Social Decision Making with Humans and Robots: Trust and Approachability Mediate Economic Decision Making

Stephanie Tulk & Eva Wiese
George Mason University

As humanoid robots become more advanced and commonplace, the average user may find themselves in the position of wondering if their robotic companion truly possesses a mind. It is important for scientists and designers to consider how this will likely affect our interactions with social robots. The current paper explores how social decision making with humanoid robots changes as the degree of their human-likeness changes. For that purpose, we created a spectrum of human-like agents via morphing that ranged from very robot-like to very human-like in physical appearance (in increments of 20%) and measured how this change in physical humanness affected decision-making in two economic games: Ultimatum Game (Experiment 1) and Trust Game (Experiment 2). We expected increases in human-like appearance to lead to a higher rate of punishment for unfair offers in the Ultimatum Game, and to a higher rate of trust in the Trust Game. While physical humanness did not have an impact on economic decisions in either of the experiments, follow-up analyses showed that both subjective ratings of trust and agent approachability mediated the effect of agent appearance on decision-making in both experiments. Possible consequences of these findings for human-robot interactions are discussed.

INTRODUCTION

While some still believe that a world filled with robots is still a distant science fiction dream, social robots are already becoming incorporated in our society (Fortunati et al., 2015). Aside from entertainment purposes, social robots are desirable for their ability to assist clinicians in providing specialized care (Rabbitt et al., 2015; Costa et al., 2015), as well as constant care within individual users' homes (Zaga et al., 2017). Social robots could also be used as teachers and personal assistants at home and work, or public assistants in environments such as stores or airports. However, not all people view these potential expansions positively (Flandorfer, 2012; Bartneck & Reichenbach, 2005; Scopelliti et al., 2005). It is possible that some users are not ready to accept social robots as proxies for human interaction partners in certain domains. However, just as humans have developed relationships with computers and smart phones which have become commonplace in most aspects of life, it is possible that humans will adjust to our new counterparts as they become further embedded into our world. The way we adjust to social robots, and the relationships we develop with them will likely be affected by a robot's perceived capacity as a social agent and how likely it is that people perceive them as being capable of having a mind.

Mind perception and Rationality

What does it mean to possess a mind? To reach a verdict of whether or not we believe an entity possesses a mind, there are two dimensions these entities are judged on: experience (awareness and the capacity to feel), and agency (the capacity to think and choose for oneself) (Gray et al., 2007). Where adult humans are seen as possessing high experience and agency, suggesting that they are naturally perceived as possessing a mind, typical robots are seen as having some agency and low experience. One way for non-humans to gain perceived experience would be to design them to emulate human-likeness (in appearance or behavior) in order to increase the likelihood that people attribute human-like characteristics and intentions to these non-human agents through anthropomorphism, and are therefore more likely to perceive them as rational social agents who are capable of having and carrying

out intentions as a human can. It has been shown that perception of a robot as a social agent is facilitated by humanlike appearance (Kiesler et al., 2008; Looser and Wheatley, 2010; Admoni et al., 2011; Martini, Gonzalez & Wiese, 2016), and/or behavior (Morewedge, 2009; Waytz et al., 2010; Short, Hart, Vu & Scassellati, 2010; Wiese et al., 2014); see Wiese et al. (2017); for a review. Importantly, perceiving a social agent as humanlike can cause them to be seen as entities that deserve to be treated respectfully and fairly (Epley et al., 2007), therefore perception of human-likeness may also affect social decision making with social robots.

Economic Games and Decision-Making

Behavioral Economics has adapted iconic 2-person interaction models from Game Theory (frequently called "economic games") in order to study social decision making in the presence of external motivation. A typical implementation leads individuals to try to maximize personal gains while coordinating with another player. In the Ultimatum Game (UG), Player 1 first decides how to split up some amount of money with a partner and making an offer to Player 2. Player 2 then decides to either accept the offer (each player ends up with the split Player 1 proposed), or reject the offer (each player ends up with nothing). A truly "rational" decision would be for player 2 to always accept any offer above \$0, since any amount above \$0 is still a gain. However, people will often reject offers they perceive as unfair as a means of punishing Player 1 for this unfairness (Luce & Raiffa, 2012). A neuroimaging study (Sanfey et al., 2003) had participants play the UG where they responded to offers made by players who they believed to be either a human or a computer. Unfair offers from the computer were rejected less than unfair offers from the human partners, and unfair offers from the human partners resulted in stronger activation in the bilateral anterior insula, a region associated with the feeling of disgust, than did the unfair offers from the computer. This suggests that a violation in fairness from a non-human results in less negative emotions as well as less impulse to punish this unfairness. It is possible that since the computer is not seen as having a mind, unfair offers from the computer occur randomly without any inten-

tion to present something unfair, which doesn't produce strong emotions. A follow up study incorporated facial expressions of a social avatar, which was assumed to be controlled either by a human or a computer, in order to manipulate the belief of emotions and intent with regard to fairness in economic games (DeMelo, 2015). While people responded in the same way to emotional displays from humans and computers, humanness has a stronger impact on social perceptions and decision-making. That raises the possibility that robots are treated as human like social agents, but ones whose feelings are less compelling than those of other humans.

When engaging in a social decision making task with a robot, it is possible that people adopt the same trust model as with human-human interactions, but that different qualities are used to judge trustworthiness or fairness in robots than those used to judge humans. Given this uncertainty, when playing the Trust Game with robots, some chose to interact with robots in order to better understand how to treat them (Mota et al., 2016). This suggests that while the human trust model may be adopted as a default, people might seek information about robots in order to update their model. It is important for researchers and designers to investigate how users form their human-robot trust models and how this impacts social perception and treatment of robotic agents.

This study attempts to understand how the degree of humanlike appearance affects social perception of agents and strategies for interacting with agents in social decision tasks, specifically repeated Ultimatum Games and Trust Games.

METHODS & MATERIALS

Experiment 1: Ultimatum Game (UG)

Participants

Participants were recruited and participated via Amazon's Mechanical Turk (MTurk) website and were paid \$0.30 for approximately 30 minutes of participation. The only qualification for participation was that participants be at least 18 years old and currently residing in the United States. There were 52 individuals that had participated in Experiment 1 (29 women, mean age: 37 years old). 10 participants were excluded from Experiment 1 for responding with the same answer for at least 63 of 64 trials (4 practice and 60 real).

Apparatus

This experiment was conducted online through MTurk. The experiment was programmed in javascript via an open source plugin to the Qualtrics survey site, called the QRTEngine (Barnhoorn et al., 2015). The QRTEngine allowed for precise timing of stimuli presentation, RT recording and also provided a speed test for individual computers to ensure that timing was accurate and reliable.

Stimuli

The stimuli for this experiment were a series of 6 images of faces that were robotic, human, or somewhere in between. The robot face that was used came from a Meka humanoid robot head. The human face was obtained with full permission from a 35 year-old female volunteer. Morphing software was used to create the spectrum from 100% robotic (i.e., 0% human) to 100% human with 20% increments. These faces constituted 6 agents whom participants interacted with throughout the experiment.

For each trial in experiments 1 and 2, an agent would make an offer of how to divide a sum of \$10 between them-

selves and the participant. The offers made by each agent were fixed (consistent with Sanfey paper) such that each agent made an offer of \$5 five times (i.e., totally fair: "I get \$5, you get \$5"), \$3 (somewhat fair) one time, \$2 (somewhat unfair) two times and \$1 (totally unfair) two times, for a total of 10 offers from each of the 6 agents. All subjects saw a random presentation of all 60 trials (within subjects design).

Procedure

Upon navigating to our experiment from the MTurk site, an Internet connectivity test was conducted to verify that the participant's computer had adequate signal to ensure accurate results. Once this verification was complete, participants gave informed consent and completed a basic demographic questionnaire (including age, native language and education). Participants were then instructed to ensure that they were in a quiet room without distraction, and to maximize the browser window to fill their entire screen.

Participants first completed 4 practice rounds with only the 100% human and 100% robot images, each presenting the totally fair (\$5, \$5) offers twice. For the practice trials, participants first saw the expression "New Game" in the center of the screen for 2 sec, followed by the words "Player 1:" in red text (which they were informed would only be present during the practice trials) for 2 sec. Next, the image of the agent was presented in the center of the screen for 5 sec. After this, the words "Player 1 says" in red text (also only present during the practice trials) was presented above the offer that was made by the agent for that trial (presented in black text in center screen) for 4 sec. Next, the words "Player 2:" was presented in red (only for the practice blocks) for 2 sec. On the next screen, the participant saw the text "Press 'a' to accept, or press 'l' to reject"; participants would only proceed to the next screen after responding, and the Reaction Time (RT) and response were recorded. Finally, participants were shown the outcome for 1.5 sec: either the same offer shown to them previously in the trial if they chose to accept the offer, or, "I get \$0, you get \$0", if they chose to reject the offer, and the next trial would begin. Once the practice block was completed, participants had a second chance to review the instructions before beginning the experimental block. The experimental block consisted of a random presentation of all agents and offers with the same trial sequence as the practice block, except that the red text was removed in order to simplify the display once the rules had been learned. After the experimental block, participants explicitly rated each of the 6 agents (from 1 - 7) on how trustworthy they believed the agents to be (Trust) and how readily they would approach each agent (Approachability).

Figure 1. Trial Sequence in the UG: Every trial begins with the presentation of the words "New Game" for 2 sec, followed by the image of the agent for 5 sec. Next, the offer is presented as a statement coming from the agent, displayed for 4 sec. Participants then respond with "A" to accept or "L" to reject. After the response, the result (either the same proposed amounts or \$0, \$0) is displayed for 1.5 sec. This sequence is repeated 60 times over the experimental block. Presentation of agents and offers are random.

Experiment 2: Trust Game (TG)

Participants

46 individuals participated in Experiment 2 (29 women, mean age: 40 years old). As with experiment 1, participants were

recruited and participated via MTurk and were paid \$0.30 for approximately 30 minutes of participation. The only qualification for participation was to be at least 18 years old and currently residing in the United States. 8 participants were removed for responding with the default answer on more than 90% of trials.

Figure 2. Trial Sequence in the TG: Every trial begins with the presentation of the words “New Game” for 2 sec, followed by the image of the agent for 5 sec. Next, the offer is presented as a statement coming from the agent, displayed for 4 sec. Participants then see the tripled offer and indicate how much they wish to return by dragging the slide rule and pressing the space key. This sequence is repeated 60 times over the experimental block. Presentation of agents and offers are random.

Procedure

Experiment 2 had the same agents, offers and trial format except that they played the Trust Game, where instead of simply accepting or rejecting the offer that player 1 made, participants chose an amount of money to return to player 1 after that offer was tripled (as is a typical setup with the trust game). Participants can then choose to return any integer amount from \$0 – 3 x Player 1’s offer, knowing that they kept the remainder of what they did not return to player 1 (i.e., if Player 1 offered \$3, the participant can offer to return anything from \$0-\$9. If they offer \$1, they keep \$8). This response screen showed the tripled offer with the text “Please indicate how much to return:”, above a horizontal slider scale which the participant could move right or left to indicate how much of the endowment to keep and to return to player 1. The relative amounts to keep and to return were displayed above the slider and were automatically adjusted as the slider was moved left or right. Participants would press space once they indicated the desired response via the slider, and begin the next trial.

ANALYSIS

Ultimatum Game

For the decisions to accept or reject offers made by the 6 agents, two way ANOVAs were used to investigate the relationship between agent appearance, fairness of the offer and the mean number of trials that were accepted. Two way repeated measures ANOVAs were also used to determine the relationship between RT, appearance and offers. T-tests were used to compare explicit ratings of trust and approachability for each of the agents. As a follow-up analysis, a Sobel test using regression was conducted to determine if trust and approachability mediated the relationship between appearance and response or RT.

Trust Game

Two-way repeated measures ANOVAs were used to analyze the effect of appearance and fairness on the amount of the tripled offer participants decided to return to P1. For RT data, all trials where participants returned \$0 were excluded (mouse task was avoided thus the RT was much shorter), and repeated measures ANOVAs were again used to determine the impact of appearance and fairness on RT. Again, t-tests were used to compare explicit ratings of trust and approachability for each

of the agents and mediation analyses were conducted to determine if trust and approachability mediated the effect of appearance on response or RT.

RESULTS

Ultimatum Game (UG)

Data was filtered by RT, where trials that were more than 2 standard deviations above the mean RT, or faster than 100 ms were pairwise deleted. This was done to remove trials where participants may not have been paying attention, since they were at home and not in a lab setting. This brought the number of total observations from 2520 to 2406. For responses, offer significantly affected the mean number of offers that were accepted ($F(1,976) = 946.44, p < .001$) but agent appearance did not ($F(1,976) = 0.15, p = 0.979$). Appearance affected explicit ratings of trust ($t(251) = -3.7, p = .002$) and approachability ($t(251) = -2.11, p = .036$), and on further analysis, approachability significantly affected the mean acceptance of offers ($F(1,248) = 6.16, p = .014$), though trust did not ($F(1,248) = 0.23, p = .635$). While ratings of trust and approachability seem to have similar trend with visual inspection (see Figure 3), they are only marginally correlated ($t(251) = -1.76, p = .079$). A Sobel mediation test indicated that approachability mediated the effect of appearance on mean offer acceptance ($z' = 2.36, p = .018$).

RT for each response was not significantly affected by agent ($F(1,2398) = 0.76, p = .383$) or offer ($F(1,2398) = 2.25, p = .134$). Approachability had a significant effect on RT ($F(1,2398) = 7.66, p = .006$) and was determined to mediate the relationship between appearance and RT ($z' = 2.84, p = .005$).

Figure 3. (A) Average amount of offers accepted by participant over approachability ratings in the UG. A boxplot is overlaid on each rating level. Each dot represents a rating of approachability given to 1 of the 6 agents and the average % of offers from the agent that were accepted. The red line is the linear best fit line, which has a positive slope. (B) Explicit ratings of trust and approachability over the percent of human likeness for the 6 agents in the UG. While the overall trends are similar, there is only a marginal correlation between trust and approachability.

Trust Game (TG)

Data was filtered by the total overall amount of money that was returned to Player 1. Participants who chose to return \$0 on a significant amount of trials were excluded from analysis, or those whose overall returns were outside the 1st quartile. A total of 8 participants were removed, bringing the number of participants to 35. Again, RT was log transformed to normalize the data. The amount of money returned to P2 was significantly correlated with offer ($F(1,2092) = 474.81, p < .001$) but not agent appearance ($F(1,2092) = 1.42, p = .234$). Again, trust ($t(209) = 19.80, p < .001$) and approachability ($t(209) = 19.90, p < .001$) were significantly correlated with agent appearance. Trust and approachability were significantly correlated ($t(209)$

= 2.53, $p = .012$). Trust significantly impacted the amount returned to P2 ($F(1,2092) = 19.47$, $p < .001$), but approachability did not ($F(1,2092) = .04$, $p < .847$). A Sobel test did not show that trust mediated the effect of appearance on amount returned ($z' = 1.16$, $p = .245$).

Figure 4. (A) Average amount returned to Player 1 by participant over trust ratings in the TG. A boxplot is overlaid on each rating level. Each dot represents a rating of trust given to 1 of the 6 agents and the average dollar amount of the tripled offer from the agent that was returned. The red line is the linear best fit line, which has a positive slope. (B) Explicit ratings of trust and approachability over the percent of humanlikeness for the 6 agents in the TG. While the overall trends are similar, there is only a marginal correlation between trust and approachability.

Offer impacted RT ($F(1,2092) = 4.44$, $p = .035$), but agent did not ($F(1,2092) = 0.12$, $p = .734$). When all returns of \$0 were excluded (significantly faster responses since there was no need to use the mouse), offer ($F(1,1214) = 3.25$, $p = .021$), still significantly impacted RT but agent did not ($F(1,1214) = 1.51$, $p = .219$). Trust significantly impacted RT ($F(1,1214) = 4.84$, $p = .028$), but approach did not ($F(1,1214) = 4.84$, $p = .028$). There was a significant interaction between trust and approachability ($F(1,1214) = 3.90$, $p = .049$). A Sobel mediation test indicated that trust mediated the effect of appearance on RT ($z' = 2.25$, $p = 0.024$).

DISCUSSION

The goal of this study was to investigate how humanlike appearance affected people's treatment of agents in social decision making tasks. We hypothesized that increased humanlikeness would result in a greater expectation of fairness and more punishments for unfair offers (i.e., rejecting unfair offers in the UG or offering unfair returns in the TG), but the results show that the story is not so straightforward. In both games, participants made the distinction between fair and unfair offers from Player 1, but degree of human-like appearance did not uniformly impact participants' decisions or the time it took to make these decisions. Instead, after having the same exposure and overall experience with each agent in the experimental block, participants' explicit levels of trust and approachability were rated significantly differently for each agent, on average increasing with humanlikeness. In a follow up analysis, it was determined that the more approachable an agent was perceived to be, the more offers were accepted in the UG and the less time it took to respond, and the more participants trusted the agent, the higher their overall returns were in the TG, and the more time it took to respond. This suggests that treatment of agents in social decisions (either in the decision or how much time it takes to make the decision) varies by how much individuals perceive that they can trust different agents and how approachable they perceive agents to be, and humanlikeness is one factor that affects these perceptions. Demonstrable evidence of mediation was found in both games.

In this study, measurements of trust and approachability were taken after participants had experiences with each agent

to draw from. It is possible that the negative experience of being presented with unfair offers is more salient for agents who are on the perceptual threshold of being sufficiently human, and that these experiences are closely attended to in order to understand how to treat the agent in future interactions, which potentially affected explicit ratings after the experimental block. It is also possible that paying more attention to unfair offers caused participants to make judgements about these agents early on, which could have impacted their performance more strongly over the block. This could be further explored by asking for pre and post ratings of trust and approachability and determining if these ratings change after playing some economic game.

It is possibly meaningful that approachability mediated response and RT in the UG while trust mediated the RT in the TG. While the overall trends of trust and approachability for the 6 agents are similar in both games, the concept of approachability is more superficial than trust. Since making decisions in the UG (accept fair offers and reject unfair ones) is simpler than deciding a fair amount to return in the TG, it's possible that the TG required more consideration of the agent since participants were required to consider if and how fairly they should act towards the agents, possibly even considering that Player 1's perception of fairness from the participant could affect what offers they would receive in the future. In this case, feeling like the agent was trustworthy might inspire empathy that resulted in more consideration of true fairness.

Overall, it is important to consider how human-robot trust models develop as people have more experiences with robots to draw from. If users begin to have more experience to draw from, it is unclear how the resulting model will differ from the human-human trust model. If the human-robot model is generalized to all interactions a person has with robots, the overall success of social robotics as a mainstream technology might be impacted by the initial experiences users have. Furthermore, while individual humans have developed tendencies towards fairness and honesty in order to exist harmoniously within society, a robot will not necessarily have the same motivations. Since robotic interfaces are intentionally developed by humans, the designer has the power to make honest robots to benefit others, or deceptive ones that can take advantage of other people, for instance, by spreading misinformation or attempting to trick users into giving away private information. To address this concern, future work will attempt to address how well people can detect deception and other untrustworthy qualities from agents as humanlikeness varies.

REFERENCES

- Admoni, H., Bank, C., Tan, J., Toneva, M., and Scassellati, B. (2011). "Robot gaze does not reflexively cue human attention," in Proceedings of the 33rd Annual Conference of the Cognitive Science Society, eds L. Carlson, C. Hölscher, and T. Shipley (Austin, TX: Cognitive Science Society), 1983–1988.
- Barnhoorn, J. S., Haasnoot, E., Bocanegra, B. R., & van Steenbergen, H. (2015). QRTEngine: An easy solution for running online reaction time experiments using Qualtrics. *Behavior research methods*, 47(4), 918–929.
- Bartneck, C., and Reichenbach, J. (2005). Subtle emotional expressions of synthetic characters. *Int. J. Hum. Comput. Stud.* 62, 179–192. doi: 10.1016/j.ijhcs.2004.11.006
- Costa, S., Lehmann, H., Dautenhahn, K., Robins, B., & Soares, F. (2015). Using a humanoid robot to elicit body awareness and appropriate physical interaction in children with autism. *International journal of social robotics*, 7(2), 265–278.
- De Melo, C. M., Gratch, J., & Carnevale, P. J. (2014, June). The Importance of Cognition and Affect for Artificially Intelligent Decision Makers. In AAAI (pp. 336–342).
- De Melo, C. M., Gratch, J., & Carnevale, P. J. (2015). Humans versus Computers: Impact of Emotion Expressions on People's Decision Making. *IEEE Transactions on Affective Computing*, 6(2), 127–136.
- Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: a three-factor theory of anthropomorphism. *Psychological review*, 114(4), 864.

- Flandorfer, P. (2012). Population ageing and socially assistive robots for elderly persons: the importance of sociodemographic factors for user acceptance. *Int. J. Popul. Res.* 2012, 13. doi: 10.1155/2012/829835
- Fortunati, L., Esposito, A., & Lugano, G. (2015). Introduction to the special issue "Beyond industrial robotics: Social robots entering public and domestic spheres".
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, 315(5812), 619-619.
- Gallo Jr, P. S., & McClintock, C. G. (1965). Cooperative and competitive behavior in mixed-motive games. *Journal of Conflict Resolution*, 9(1), 68-78.
- Kiesler, S., Powers, A., Fussell, S. R., and Torrey, C. (2008). Anthropomorphic interactions with a robot and robot-like agent. *Soc. Cogn.* 26, 169–181. doi: 10.1521/soco.2008.26.2.169
- Looser, C. E., and Wheatley, T. (2010). The tipping point of animacy. How, when, and where we perceive life in a face. *Psychol. Sci.* 21, 1854–1862. doi: 10.1177/0956797610388044
- Luce, R. D., & Raiffa, H. (2012). *Games and decisions: Introduction and critical survey*. Courier Corporation.
- Martini, M. C., Gonzalez, C. A., & Wiese, E. (2016). Seeing Minds in Others—Can Agents with Robotic Appearance Have Human-Like Preferences?. *PloS one*, 11(1), e0146310.
- Morewedge, C. K. (2009). Negativity bias in attribution of external agency. *J. Exp. Psychol. Gen.* 138, 535–545. doi: 10.1037/a0016796
- Mota, R. C. R., Rea, D. J., Le Tran, A., Young, J. E., Sharlin, E., & Sousa, M. C. (2016, August). Playing the 'trust game' with robots: Social strategies and experiences. In *Robot and Human Interactive Communication (RO-MAN)*, 2016 25th IEEE International Symposium on (pp. 519-524). IEEE.
- Rabbitt, S. M., Kazdin, A. E., & Scassellati, B. (2015). Integrating socially assistive robotics into mental healthcare interventions: Applications and recommendations for expanded use. *Clinical psychology review*, 35, 35-46.
- Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science*, 300(5626), 1755-1758.
- Scopelliti, M., Giuliani, M. V., and Fornara, F. (2005). Robots in a domestic setting: a psychological approach. *Univers. Access Inf. Soc.* 4, 146–155. doi: 10.1007/s10209-005-0118-1
- Short, E., Hart, J., Vu, M., & Scassellati, B. (2010, March). No fair!! an interaction with a cheating robot. In *Human-Robot Interaction (HRI)*, 2010 5th ACM/IEEE International Conference on (pp. 219-226). IEEE.
- Waytz, A., Morewedge, C. K., Epley, N., Monteleone, G., Gao, J.-H., and Cacioppo, J. T. (2010). Making sense by making sentient: effectance motivation increases anthropomorphism. *J. Pers. Soc. Psychol.* 99, 410–435. doi: 10.1037/a0020240
- Wiese, E., Wykowska, A., Zwickel, J., & Müller, H. J. (2012). I see what you mean: how attentional selection is shaped by ascribing intentions to others. *PloS one*, 7(9), e45391.
- Wiese, E., Wykowska, A., and Müller, H. J. (2014). What we observe is biased by what other people tell us: beliefs about the reliability of gaze behavior modulate attentional orienting to gaze cues. *PLoS ONE* 9:4. doi: 10.1371/journal.pone.0094529
- Wiese, E., Metta, G., & Wykowska, A. (2017). Robots as intentional agents: Using neuroscientific methods to make robots appear more social. *Frontiers in psychology*, 8, 1663.
- Wykowska, A., Wiese, E., Prosser, A., & Müller, H. J. (2014). Beliefs about the minds of others influence how we process sensory information. *PLoS One*, 9(4), e94339.
- Zaga, C., de Vries, R. A., Li, J., Truong, K. P., & Evers, V. (2017, May). A Simple Nod of the Head: The Effect of Minimal Robot Movements on Children's Perception of a Low-Anthropomorphic Robot. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (pp. 336-341). ACM.