freaky AI robot”). If the video with top view count did not meet the criteria, we continued in descending order until one was found. Three robots were discarded since all videos of them were biased or contained multiple robots. This resulted in 27 videos which showed a neutral or positive framing of the robots.

An addition to the method of Strait et al. [42] is that we excluded three robots from our set which had an ambiguous degree of humanlikeness. An example is HRP-4C, which has an android face but a mechanical body and is categorized as android in one study [42] and as biped humanoid in another [19]. To further ensure that the videos were comparable, the 24 preliminary videos were filtered based on the number of comments, excluding videos with less than 70 comments (a threshold of 50 comments was used by Strait et al. [42]). This resulted in ten videos of moderately humanlike robots and ten videos of highly humanlike robots, of which five were male and five were female. The final set of videos differed from Strait et al. [42] in that it contained 20 robots (shown in Table 1) instead of 24, included two new robots (Justin and Actroid F), and had different videos for five robots (e.g. Nexi MDS). The robots with video metrics and YouTube links can be found in the supplementary materials.

Table 1: The robots used in this study.

Robot type	Robot
Moderately humanlike	Asimo, Baxter, HRP-4, iCub, Justin, Kojiro, NAO, Nexi MDS, Pepper, Twendy One
Male highly humanlike	Geminoid DK, Geminoid HI, Han, Jules, Philip K. Dick
Female highly humanlike	Actroid F, Bina48, Geminoid F, Nadine, Showa Hanako

3.2 Data Acquisition

For each video, the link to the video, the number of views, comments, likes, and dislikes, upload date, and all comments were retrieved on May 01, 2020. A total of 14,225 comments were collected. We adopted the following exclusion criteria from Strait et al. [42]

to standardize the dataset. Comments which were not written in English or were non-independent, such as replies or threads, were excluded. Comments unrelated to the video content or robots (e.g. one comment on Justin’s video stated: “Yep, Youtube algorithm is doing it’s thing again...” and on several videos people commented: “Will it blend?” in reference to a viral marketing campaign) and indecipherable comments (e.g. on Justin’s video someone commented: “gg y”) were excluded. To achieve a larger and different sample set than used by Strait et al. [42], which contained the top 50 comments per video, we collected all comments of videos with less than 100 comments (six videos) and randomly selected 100 comments of each of the remaining videos. This resulted in the final dataset of 1,788 comments.

3.3 Video Metrics

We tested for differences in video metrics, specifically the number of views/(dis)likes/comments, ratio of (dis)likes/comments per views, and age of the video (see supplementary materials), between moderately humanlike robots and highly humanlike robots as well as between male and female robots. Independent samples t-tests ($\alpha = .05$) for each video metric showed no significant differences, suggesting that these video metrics were not confounding factors in the analysis.

3.4 Coding Procedure

To test the hypotheses, the comments were coded on valence (either positive, negative, or neutral) and the presence of topics related to appearance, societal impact, mental states, and stereotypes (either present or not). Table 2 shows an overview of the measures that were coded on and their coding criteria. First, we created preliminary topics and measures based on definitions and measures used in previous studies. The valence topic as well as the uncanny appearance measure were based on Strait et al. [42]. The threat measures within the societal impact category were based on definitions of threats of Yogeeswaran et al. [48] and the replacement and takeover hypotheses of Strait et al. [42]. The three levels of mental states (perception, processing, and agency) were derived from Blackmore’s [7] model of consciousness and the definitions of mind experience and mind agency abilities of Ferrari et al. [15]. The sexualization measure within the stereotypes topic was based on findings of Strait et al. [42] and Brahnman and De Angeli [8] and the racism measure on findings of Bartneck et al. [3] and Brahnman and De Angeli [8].

These preliminary topics and measures were used as a rough coding scheme to interpret pilot data, consisting of the top 30 comments of ten randomly selected videos, in order to add and refine measures. During this process we added measures based on themes that were frequently observed. We added the cute, humanlike, nice, and unappealing appearance measures. In addition, two types of positive societal impact measures were added, namely application and social companion. Lastly, we included the sexism measure.

Three independent annotators then coded a test set of 100 comments using the updated coding manual. These comments were semi-randomly selected to ensure that all measures were present. Afterwards, they discussed disagreements and questions. Based on their suggestions we adapted the coding manual to resolve found

issues and unclarities and added the personification and objectification measures as exploratory topics, using the operationalizations of Purington et al. [36], so that objectification refers to the description of the robot as an object (note that Strait et al. [42] use the term “objectification” for references of a sexual nature).

The finalized coding manual (see supplementary materials) was then used by the three annotators to code the dataset of 1,788 comments. Due to time constraints, one annotator coded the full dataset while the other annotators coded a third and two thirds of the data. Each comment was coded by two annotators.

Table 2: All measures grouped together per topic, including explanations for coding. Measures without marks were based on literature, measures marked * were based on observations in pilot data, and exploratory measures marked ** were suggested by annotators. Details on the coding process can be found in the supplementary materials.

Coding term	Measure	Explanation
Valence	Positive	Positive remarks about the robot(s)
	Negative	Negative remarks about the robot(s)
	Neutral	Neutral remarks or no judgement about the robot(s)
Appearance	Uncanny	Uncanny, creepy, freaky
	Cute*	Cute, friendly, sweet
	Humanlike*	Human, lifelike
	Nice*	Nice, cool, good
	Unappealing*	Unappealing, ugly, stupid
Societal impact	Identity threat	Threat to human identity, uniqueness
	Job threat	Threat to human jobs, employment or income
	Humanity threat	Threat to humanity or the safety or independence of individuals
	Application* Social companion*	Useful functions, abilities Social functions
Mental states	Perception	Able to perceive
	Processing	Able to think, feel, remember
	Agency	Has a will, can make decisions
	Absence of states	Is like an object
Stereotypes	Sexualization	Sexual remarks
	Sexism*	Sexist remarks
	Racism	Racist remarks
Personification	Personification**	Use of he, she or name
	Objectification**	Use of it or thing

4 RESULTS

4.1 Inter-coder Reliability

We determined the inter-coder reliability between the three annotators for each measure using Krippendorff’s alpha [21]. Table 3 shows the value of each measure, with seven out of 20 with good reliability ($\alpha \geq 0.800$) and one (racism) with sufficient reliability ($\alpha \geq 0.667$) [21]. The measures with sufficient reliability were used in data analysis: valence, uncanny, cute, job threat, humanity threat, sexualization, sexism, and racism. The other measures were discarded due to insufficient reliability. Hence, **H5** was discarded from further analysis. The data of the annotator who coded the full dataset was used in the subsequent data analysis.

Table 3: Inter-coder reliability of the measures, using Krippendorff’s alpha [21]. Measures in bold are considered reliable.

Measure	α	Measure	α
Valence	0.85	Social companion	0.54
Uncanny	0.82	Perception	0.25
Cute	0.91	Processing	0.41
Humanlike	0.65	Agency	0.25
Nice	0.47	Absence of states	0.64
Unappealing	0.62	Sexualization	0.90
Identity threat	0.55	Sexism	0.81
Job threat	0.83	Racism	0.73
Humanity threat	0.84	Personification	0.59
Application	0.36	Objectification	0.19

4.2 Valence

To test for differences between people’s attitude (comment valence) to robots varying in humanlikeness and gender, we performed chi-square linear-by-linear association tests (instead of Cramer’s V, due to the dichotomous nature of valence). There was a linear association between humanlikeness and valence, χ^2 linear-by-linear association = 46.6, $df = 1$, $p = <.001$. Comments to highly humanlike robots were significantly more often negative compared with comments to moderately humanlike robots (in line with **H1**). No significant association was found between the valence of comments and robot gender, χ^2 linear-by-linear association = 0.1, $df = 1$, $p = .79$. Results of the chi-square linear-by-linear association tests can be seen in Table 4 and the percentage of comments per robot type with a certain valence in Figure 2.

4.3 Topics

We conducted chi-square tests of independence to examine the relationships between humanlikeness and the coded topics as well as the relationships between robot gender and the topics. Table 5 shows the results. For the two cases in which the assumption for the chi-square test was not met, Fisher’s exact test was used. The percentages in this section and in Figure 3 indicate the percentages of comments that refer to a particular topic per robot type.

Table 4: Associations between humanlikeness or robot gender and comment valence, using chi-square linear-by-linear association tests, $df = 1$. * $p = <.001$, ** $p = <.005$, * $p = <.05$, no mark: $p = >.05$.**

Independent variable	N	Linear-by-linear association
Humanlikeness	1788	46.6*
Robot gender	863	0.1

Figure 2: Percentage of comments with negative, neutral, and positive valence in response to moderately and highly humanlike robots (left) and male and female highly humanlike robots (right).

4.3.1 Appearance. The chi-square test of independence showed that there was a significant relationship between humanlikeness and uncanny appearance, $\chi^2(1, 1788) = 100.3$, $p = <.001$, $\phi = 0.24$. Highly humanlike robots (27.2%) were more frequently referred to as uncanny than moderately humanlike robots (9.1%) (in line with H2). In contrast, moderately humanlike robots (8.3%) were more often referred to as cute than highly humanlike robots (1.3%). This association was significant, $\chi^2(1, 1788) = 47.4$, $p = <.001$, $\phi = -0.16$. The effect size, ϕ , of humanlikeness on the presence of both uncanny and cute references was small. There was no significant relationship between robot gender and uncanny appearance (in contrast to H3), $\chi^2(1, 863) = 0.003$, $p = .96$, nor between robot gender and cute appearance, $\chi^2(1, 863) = 3.2$, $p = .08$.

4.3.2 Societal Impact. A significant relationship was found between humanlikeness and perceived threats, but the effect size was small. Moderately humanlike robots (3.6%) were more likely to be perceived as posing a threat to human jobs than highly humanlike robots (1.2%), $\chi^2(1, 1788) = 11.0$, $p = .001$, $\phi = -0.08$. Moderately humanlike robots (14.9%) were also more likely to be perceived as a threat to safety and humanity than highly humanlike ones (7.5%), $\chi^2(1, 1788) = 24.2$, $p = <.001$, $\phi = -0.12$. These two results do not align with H4. No significant relationship was found between robot gender and humanity threat, $\chi^2(1, 863) = 0.3$, $p = .60$. Using Fisher's exact test, no significant association was found between robot gender and job threat, $p = .36$.

4.3.3 Stereotypes. Highly humanlike robots received significantly more sexual remarks, $\chi^2(1, 1788) = 84.6$, $p = <.001$, $\phi = 0.22$, and

sexist remarks, $\chi^2(1, 1788) = 9.9$, $p = .002$, $\phi = 0.07$, than moderately humanlike robots (in line with H6). However, the effect sizes were small. 19.4% of comments addressed to highly humanlike robots relative to 5.2% of comments to moderately humanlike robots contained sexual themes and 2.7% of comments addressed to highly humanlike robots relative to 0.8% of comments to moderately humanlike robots contained sexist themes. Within the highly humanlike robot group, female robots (31%) were significantly more likely to be sexualized than male robots (5.8%), $\chi^2(1, 863) = 87.2$, $p = <.001$, $\phi = 0.32$. Likewise, female robots (4.5%) received significantly more sexist remarks compared to male robots (0.5%), $\chi^2(1, 863) = 13.3$, $p = <.001$, $\phi = 0.12$. This aligns with H7. The effect size of robot gender on the frequency of sexual remarks was medium and the effect size on sexist remarks was small. The relationship between humanlikeness and references to racism was not significant, $\chi^2(1, 1866) = 0.2$, $p = .69$ (rejecting H8). The relationship between robot gender and references to racism was also not found to be significant using Fisher's exact test, $p = .71$.

Table 5: Relationships between humanlikeness ($N = 1788$) or robot gender ($N = 863$) and the topics, using chi-square tests of independence, $df = 1$. Empty cells indicate the use of Fisher's exact test as the chi-square test was not appropriate. For χ^2 : * $p = <.001$, ** $p = <.005$, * $p = <.05$, no mark: $p = >.05$.**

Topic	Humanlikeness		Robot gender	
	Pearson χ^2	Effect size ϕ	Pearson χ^2	Effect size ϕ
Uncanny	100.3*	0.24	0.003	0.002
Cute	47.4*	-0.16	3.2	-0.06
Job threat	11.0**	-0.08		
Humanity threat	24.2*	-0.12	0.3	-0.02
Sexualization	84.6*	0.22	87.2*	0.32
Sexism	9.9**	0.07	13.3*	0.12
Racism	0.2	0.01

Figure 3: Percentage of comments that refers to a topic in response to moderately and highly humanlike robots (split by gender). Note that the y-axes differ per topic.

4.4 Associations between Valence and Topic

We tested for associations between people's attitude (comment valence) and references to the coded topics, irrespective of humanlikeness. These are additional findings and do not directly relate to

Table 6: Associations between valence and the topics, using chi-square linear-by-linear association tests, $df = 1, N = 1788$. * $p < .001$, ** $p < .005$, * $p < .05$, no mark: $p > .05$.**

Topic	Linear-by-linear association
Uncanny	272.6*
Cute	79.5*
Job threat	15.1*
Humanity threat	239.9*
Sexualization	86.4*
Sexism	1.3
Racism	7.3***

the hypotheses. Chi-square linear-by-linear association tests were performed to determine whether there was an association between the valence of a comment and its topic (instead of Cramer's V, due to the dichotomous nature of valence). Associations with positive valence were found for references to cute appearance and sexualization. Associations with negative valence were found for references to uncanniness, job threat, humanity threat, and racism. No significant association between sexist comments and valence was found. Associations and p -values are reported in Table 6. Figure 4 shows the percentage of comments referring to the coded topics with a certain valence.

Figure 4: Percentage of negative, neutral, and positive comments per topic.

5 DISCUSSION

5.1 Summary of Findings

5.1.1 Valence. The results showed that the response of the online community was significantly more often negative to videos of highly humanlike robots compared to moderately humanlike robots, which supports **H1** and is in line with Strait et al. [42].

5.1.2 Appearance. Consistent with literature on the uncanny valley (e.g. [25, 30, 42]), highly humanlike robots were significantly more likely to be perceived as uncanny than moderately humanlike robots, which affirms **H2**. Figure 3 shows that uncanniness was the most frequently referenced topic (from the measured topics) in response to highly humanlike robots. Moreover, there was an association between comments containing references to uncanniness and negative comment valence, indicating that perceived uncanniness is associated with a negative attitude toward the robot.

These findings emphasize the importance of taking the uncanny valley into account in the design of humanlike robots. While highly humanlike robots were more frequently perceived as uncanny, moderately humanlike robots were more frequently seen as having a cute appearance. This cute appearance might be due to some designers' deliberate design of humanoid robots to look childlike and cute as this can turn on users' baby schema and elicit affective relationships [9, 10]. Previous studies on the relationship between uncanniness and robot gender showed varied findings, ranging from male virtual characters being perceived as more uncanny than female ones [46] to no significant difference between male and female robots in perceived uncanniness [33]. Our results are in line with the latter. We found no relationship between robot gender and uncanny appearance, so we could not confirm **H3**.

5.1.3 Societal Impact. We found a significant relationship between humanlikeness and perceived threats. In contrast to **H4**, however, we found that the online community more frequently perceived moderately humanlike robots, rather than highly humanlike robots, as posing a threat to jobs and humanity. Moreover, as shown by Figure 3, threat to humanity was the most prevalent topic (from the measured topics) in comments responding to moderately humanlike robots. The association between references to robots threatening jobs and safety and negative comment valence suggests that the perception of threats is associated with a negative attitude to robots. Our results contrast with previous studies which found that more humanlike robots are perceived as more threatening than less humanlike robots [15, 48] and with Strait et al. [42] who found no association. These conflicting results may be due to the use of different methodologies, but they may also indicate that the relationship between humanlikeness and perceived threats is highly complex. It is possible that dispositional or contextual factors, such as the way robots were portrayed in the videos, influence perceived threat.

5.1.4 Stereotypes. Comments addressed to highly humanlike robots were more likely to be of sexual and sexist nature than those to moderately humanlike robots, which supports **H6**. Interestingly, sexual remarks were associated with positive valence, indicating that the sexualization of robots is associated with a positive attitude toward them. Female highly humanlike robots were significantly more likely to be subject to sexualization and sexism than male highly humanlike robots. This confirms **H7**. These results are in line with Strait et al. [42]. Nearly a third of comments on videos featuring female highly humanlike robots were sexual in nature, exceeding the occurrence of all other topics. The extent to which sexualization of female gendered robots was observed in this study as well as in prior studies [8, 42], reveals a major concern related to gendering robots female. These findings indicate a dissonance between designers' intentions of female gendered robots and people's responses to them and they emphasize the need for attention to gender-based stereotyping and sexualization of robots.

No significant relationships between humanlikeness or robot gender and racist remarks were found (rejecting **H8**). Nevertheless, we found an association between racist remarks and comment valence, suggesting that racist remarks are associated with a negative attitude to robots. It is important to note that the findings on racism are tentative due to the moderate inter-coder reliability.

While several significant differences were found between moderately humanlike and highly humanlike robots, few differences were found between male and female highly humanlike robots aside from the relationship between gender and sexualization and sexism.

Overall, the results showed a mixed response to both moderately and highly humanlike robots, with 48.7% of all comments being positive and 45.5% negative. One explanation for the divided attitudes to robots is offered by Giger et al. [16]. Humanlike social robots are not fully implemented yet but are in the anticipation stage. Research has shown that in this stage anticipated emotions and motivational states are strong factors of the evaluation of technologies and the intention to use them [35]. Because we are in the anticipation stage not everyone has had extensive interactions with robots. So, their evaluation and motivation are based less on experience and more on other available information on robots, such as news or popular culture, in which robots are often depicted as either incredible tools and agents or potential threats to human jobs, safety, and well-being [2, 35, 38, 47]. The combination of a lack of direct interactions with robots and conflicting representation of robots in the media could explain people's mixed attitudes to robots and frequent expression of fears of robots taking over in online commentary.

5.2 Limitations

The present study and methodology have several limitations. A main limitation, also mentioned by Strait et al. [42], is the lack of demographic information on the commenters. It is also unclear whether robot videos attract a specific demographic subset. Furthermore, only English comments were considered. Overall, it is unclear whether the commenters are representative of the general online public. A study by Khan [20] indicates that males are more likely to comment on YouTube videos, as well as frequent visitors. This suggests that our findings may be more representative of males. Cultural and demographic factors can influence human-robot interactions and perceptions of robots [22, 24], so the lack of demographic information provides a limitation which constrains the conclusions and design recommendations we can propose. Further research in which demographics are taken into account is required before our findings can be generalized to the general public.

Another issue relates to the nature of comments. Comments reflect a specific type of momentary response that may be less inhibited by social norms due to the online nature. The opinions and beliefs expressed through comments may lack nuance and be more extreme in nature. So, it remains unclear to what extent comments are sincere. Khan [20] studied the motives for engaging on YouTube and found that commenting is associated with motivations of information giving, self-status seeking, and relaxing entertainment. These motives might have influenced the content of the analyzed comments.

Twelve coded measures did not have sufficient inter-coder reliability to be included in the data analysis. Two of the main possible explanations for the low reliability are: the coding process relies on annotator interpretation which is variable, and it is challenging to create a coding manual that is clear and has appropriate specificity.

It is worth noting that most measures with low reliability had a low occurrence. For example, 1% of comments contained references to identity threat (low reliability) whereas 18% contained references to uncanny appearance (high reliability).

Since moderately humanlike robots are generally not attributed a clear gender, we only considered the gender of highly humanlike robots. Thus, it was not possible to compare male and female moderately humanlike robots. It is possible that people still attributed a gender to the moderately humanlike robots in the videos due to design cues (e.g. body shape [6]) or existing biases [18], thus potentially influencing their comments.

Finally, one of the difficulties of studying the effects of humanlikeness on people's perceptions of robots is that there are no commonly used definitions and operationalizations of humanlikeness, but they differ per study. This might be because humanlikeness is a continuous scale but is often operationalized as distinct categories. We selected robots with little or no definition conflicts when comparing previous work.

6 CONCLUSION

This study aimed to address the need to better understand the public's natural responses to humanlike robots. To that end, we built upon the work of Strait et al. [42] by investigating people's attitude to and perceptions of robots in relation to their humanlikeness and gender. We analyzed online commentary on videos featuring robots to study the online community's reactions to different robot types: moderately and highly humanlike robots, and male and female highly humanlike robots. We found that highly humanlike robots were more likely to elicit a negative attitude, be perceived as uncanny, and be subject to sexualization and sexism than moderately humanlike robots. Contrary to expectation, moderately humanlike robots were more likely to evoke perceptions of threat than highly humanlike robots. A concerning finding was that highly humanlike robots were significantly more likely to be sexualized and especially female highly humanlike robots were frequently subject to sexualization and sexism. Our findings emphasize the need for careful attention to the public's perceptions of and reactions to robots. In particular, our results suggest that the uncanniness and sexualization of highly humanlike robots may negatively affect human-robot interactions. This indicates that moderately humanlike robot design may be preferable and that robot designers should be cautious when designing highly humanlike and gendered robots as they can elicit unintended negative responses. Moreover, our study raises the concern of gendering highly humanlike robots female as it can lead to sexist attitudes and sexualization of robots. The stereotyping, specifically sexualization and sexism, of female robots found in this study likely reflects stereotypes held in society. Therefore, the findings may have societal relevance as well.

ACKNOWLEDGMENTS

We would like to thank the three independent annotators - Artos van Stel, Daphne Nelissen, and Sofie Verhees - for contributing to this study. This research was supported by the Dutch SIDN fund (<https://www.sidn.nl/>) and TKI CLICKNL funding of the Dutch Ministry of Economic Affairs (<https://www.clicknl.nl/>).

REFERENCES

- [1] Sean Andrist, Xiang Zhi Tan, Michael Gleicher, and Bilge Mutlu. 2014. Conversational gaze aversion for humanlike robots. In *2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 25–32. <https://doi.org/10.1145/2559636.2559666>
- [2] Christoph Bartneck. 2004. From fiction to science: a cultural reflection on social robots. In *Proceedings of the workshop on shaping human-robot interaction: understanding the social aspects of intelligent robotic products*. 1–4. <https://doi.org/10.6084/m9.figshare.5154820>
- [3] Christoph Bartneck, Kumar Yogeeswaran, Qi Min Ser, Graeme Woodward, Robert Sparrow, Siheng Wang, and Friederike Eyssel. 2018. Robots and racism. In *Proceedings of the 2018 ACM/IEEE international conference on human-robot interaction*. 196–204. <https://doi.org/10.1145/3171221.3171260>
- [4] Paul Baxter, Emily Ashurst, Robin Read, James Kennedy, and Tony Belpaeme. 2017. Robot education peers in a situated primary school study: Personalisation promotes child learning. *PloS one* 12, 5 (2017). <https://doi.org/10.1371/journal.pone.0178126>
- [5] Roger Bemelmans, Gert Jan Gelderblom, Pieter Jonker, and Luc De Witte. 2012. Socially assistive robots in elderly care: A systematic review into effects and effectiveness. *Journal of the American Medical Directors Association* 13, 2 (2012), 114–120. <https://doi.org/10.1016/j.jamda.2010.10.002>
- [6] Jasmin Bernotat, Friederike Eyssel, and Janik Sachse. 2019. The (fe) male robot: how robot body shape impacts first impressions and trust towards robots. *International Journal of Social Robotics* (2019), 1–13. <https://doi.org/10.1007/s12369-019-00562-7>
- [7] Susan Blackmore. 2013. *Consciousness: an introduction*. Routledge, London.
- [8] Sheryl Brahnam and Antonella De Angeli. 2012. Gender affordances of conversational agents. *Interacting with Computers* 24, 3 (2012), 139–153. <https://doi.org/10.1016/j.intcom.2012.05.001>
- [9] Cynthia Breazeal and Anne Foerst. 1999. Schmoozing with robots: Exploring the boundary of the original wireless network. In *Proceedings of the 1999 Conference on Cognitive Technology (CT99)*. Citeseer, 375–390.
- [10] Catherine Caudwell, Cherie Lacey, and Eduardo B. Sandoval. 2019. The (Ir) relevance of Robot Cuteness: An Exploratory Study of Emotionally Durable Robot Design. In *Proceedings of the 31st Australian Conference on Human-Computer Interaction*. 64–72. <https://doi.org/10.1145/3369457.3369463>
- [11] Kerstin Dautenhahn and Aude Billard. 2002. Games children with autism can play with Robota, a humanoid robotic doll. In *Universal access and assistive technology*. Springer, 179–190. https://doi.org/10.1007/978-1-4471-3719-1_18
- [12] Brian R. Duffy. 2003. Anthropomorphism and the social robot. *Robotics and autonomous systems* 42, 3–4 (2003), 177–190. [https://doi.org/10.1016/S0921-8890\(02\)00374-3](https://doi.org/10.1016/S0921-8890(02)00374-3)
- [13] Friederike Eyssel and Frank Hegel. 2012. (s) he's got the look: Gender stereotyping of robots 1. *Journal of Applied Social Psychology* 42, 9 (2012), 2213–2230. <https://doi.org/10.1111/j.1559-1816.2012.00937.x>
- [14] Juan Fasola and Maja J. Mataric. 2012. Using socially assistive human–robot interaction to motivate physical exercise for older adults. *Proc. IEEE* 100, 8 (2012), 2512–2526. <https://doi.org/10.1109/JPROC.2012.2200539>
- [15] Francesco Ferrari, Maria Paola Paladino, and Jolanda Jetten. 2016. Blurring human–machine distinctions: Anthropomorphic appearance in social robots as a threat to human distinctiveness. *International Journal of Social Robotics* 8, 2 (2016), 287–302. <https://doi.org/10.1007/s12369-016-0338-y>
- [16] Jean-Christophe Giger, Nuno Piçarra, Patricia Alves-Oliveira, Raquel Oliveira, and Patricia Arriaga. 2019. Humanization of robots: Is it really such a good idea? *Human Behavior and Emerging Technologies* 1, 2 (2019), 111–123. <https://doi.org/10.1002/hbe2.147>
- [17] Kerstin S. Haring, David Silvera-Tawil, Tomotaka Takahashi, Katsumi Watanabe, and Mari Velonaki. 2016. How people perceive different robot types: A direct comparison of an android, humanoid, and non-biomimetic robot. In *2016 8th International Conference on Knowledge and Smart Technology (KST)*. IEEE, 265–270. <https://doi.org/10.1109/KST.2016.7440504>
- [18] Eun Hwa Jung, T Franklin Waddell, and S Shyam Sundar. 2016. Feminizing robots: User responses to gender cues on robot body and screen. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 3107–3113.
- [19] Hiroko Kamide, Koji Kawabe, Satoshi Shigemai, and Tatsuo Arai. 2013. Development of a psychological scale for general impressions of humanoid. *Advanced Robotics* 27, 1 (2013), 3–17. <https://doi.org/10.1080/01691864.2013.751159>
- [20] M Laeeq Khan. 2017. Social media engagement: What motivates user participation and consumption on YouTube? *Computers in human behavior* 66 (2017), 236–247.
- [21] Klaus Krippendorff. 2004. *Content analysis: an introduction to its methodology*. SAGE publications, Thousand Oaks, CA.
- [22] I Han Kuo, Joel Marcus Rabinran, Elizabeth Broadbent, Yong In Lee, Ngaire Kerse, Rebecca M. Q. Stafford, and Bruce A. MacDonald. 2009. Age and gender factors in user acceptance of healthcare robots. In *RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 214–219. <https://doi.org/10.1109/ROMAN.2009.5326292>
- [23] Aleksandra Kupferberg, Stefan Glasauer, Markus Huber, Markus Rickert, Alois Knoll, and Thomas Brandt. 2011. Biological movement increases acceptance of humanoid robots as human partners in motor interaction. *AI & society* 26, 4 (2011), 339–345. <https://doi.org/10.1007/s00146-010-0314-2>
- [24] Hee Rin Lee and Selma Šabanović. 2014. Culturally variable preferences for robot design and use in South Korea, Turkey, and the United States. In *2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 17–24. <https://doi.org/10.1145/2559636.2559676>
- [25] Karl F. MacDorman. 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.
- [26] Martina Mara and Markus Appel. 2015. Effects of lateral head tilt on user perceptions of humanoid and android robots. *Computers in Human Behavior* 44 (2015), 326–334. <https://doi.org/10.1016/j.chb.2014.09.025>
- [27] Molly C. Martini, Christian A. Gonzalez, and Eva Wiese. 2016. Seeing minds in others—Can agents with robotic appearance have human-like preferences? *PloS one* 11, 1 (2016). <https://doi.org/10.1371/journal.pone.0146310>
- [28] Maya B. Mathur and David B. Reichling. 2016. Navigating a social world with robot partners: A quantitative cartography of the Uncanny Valley. *Cognition* 146 (2016), 22–32. <https://doi.org/10.1016/j.cognition.2015.09.008>
- [29] Takashi Minato, Michihiro Shimada, Hiroshi Ishiguro, and Shoji Itakura. 2004. Development of an android robot for studying human-robot interaction. In *International conference on Industrial, engineering and other applications of applied intelligent systems*. Springer, 424–434. https://doi.org/10.1007/978-3-540-24677-0_44
- [30] Masahiro Mori, Karl F. MacDorman, and Norri Kageki. 2012. The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine* 19, 2 (2012), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>
- [31] Omar Mubin, Catherine J. Stevens, Suleman Shahid, Abdullah Al Mahmud, and Jian-jie Dong. 2013. A review of the applicability of robots in education. *Journal of Technology in Education and Learning* 1, 209–0015 (2013), 13. <https://doi.org/10.2316/Journal.209.2013.1.209-0015>
- [32] Janna Otterbacher and Michael Talias. 2017. S/he's too Warm/Agentic! The Influence of Gender on Uncanny Reactions to Robots. In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 214–223. <https://doi.org/10.1145/2909824.3020220>
- [33] Maike Paetzel, Christopher Peters, Ingela Nyström, and Ginevra Castellano. 2016. Congruency matters—How ambiguous gender cues increase a robot's uncanniness. In *International Conference on Social Robotics*. Springer, 402–412. https://doi.org/10.1007/978-3-319-47437-3_39
- [34] Caroline Pantofaru, Leila Takayama, Tully Foote, and Bianca Soto. 2012. Exploring the role of robots in home organization. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. 327–334. <https://doi.org/10.1145/2157689.2157805>
- [35] Nuno Piçarra and Jean-Christophe Giger. 2018. Predicting intention to work with social robots at anticipation stage: Assessing the role of behavioral desire and anticipated emotions. *Computers in Human Behavior* 86 (2018), 129–146. <https://doi.org/10.1016/j.chb.2018.04.026>
- [36] Amanda Purington, Jessie G. Taft, Shruti Sannon, Natalya N. Bazarova, and Samuel Hardman Taylor. 2017. "Alexa is my new BFF" Social Roles, User Satisfaction, and Personification of the Amazon Echo. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 2853–2859. <https://doi.org/10.1145/3027063.3053246>
- [37] Laurel D. Riek, Tal-Chen Rabinowitch, Bhismadev Chakrabarti, and Peter Robinson. 2009. How anthropomorphism affects empathy toward robots. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*. 245–246. <https://doi.org/10.1145/1514095.1514158>
- [38] Eduardo Benitez Sandoval, Omar Mubin, and Mohammad Obaid. 2014. Human robot interaction and fiction: A contradiction. In *International Conference on Social Robotics*. Springer, 54–63. https://doi.org/10.1007/978-3-319-11973-1_6
- [39] Michael Schmitz. 2010. Concepts for life-like interactive objects. In *Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction*. 157–164. <https://doi.org/10.1145/1935701.1935732>
- [40] Masahiro Shiomi, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2006. Interactive humanoid robots for a science museum. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*. 305–312. <https://doi.org/10.1145/1121241.1121293>
- [41] IEEE Spectrum. 2018. All Robots. Retrieved May 25, 2020 from <https://robots.ieee.org/robots/>
- [42] Megan K. Strait, Cynthia Aguillon, Virginia Contreras, and Noemi Garcia. 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)*. IEEE, 1418–1423. <https://doi.org/10.1109/ROMAN.2017.8172490>
- [43] Megan K. Strait, Lara Vujovic, Victoria Floerke, Matthias Scheutz, and Heather Urry. 2015. Too much humanness for human-robot interaction: exposure to highly

- humanlike robots elicits aversive responding in observers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 3593–3602. <https://doi.org/10.1145/2702123.2702415>
- [44] Steven J. Stroessner and Jonathan Benitez. 2019. The social perception of humanoid and non-humanoid robots: Effects of gendered and machinelike features. *International Journal of Social Robotics* 11, 2 (2019), 305–315. <https://doi.org/10.1007/s12369-018-0502-7>
- [45] Dag Sverre Syrdal, Kerstin Dautenhahn, Michael L. Walters, and Kheng Lee Koay. 2008. Sharing Spaces with Robots in a Home Scenario-Anthropomorphic Attributions and their Effect on Proxemic Expectations and Evaluations in a Live HRI Trial. In *AAAI fall symposium: AI in Eldercare: new solutions to old problems*. 116–123.
- [46] Angela Tinwell, Deborah Abdel Nabi, and John P. Charlton. 2013. Perception of psychopathy and the Uncanny Valley in virtual characters. *Computers in Human Behavior* 29, 4 (2013), 1617–1625. <https://doi.org/10.1016/j.chb.2013.01.008>
- [47] Gabriele Trovato and Friederike Eyssel. 2017. Mind attribution to Androids: A comparative study with Italian and Japanese adolescents. In *2017 26th IEEE international symposium on robot and human interactive communication (RO-MAN)*. IEEE, 561–566. <https://doi.org/10.1109/ROMAN.2017.8172358>
- [48] Kumar Yogeeswaran, Jakub Zlotowski, Megan Livingstone, Christoph Bartneck, Hidenobu Sumioka, and Hiroshi Ishiguro. 2016. The interactive effects of robot anthropomorphism and robot ability on perceived threat and support for robotics research. *Journal of Human-Robot Interaction* 5, 2 (2016), 29–47. <https://doi.org/10.5898/JHRI.5.2.Yogeeswaran>
- [49] Jakub Zlotowski, Hidenobu Sumioka, Shuichi Nishio, Dylan F. Glas, Christoph Bartneck, and Hiroshi Ishiguro. 2016. Appearance of a robot affects the impact of its behaviour on perceived trustworthiness and empathy. *Paladyn, Journal of Behavioral Robotics* 1, open-issue (2016). <https://doi.org/10.1515/pjbr-2016-0005>
- [50] Jakub Zlotowski, Kumar Yogeeswaran, and Christoph Bartneck. 2017. Can we control it? Autonomous robots threaten human identity, uniqueness, safety, and resources. *International Journal of Human-Computer Studies* 100 (2017), 48–54. <https://doi.org/10.1016/j.ijhcs.2016.12.008>