

Perceptions of Cognitive and Affective Empathetic Statements by Socially Assistive Robots

Christopher Birmingham
Computer Science Dept.
University of Southern California
Los Angeles, California
cbirming@usc.edu

Ashley Perez
Computer Science Dept.
University of Southern California
Los Angeles, California
asperez@usc.edu

Maja Matarić
Computer Science Dept.
University of Southern California
Los Angeles, California
mataric@usc.edu

Abstract—Communicating empathy is important for building relationships in numerous contexts. Consequently, the failure of robots to be perceived as empathetic by human users could be detrimental to developing effective human-robot interaction. Work on computational models of empathy has been growing rapidly, reflecting the importance of this ability for machines. Despite growing recent work, there remain unanswered questions about how users perceive different forms of empathetic expression by robots and how attitudes towards robots may mediate perceptions of robot empathy. Do people really believe that robots can feel or understand emotions? This work studied the difference in viewers’ perceptions of cognitive and affective empathetic statements made by a robot in response to human disclosure. In a within-subjects study, participants (n=111) watched videos in which a human disclosed negative emotions around COVID-19, and a robot responded with either affective or cognitive empathetic responses. Using an adapted version of the Robot’s Perceived Empathy (RoPE) scale, participants rated their perceptions of the robot’s empathy in both cases. We found that participants perceived the robot that made affective empathetic statements as being more empathetic than the robot that made cognitive empathetic statements; we also found that participants with more negative attitudes toward robots were more likely to rate the cognitive condition as more empathetic than the affective condition. These results inform HRI in general and future work into developing robots that will be perceived as empathetic and could personalize empathetic responses to each user.

Index Terms—Human-Robot Interaction, Empathy, Socially Assistive Robotics

I. INTRODUCTION

Empathy is an individual’s capacity to respond to the emotions of others [1], an integral part of human social communication. We use empathy to build connections to one another by understanding and responding to one another’s thoughts and feelings. As robots and conversational agents begin to have complex interactions with humans, they will also need the capacity to understand and respond with empathy to the thoughts and feelings of humans. Computational empathy [2] can support the development of social relationships between humans and robots, essential for socially assistive robots and social robots, and helpful for most human-robot interaction. As with human empathy, computational empathy must be communicated clearly and perceived as empathetic in order to be effective.

While the form empathy takes can vary greatly depending on the context, empathy can generally be understood through two broad categories or dimensions: cognitive empathy and affective empathy. *Cognitive empathy* is the ability to understand another person’s thoughts and feelings. *Affective empathy* is the ability to have an affective response that is more appropriate to someone else’s situation than one’s own [3]. Humans naturally employ cognitive, affective empathy or both forms of empathy in various social interactions, but it is as yet unclear how a social or socially assistive robot should choose between these types of empathy when conveying an empathetic response.

A robot’s ability to *understand* another’s feelings (cognitive empathy) and to *feel* in response (affective empathy) is dramatically different from human abilities. This may lead the human target of a robot’s empathy to perceive the robot’s attempt at empathy as insincere, causing unintended consequences. Although there is no prior work exploring *perceptions* of cognitive and affective empathy, we hypothesized that people are more likely to believe that a robot is able to identify and understand an emotion rather than genuinely feel an emotion and would therefore rate a robot using cognitive empathy as more believable and empathetic than a robot using affective empathy.

In this work we explored the perceptions of a robot’s cognitive and affective empathetic responses to a human’s disclosure through a within-subject study with counterbalancing (n=111). We asked participants to rate their perceptions of a robot’s empathy based on short videos of the robot giving cognitive and affective empathetic responses to disclosures about stresses related to the COVID-19 pandemic made by an actor. To rate the interaction, participants filled out a short survey after watching two videos demonstrating each condition. We analyzed the relationship between the participants’ attitudes towards robots and their rating of how genuine they felt the interaction was and their rating of the robot’s empathy in each condition.

Our surprising results showed that the robot that responded with affective empathy was perceived as more empathetic. Further, we found that participants’ attitudes towards robots were correlated with their ratings of the robot’s empathy and their belief that the interaction between the robot and the

actors were believable, natural, and genuine, and that negative attitudes towards robots were associated with higher ratings of cognitive empathy over affective empathy. These results serve to inform work into human-centered HRI and into developing robots that will be perceived to be appropriately empathetic and could personalize their empathetic responses to each user.

This paper is structured as follows. Section II presents a summary of commonly used measures of empathy in both humans and robots and prior work studying empathy in HRI. Section III provides a detailed description of our online, video-based study including the all of the details needed to replicate our work. Section IV includes all findings pertaining to the different perceptions of cognitive and affective empathy, as well as the relationship between attitudes towards robots, empathy, and belief. Section V discusses the significance of these findings along with our planned future work. Finally, Section VI concludes the paper.

II. RELATED WORK

A. Measuring Empathy

Empathy involves an individual experiencing a change in their own thoughts or feelings in response to observations of another individual. Measuring this internal phenomenon has taken many forms, including the first person perspective (the trait of empathy), the second person perspective (the recipient's rating), and the third person perspective (expert rating of empathetic actions) [4]. Measuring the trait of empathy has evolved. In the past, it was done by separately measuring the two dimensions of empathy, cognitive empathy through the use of the Hogan Empathy Scale [5] and affective empathy through the Interpersonal Reactivity Index [6]). Subsequently, increasingly comprehensive measures have been applied, such as the Toronto Empathy Scale [7] and the Basic Empathy Scale [8]. More recent work, such as the Empathy Assessment Index [9], has introduced subscales for measuring tendencies towards empathic action in order to incorporate a pro-social dimension into the measurement of empathy.

Measuring perceptions of empathy is used in clinical practice, typically through patient ratings of the empathy of their doctor or therapist, such as the Consultation and Relational Empathy (CARE) [10] and the Jefferson Scale of Patient Perceptions of Physician Empathy (JSPPE) [11]. Most common in healthcare settings are third person scales completed by an observer of the interaction between two individuals, typically a healthcare provider and a patient. Examples include the Empathic Communication Coding System (ECCS) [12] and the Global Rating Scale [13].

Prior work on measuring empathy or perceived empathy of robots is limited. One approach to measuring the success of an empathetic robot is the use of the McGill Friendship Questionnaire along with the CARE measure [14], [15]. More recently, Charrier *et al.* [16], [17] introduced the Robot's Perceived Empathy (RoPE) scale to provide a means of directly measuring the perceived empathy of a robot. This scale provides a measure of the second person perspective of the robot. The RoPE scale has been used to measure

the perceived empathy of an empathetic chatbot response for medical assistance [18] and of a virtual robot used for genetic counseling [19].

Our work measured the perceived empathy of a robot shown in a video interaction. Because we used a scripted interaction, instead of applying traditional third person measures adapted for robot interaction, we modified the RoPE scale for use in the third person perspective.

B. Empathy in HRI

Two distinct research areas address empathy in social agents [20]. In the first, the human interaction partner is the observer of the robot and the robot is the target of the human's empathy. In the second, the robot is the observer of the human and is displaying empathy to the human. There has been significant progress on both areas in recent years [21]; in this work, we focus on the second area. The first model of empathy in socially assistive robotics was proposed by Tapus and Matarić [22] and included the following features: recognizing, understanding and interpreting other's emotional state, processing and expressing emotions across modalities, communicating with others, and perspective taking.

Empathy in HRI is particularly challenging because accurately perceiving a user's affective states in an open domain or environment has not been technologically possible until recently. Early studies in empathy used mimicry to approximate empathy, either by mirroring the emotion inferred through speech intonation [23] or through copying the user's mouth and head movements [24]. Other work has explored how empathy during game play affects children, such as in two studies where children played chess against an iCat robot [25], [26]. In one study the robot utilized an SVM-based Affect Recognitions system to return probabilities of the children's valence and displayed social supportive behaviors if the child's affective state was negative [25]. In the other, the robot analyzed the game from the child's position to determine how they might be feeling, and reacted according to the assumed feelings of the child. Cramer *et al.* [26] examined how a robot's empathy could affect people's attitudes towards the robot, including exploring how the accuracy of the robot's empathy (in this case when an iCat robot was playing a cooperative game with an actor) affected attitudes towards the robot. Results showed that inaccurate empathy led to a significant decrease in user trust in the robot.

Our work builds on Cramer *et al.* [26] by examining how different types of empathy (affective and cognitive) affect user perceptions of the overall empathy of the robot, and how attitudes toward robots affect that perception.

III. METHODS

A. Hypotheses

H₁: Participants will rate the empathy of a robot that makes cognitive-based empathy statements higher than a robot that makes affective-based empathy statements.

H₂: More negative attitudes towards robots will be correlated with higher ratings in the cognitive condition compared

to the affective condition for both:

H_{2A} : empathy, and

H_{2B} : belief in the interaction.

B. Study Design

We utilized a within-subject design with counterbalancing. The same participants evaluated both the cognitive and affective conditions in a random order. This introduced less variation in participant scores and required fewer participants than a between-subjects design.

Prior to running the study we chose a desired significance level of **0.05**, power of **0.8**, and a medium effect size: Cohen’s $d=.4$. To determine an appropriate sample size we needed an estimate of the standard deviation, which we obtained by running a pilot study on Amazon Mechanical Turk (AMT) with 15 participants under conditions identical to the main study. We found the standard deviation in scored empathy in the pilot study to be approximately **0.7**. To calculate the necessary sample size, we utilized Kadam *et al.* [27]:

$$n = 2(Z_{\alpha} + Z_{1-\beta})^2 * \sigma^2 / \delta^2$$

where Z_{α} is a constant based on a 5% α error in a two-sided effect, $Z_{1-\beta}$ is a constant set according to 80% power, σ is the standard deviation and δ is the estimated effect size. From this, the minimum study population required to test H_1 was determined to be 48.

For evaluating H_2 , we utilized the commonly used heuristic for testing individual predictors proposed in Green [28]: $N > 104 + m$ where m is the number of independent variables. From this we concluded that the minimum study population should be increased to 106.

C. Participants

Participants were recruited through AMT. Inclusion criteria were: being from the US, having an approval rate greater than 98%, having at least 1,000 prior approved tasks, and not having taken this task before. Exclusion criteria were: age under 18, not fluent in English, and being visually impaired (because of the visual nature of the study stimuli). In total, 181 participants accepted the task and 111 completed the full survey. Only those who completed the full survey were included in our analysis.

Participants who completed the full survey identified as: 75 Male, 33 Female, 3 Prefer not to specify; 76 White, 15 Asian, 14 Hispanic, 8 African American, 2 American Indian or Alaskan Native, and 1 Native Hawaiian or Other Pacific Islander; the ages ranged from 20 to 64, with a mean of 36; the highest level of education completed was: 12 High School Equivalent or less, 26 Associate’s or some college, 56 Bachelor, 14 Master’s, 3 Professional degree. All participants consented to be a part of the study and were paid for their time.

Although the cognitive and affective conditions were balanced due to the nature of the within-subject design, the condition order assignment was random and therefore was not balanced by participant gender, age, or education.



(a) Camera views of the interaction with the female actor.



(b) Camera views of the interaction with the male actor.

Fig. 1: Camera views used to record the interaction: the view capturing the front of the actor (left) and the view capturing the front of the robot.

Negative Emotion (X)	Difficulty Related to the Pandemic (Y)
Depressed	I don’t have the energy to do anything anymore.
Overwhelmed	I haven’t been able to go outside and do anything in so long.

TABLE I: Actor disclosures were structured in the format of: “I feel X because Y”

D. Recording Setup

The recording space was chosen and arranged to give the impression of a therapy session between the robot and the actor. The robot was a tabletop humanoid QTRobot [29], set up across the table from the actor. The recording equipment was structured to capture an interview between the robot and the actor. One camera recorded the front of the robot over the shoulder of the actor, and another recorded the front of the actor over the shoulder of the robot, as shown in Figure 1. A microphone was placed between the robot and the actor to capture the speech of both. The robot was controlled with the open-source HARMONI software [30] that utilized CoRDial [31] for generating the robot’s facial expressions and Google text-to-speech [32] for generating the robot’s speech. All software necessary for reproducing this study can be found at: <https://github.com/interaction-lab/HARMONI/releases/tag/Empathy-Study>.

The dialog in each video consisted of three parts. As the actor sat down, the robot greeted them and asked, “How are you doing today?” The actor then made a disclosure related to the COVID-19 pandemic. The disclosures were structured as follows: “To be honest, I am not doing so well, I feel X because Y” where X is a negative emotion and Y is a difficulty

Belief In	Statement
Robot Response	When the robot responded to Brandon and Isabel the response was believable.
Natural Interaction	The interaction between Brandon and Isabel and the robot appears natural.
Actor Distress	Brandon and Isabel appear genuinely distressed.

TABLE II: The participant belief and the corresponding statement used to gauge participant belief in the interaction. The participants rated their agreement with each statement on a five-point Likert scale.

related to the pandemic. The robot then made an empathetic response that consisted of three parts: 1) an acknowledgement of the actor sharing information: “Thank you for sharing that”; 2) an affective or cognitive empathetic statement (depending on the condition, described below); and 3) an invitation to share more information: “Would you like to expand on that? We can talk more about why you are feeling this way.”

The affective and cognitive empathetic statements were designed to convey the difference between the two types of empathy while maintaining a balanced length and structure of the utterance. In the cognitive empathy statement, the robot communicated that it understood the actor by saying: “It sounds like you are having a hard time feeling **X** right now”, where **X** is the emotion that the actor stated they were feeling. In the affective empathy statement, the robot communicated that it related to the emotion identified by the actor by saying: “I can totally relate to what you are feeling, I also feel **X** sometimes”, where **X** is the emotion that the actor stated they were feeling.

In addition to speaking, the robot used a set of three gestures. The first gesture was a two-arm wave; the robot used it while it was greeting the actor, to express welcoming and excitement. The second gesture involved the robot bringing its right hand to its chest in a gesture of sincerity and thanks; the robot used it when it said “Thank you for sharing” and also when it invited the participant to talk more. In the third gesture, the robot brought its hands to its hips to emphasize the point it was making whenever it made the affective or cognitive empathetic statements.

In both study conditions (affective and cognitive empathy statements) the robot maintained a neutral expression and tone of voice, and used the same gestures and the same turn-taking timing (when it started to speak after the actor stopped speaking). The video shown to the study participants did not include a response by the actors to avoid biasing the participants. A video of the full interaction is available here: <https://youtu.be/nUf3SDUtBZ8>.

E. Measures

We collected quantitative and qualitative data to measure the participants’ perception of the robot’s empathy and explore what factors that affected that perception.

1) *Quantitative Data*: Participants’ attitudes toward robots were measured with the Negative Attitudes Towards Robots

Scale (NARS) [33] that was shortened to include only questions relevant to socially assistive robots. The participants’ perceptions of the robot’s empathy to the actor were measured with the Robot’s Perceived Empathy (RoPE) scale, modified from a first person questionnaire to a third person questionnaire to account for the fact that participants were watching an actor instead of interacting themselves with the robot. Specifically, we replaced the first person pronouns “I” and “me” with the third person reference to the actors’ names, and modified the questions to ask if the robot “appears to” meet the criteria of each question. Three additional questions were included to assess the participants’ belief in the interaction scenario, as shown in Table II. All three used a five-point Likert scale ranging from ‘Strongly Disagree’ to ‘Strongly Agree’.

2) *Qualitative Data*: Participants were asked if they would like to interact with the robot and to explain why or why not. Participants were also asked to describe the differences between the robot responses in the two conditions in their own words and explain which felt more empathetic. The participants provided answers to these write-in questions in a blank text box.

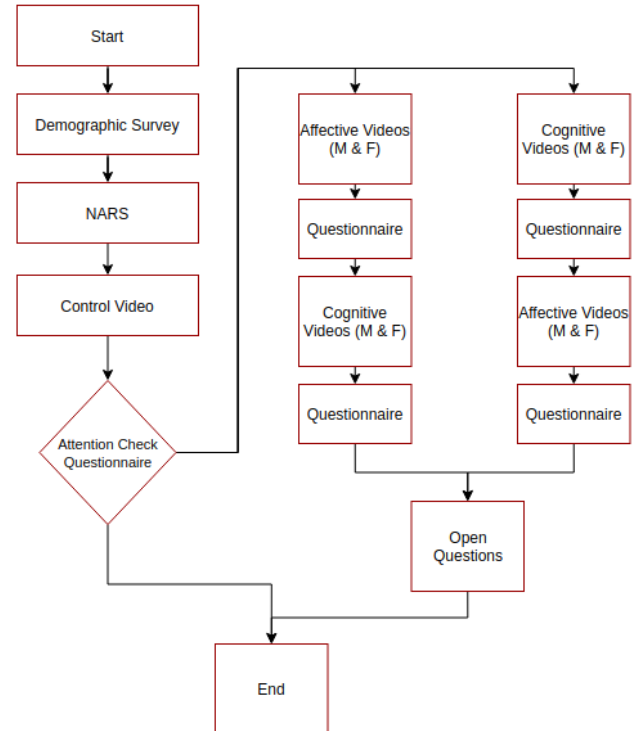


Fig. 2: Flow chart of steps participants completed during the study. Participants were considered to have completed the full study if they correctly answered all questions in the Attention Check Questionnaire.

F. Procedure

This study was approved by the University’s Institutional Review Board (IRB). Participants were recruited through AMT. As seen in the flowchart in Figure 2, participants first completed a demographic questionnaire and the NARS. They then watched a video of the short interaction between the robot and an actor and completed an attention check questionnaire that asked three simple multiple choice questions about the video they had just watched. If they answered any of the questions incorrectly, they were directed to the end of the survey. If they answered the questions correctly, they were randomly assigned to a condition and either watched the affective or cognitive condition videos first. Each condition involved watching two videos of actors interacting with the robot, one male actor and one female actor. After they completed both videos, they were directed to complete a short questionnaire about the robot they had just watched, including an attention check question, the modified RoPE scale questions, the three belief questions, and the short-form question about whether they would like to interact with the robot. Participants then completed the other condition of the study. After completing both conditions, participants were asked to describe the differences between the robots and explain which robot was more empathetic.

Participants were paid proportionally to the time spent on the task. Those who failed to complete the first attention check spent about 4 minutes on the study and were paid US\$1 for their time. Those who completed the full survey, which took approximately 24 minutes, were paid US\$6.

G. Analyses

This study measured participants’ attitudes toward the two types of robot-expressed empathy. The dependent variables were **empathy** (the RoPE scale score) and **belief** (how the participants rated their belief in the interaction scenario, shown in Table II). The independent variables were the condition viewed (affective or cognitive), participants’ attitudes towards robots (the NARS score), participants’ demographics, and the ordering of conditions.

The five-point Likert scale questions were coded from -2 to 2 , and the net score was calculated as the average of the individual question scores, with the appropriate questions reverse-coded. The dependent variables were tested with the Shapiro-Wilk test; they were found to not be normally distributed, so non-parametric tests were used. To test the difference between the affective and cognitive conditions, the Wilcoxon signed-rank test was used. To establish the relationship between the dependent and independent variables, a two-sided linear least-squares regression was used, with the correlation and significance determined with the Spearman’s rank correlation.

IV. RESULTS

A. Cognitive vs. Affective Conditions

1) *Empathy*: The perceived empathy was higher in the affective condition ($M = 0.62$, $SD = 0.67$) than in




Fig. 3: Differences between the cognitive and affective conditions in the rated empathy and belief. Significant differences are indicated with asterisks: $*=p < .05$, $****=p < .0001$.

Belief In	Affective Condition	Cognitive Condition
Robot Response	0.76	0.60
Natural Interaction	0.54	0.36
Actor Distress	0.78	0.72

TABLE III: Mean rating of agreement with each belief statement for each condition.

the cognitive condition ($M = 0.38$, $SD = 0.71$), as can be seen in Figure 3. The affective condition was rated significantly higher than the cognitive condition ($p < .001$), showing a small-to-medium effect (Cohen’s $d = 0.35$, effect-size $r = 0.17$). Thus, H_1 is not supported.

2) *Belief*: The perceived net belief was higher in the affective condition ($M = 0.69$, $SD = 0.82$) than in the cognitive condition ($M = 0.56$, $SD = 0.82$), as shown in Figure 3. The affective condition was rated significantly higher than the cognitive condition ($p < .05$), showing a small effect (Cohen’s $d=0.16$, effect-size $r=0.08$). Table III provides the data demonstrating that the affective condition was rated as more believable on all questions. The largest difference was found in the question asking if the interaction appears natural, where the mean participant rating for the affective condition was 0.18 greater than the mean for the cognitive condition.

3) *Interest In Future Interaction*: After both conditions, participants were asked if they would be interested in interacting with the robot in the future. In the affective condition, 47% of participants said they would like to interact with a robot like the one shown in the video, while 32% said they would not, and 22% were unsure. In the cognitive condition, 38% said they would be like to, 36% said they would not, and 26% were unsure. We found no statistical test for significance of these findings given their nominal scale (yes-no-maybe). While this is not a direct test of H_1 , it provides further evidence that H_1 is not supported.

B. Attitude Towards Robots vs. Empathy and Belief

As can be seen in Figure 4, the NARS scores were negatively correlated with empathy (RoPEs) and belief measures,




Fig. 4: The dependent measures plotted against the NARS scores, including the linear regression line and the shaded 95% confidence interval. Empathy is plotted in the top row, belief in the bottom. Ratings from the cognitive condition are on the left, the affective condition in the center, and the differences between the two are plotted on the right.

although they were positively correlated with the difference between the cognitive and affective conditions for both measures. In the affective condition, the correlation between NARS and empathy ($r(111)=-.23$, $p = .016$) and between NARS and belief measures ($r(111)=-.30$, $p = 0.002$) was significant. In the cognitive condition, it was significantly correlated with the belief measure ($r(111)=-.20$, $p = 0.037$) but not with the empathy measure ($r(111)=-.11$, $p = 0.244$). For the difference between the two conditions (cognitive and affective) the correlation with the empathy measure was significant ($r(111)=.23$, $p = 0.015$) but the belief measure was not ($r(111)=.11$, $p = 0.251$). Therefore, H_{2A} is supported and H_{2B} is not supported.

C. Belief vs. Empathy



Fig. 5: Correlation between rated belief in the interaction and perceived empathy of the robot, including the linear regression line and the shaded 95% confidence interval.

As can be seen in Figure 5, there was a positive correlation between the participants' rating of belief in the interaction

and their ratings of the robot's empathy. This correlation was significant for both the cognitive ($r(111)=.68$, $p < 0.001$) and affective ($r(111)=.67$, $p < 0.001$) conditions.

D. Order Effects




Fig. 6: Difference in participant ratings of dependent measures between the cognitive first and affective first ordering. Significant differences are indicated with asterisks: ns= $p > .05$, $*$ = $p < .05$, $**$ = $p < .01$.

There were no significant differences in participant ratings of the robot's empathy in either the affective first or the cognitive first condition, as shown in Figure 6. Participant ratings of their belief in the interaction were significantly higher in the cognitive first order for both the cognitive ($p = 0.033$) and affective conditions ($p = 0.002$).

V. DISCUSSION

This work examined participant ratings of robot empathy after watching a socially assistive robot deliver two types of empathetic statements (cognitive and affective) to an actor. We found that the participants rated the affective statements higher than cognitive ones. The participants' attitude toward the robot was correlated with their rating of the robot's empathy and with the difference between the two conditions.

A. Participants rated affective empathy higher

We expected participants to find the idea that a robot could feel an emotion to be less believable than the idea that a robot could understand the feelings of a human, which we hypothesized (H_1) would lead participants to rate the cognitive condition higher than the affective condition. Contrary to this expectation, we found that participants rated the robot that made affective empathetic statements as significantly more empathetic than the robot that made cognitive empathetic statements across all measures. Participants also expressed more interest in interacting with the robot seen in the affective condition. These results contradict H_1 .

One possible explanation for these surprising results can be found in the breakdown of our expectations. As shown in Figure 5, the expectation that belief was strongly correlated with the perceived empathy is correct; however, the expectation that the belief in the cognitive condition would be higher was incorrect, as shown in Figure 3. Although this may explain why H_1 was not supported, it does not explain why the affective empathy condition was rated higher.

The higher rating of the affective condition may be an indication that affective empathy (expressing feeling similarly) is generally perceived as more powerful than cognitive empathy (expressing understanding); however, we found no evidence for this in the literature on empathy and therefore this requires further investigation. Alternatively, the fact that the affective condition was rated higher may be explained by the correlation between the negative attitudes toward robots and the difference between the ratings of empathy in the cognitive and affective condition. Participants in our study generally scored low on the NARS, which correlated with higher ratings of empathy in the affective condition compared to the cognitive condition.

Participant interest in a future interaction with the robot matched their ratings of empathy and belief in the two conditions. In the affective condition, almost half of the participants said they would be interested in interacting with the robot shown in the video, while another third of participants were not sure. This is higher than in the cognitive condition, where only 38% of participants were interested in interacting with the robot and about a quarter described themselves as unsure.

B. More negative attitudes toward robots were correlated with empathy and belief ratings

As presented in Section IV-B, participants' prior attitude toward robots was correlated with the ratings of belief and empathy in the affective condition and with belief in the cognitive condition. This correlation shows how a more negative attitude

toward robots predisposes participants to be less likely to believe in the interaction scenario and to rate the robot as less empathetic, regardless of condition. This correlation is stronger for ratings in the affective condition than in the cognitive condition, and there is a positive correlation between the NARS measure and difference between the empathy ratings in the two conditions, as shown in Figure 5. This means that participants with more negative views of robots rated the empathy of the robot in the cognitive condition higher, as expected. Therefore, H_{2A} is supported. While there is a small positive slope to the regression between NARS and the difference in ratings of belief between conditions, it is not significant, so H_{2B} is not supported by the data. That H_{2A} is supported while H_{2B} is not supported suggests that participants' prior attitude towards robots was more important than the specific type of empathy the robot was trying to convey in determining their belief in the interaction. This is intuitive, in that the type of empathy conveyed was only one factor among many that participants may have incorporated into their belief rating, while it was likely the dominant factor in their rating of the robot's empathy.

C. Ethics and trade-offs of deception in HRI

Having robots express empathy inherently involves some deception through anthropomorphization, which has been identified as a potential ethical hazard in the world's first explicitly ethical standard in robotics: BS8611-2016 Guide to the ethical design and application of robots and robotic systems [34]. In Winkle *et al.* [35], the authors establish that this anthropomorphization may be a necessary component for effective HRI, but also explore the use of strategies to mitigate the potential hazard. One strategy is to "minimize the use of affective social interaction that suggests robots have 'emotional states' or social agency". This strategy mirrors the difference between the cognitive and affective conditions in our study, so it is unsurprising that Winkle *et al.* also found that reducing anthropomorphism (and thus the ethical risk) would make HRI more effective for some users and less effective for others. We built on this result by describing how prior negative attitudes towards robots can explain some of the differences in this effectiveness.

D. Limitations

There are several limitations to this study. This work was conducted during the COVID-19 pandemic, making in-person data collection at scale difficult. A natural limitation of studying interaction preferences for a robot via AMT is that participants were not able to interact with the robot directly. The resulting reported preferences may differ from those that would be given as a result of in-person interactions. Additionally, the demographics of the participants who selected the task on AMT are not balanced, with almost three-quarters identifying as white and three-quarters as male, preventing generalization to more diverse user populations.

Another limitation of the study was the short and constrained nature of the dialog shown to participants. In order

to balance the gender of the actor (one male and one female), two videos were shown to each participant for each condition. This meant each participant watched a total of four videos, which had to be kept short in order to ensure participants would stay engaged with each video and complete the study in a reasonable amount of time. To avoid creating confounds in the study dialog, the structure of what the robot said was kept constant, with only the relevant emotion and the type of empathy expressed changing in each video. This constrained interaction structure was necessarily repeated in each video and may also have affected the participants' views of the robot.

E. Future Work

In this work we explored user preferences for different types of empathetic statements by socially assistive robots. The results and insights aim to inform future work on designing empathetic robot utterances that will be perceived as appropriately empathetic and believable by users. Future work can also explore customizing the type of empathetic response to each user based on their prior attitude towards robots. Finally, additional studies should verify that the reported AMT results transfer to in-person human-robot interactions.

VI. CONCLUSION

This work contributes a novel study of the differences in user perception of affective and cognitive empathetic statements made by a robot to an actor's disclosure. We report several findings, including that participants rated affective empathetic statements as both more empathetic and more believable than cognitive empathetic statements, that the participant's belief in the interaction was highly correlated with the perceived empathy of the robot in both the affective and cognitive conditions, and that more negative attitudes towards robots were correlated with a higher rating of empathy in the cognitive condition than the affective condition. We hope that this work will inspire continued development of ethical human-centered HRI.

ACKNOWLEDGMENT

This work was supported by the National Science Foundation National Robotics Initiative, IIS-1925083. We thank the actors who volunteered their time to record our study videos.

REFERENCES

- [1] N. Eisenberg, C. L. Shea, G. Carlo, and G. P. Knight, "Empathy-related responding and cognition: A "chicken and the egg" dilemma," in *Handbook of moral behavior and development*. Psychology Press, 2014, pp. 85–110.
- [2] Ö. N. Yalçın, "Evaluating empathy in artificial agents," *arXiv preprint arXiv:1908.05341*, 2019.
- [3] M. L. Hoffman, *Empathy and moral development: Implications for caring and justice*. Cambridge University Press, 2001.
- [4] G. Sanchez, M. W. Peterson, E. D. Musser, I. Galynker, S. Sandhu, and A. E. Foster, "Measuring empathy in health care," in *Teaching Empathy in Healthcare*. Springer, 2019, pp. 63–82.
- [5] R. Hogan, "Development of an empathy scale," *Journal of Consulting and Clinical Psychology*, vol. 33, no. 3, p. 307–316, 1969.
- [6] M. H. Davis, "Measuring individual differences in empathy: Evidence for a multidimensional approach," *Journal of personality and social psychology*, vol. 44, no. 1, p. 113, 1983.
- [7] R. N. Spreng*, M. C. McKinnon*, R. A. Mar, and B. Levine, "The toronto empathy questionnaire: Scale development and initial validation of a factor-analytic solution to multiple empathy measures," *Journal of personality assessment*, vol. 91, no. 1, pp. 62–71, 2009.
- [8] D. Jolliffe and D. P. Farrington, "Development and validation of the basic empathy scale," *Journal of adolescence*, vol. 29, no. 4, pp. 589–611, 2006.
- [9] C. A. Lietz, K. E. Gerdes, F. Sun, J. M. Geiger, M. A. Wagaman, and E. A. Segal, "The empathy assessment index (eai): A confirmatory factor analysis of a multidimensional model of empathy," *Journal of the Society for Social Work and Research*, vol. 2, no. 2, pp. 104–124, 2011.
- [10] S. W. Mercer, M. Maxwell, D. Heaney, and G. Watt, "The consultation and relational empathy (care) measure: development and preliminary validation and reliability of an empathy-based consultation process measure," *Family practice*, vol. 21, no. 6, pp. 699–705, 2004.
- [11] G. C. Kane, J. L. Gotto, S. West, M. Hojat, and S. Mangione, "Jefferson scale of patient's perceptions of physician empathy: preliminary psychometric data," *Croatian medical journal*, vol. 48, no. 1, pp. 81–86, 2007.
- [12] C. L. Bylund and G. Makoul, "Examining empathy in medical encounters: an observational study using the empathic communication coding system," *Health communication*, vol. 18, no. 2, pp. 123–140, 2005.
- [13] W. Tavares, S. Boet, R. Theriault, T. Mallette, and K. W. Eva, "Global rating scale for the assessment of paramedic clinical competence," *Prehospital emergency care*, vol. 17, no. 1, pp. 57–67, 2013.
- [14] D. L. Johanson, H. S. Ahn, B. A. MacDonald, B. K. Ahn, J. Y. Lim, E. Hwang, C. J. Sutherland, and E. Broadbent, "Pay attention! the effect of attentional behaviours by a robotic receptionist on user perceptions and behaviors,"
- [15] I. Leite, S. Mascarenhas, A. Pereira, C. Martinho, R. Prada, and A. Paiva, "“why can't we be friends?” an empathic game companion for long-term interaction," in *International Conference on Intelligent Virtual Agents*. Springer, 2010, pp. 315–321.
- [16] L. Charrier, A. Rieger, A. Galdeano, A. Cordier, M. Lefort, and S. Has-sas, "The rope scale: a measure of how empathic a robot is perceived," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 656–657.
- [17] L. Charrier, A. Galdeano, A. Cordier, and M. Lefort, "Empathy display influence on human-robot interactions: a pilot study," in *Workshop on Towards Intelligent Social Robots: From Naive Robots to Robot Sapiens at the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2018)*, 2018, p. 7.
- [18] K. Daher, J. Casas, O. A. Khaled, and E. Mugellini, "Empathic chatbot response for medical assistance," in *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*, 2020, pp. 1–3.
- [19] A. Reghunath, "Expression of empathy in social virtual bots used for genetic counseling," 2021.
- [20] A. Paiva, "Empathy in social agents," *International Journal of Virtual Reality*, vol. 10, no. 1, pp. 1–4, 2011.
- [21] A. Paiva, I. Leite, H. Boukricha, and I. Wachsmuth, "Empathy in virtual agents and robots: A survey," *ACM Transactions on Interactive Intelligent Systems (TiIS)*, vol. 7, no. 3, pp. 1–40, 2017.
- [22] A. Tapus and M. J. Mataric, "Emulating empathy in socially assistive robotics," in *AAAI spring symposium: multidisciplinary collaboration for socially assistive robotics*, 2007, pp. 93–96.
- [23] F. Hegel, T. Spexard, B. Wrede, G. Horstmann, and T. Vogt, "Playing a different imitation game: Interaction with an empathic android robot," in *2006 6th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2006, pp. 56–61.
- [24] L. D. Riek, P. C. Paul, and P. Robinson, "When my robot smiles at me: Enabling human-robot rapport via real-time head gesture mimicry," *Journal on Multimodal User Interfaces*, vol. 3, no. 1, pp. 99–108, 2010.
- [25] I. Leite, G. Castellano, A. Pereira, C. Martinho, and A. Paiva, "Empathic robots for long-term interaction," *International Journal of Social Robotics*, vol. 6, no. 3, pp. 329–341, 2014.
- [26] H. Cramer, J. Goddijn, B. Wielinga, and V. Evers, "Effects of (in) accurate empathy and situational valence on attitudes towards robots," in *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2010, pp. 141–142.
- [27] P. Kadam and S. Bhalerao, "Sample size calculation," *International journal of Ayurveda research*, vol. 1, no. 1, p. 55, 2010.
- [28] S. B. Green, "How many subjects does it take to do a regression analysis," *Multivariate behavioral research*, vol. 26, no. 3, pp. 499–510, 1991.

- [29] “Luxai - award winning social robots for autism and special needs education,” Aug 2021. [Online]. Available: <https://luxai.com/>
- [30] Interaction-Lab, “interaction-lab/harmoni: Controller code for human and robot modular open interactions.” [Online]. Available: <https://github.com/interaction-lab/HARMONI>
- [31] —, “Interaction-lab/cordial-public: Robot control for speech and synchronized behaviors, phone-based javascript face, and tablet interface.” [Online]. Available: <https://github.com/interaction-lab/cordial-public>
- [32] “Text-to-speech: Lifelike speech synthesis — google cloud.” [Online]. Available: <https://cloud.google.com/text-to-speech>
- [33] T. Nomura, T. Kanda, and T. Suzuki, “Experimental investigation into influence of negative attitudes toward robots on human–robot interaction,” *Ai & Society*, vol. 20, no. 2, pp. 138–150, 2006.
- [34] BSI-2016, “Bs 8611: 2016 robots and robotic devices: Guide to the ethical design and application of robots and robotic systems,” 2016.
- [35] K. Winkle, P. Caleb-Solly, U. Leonards, A. Turton, and P. Bremner, “Assessing and addressing ethical risk from anthropomorphism and deception in socially assistive robots,” in *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 2021, pp. 101–109.