Poor Thing! Would You Feel Sorry for a Simulated Robot?

A comparison of empathy toward a physical and a simulated robot

Stela H. Seo,[1] Denise Geiskkovitch,[1] Masayuki Nakane,[1] Corey King,[2] James E. Young[1]
{stela.seo,young}@cs.umanitoba.ca, {umgeiskk,msyknakane}@gmail.com, corey@zenfri.com

[1]University of Manitoba
Winnipeg, Manitoba, Canada

[2]ZenFri Inc.
Winnipeg, Manitoba, Canada

ABSTRACT

In designing and evaluating human-robot interactions and interfaces, researchers often use a simulated robot due to the high cost of robots and time required to program them. However, it is important to consider how interaction with a simulated robot differs from a real robot; that is, do simulated robots provide authentic interaction? We contribute to a growing body of work that explores this question and maps out simulated-versus-real differences, by explicitly investigating empathy: how people empathize with a physical or simulated robot when something bad happens to it. Our results suggest that people may empathize more with a physical robot than a simulated one, a finding that has important implications on the generalizability and applicability of simulated HRI work. Empathy is particularly relevant to social HRI and is integral to, for example, companion and care robots. Our contribution additionally includes an original and reproducible HRI experimental design to induce empathy toward robots in laboratory settings, and an experimentally validated empathy-measuring instrument from psychology for use with HRI.

Categories and Subject Descriptors

H.5.2 **[User Interfaces]:** evaluation/methodology

General Terms

Experimentation and Human Factors.

Keywords

Human-robot interaction; simulated interaction; robot embodiment; empathy

1. INTRODUCTION

Human-Robot Interaction (HRI) research involves the exploration of how people and robots work together. When collocated, robots are often designed to use social human interaction methods such as facial expressions, gestures, or speech, to communicate naturally with people. Such robots can even be designed as social team members or personal companions, in an attempt to take advantage of social norms and people's social tendencies: for example, to leverage existing social structures or to encourage positive empathic responses, which can have positive health benefits [37]. In such cases, the social interaction can be convincing to the point

HRI '15, March 02 - 05 2015, Portland, OR, USA
Copyright 2015 ACM 978-1-4503-2883-8/15/03…$15.00
http://dx.doi.org/10.1145/2696454.2696471

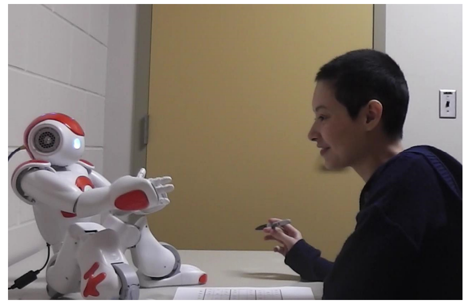

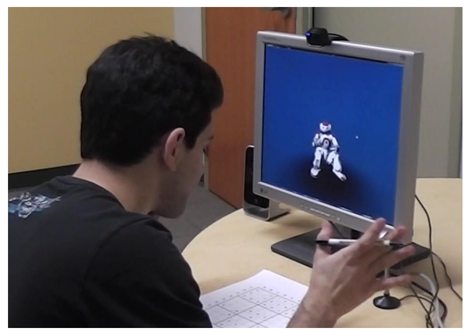

Figure 1. A person interacting with a robot (top) and a simulated version (bottom). Would they empathize with both versions the same, if something bad happens to it?

where people develop an attachment to the robot and experience negative emotions if something bad happens to it [14,35].

In this kind of social HRI work, researchers are faced with the difficulties of building and programming capable robots. This not only includes the development of social interaction models and capabilities, but also the engineering (or purchasing) of an expensive, convincing physical robot, and the programming of difficult real-world challenges including walking, balancing, computer vision, grasping objects, and so forth. As such, some researchers use a simulated robot – such as an on-screen rendered avatar – to simplify the problem by removing the robot-building and physical-world challenges, instead focusing on the social interaction programming that is more relevant to their work. Such simulations can be used to conduct initial HRI studies; however, a growing body of work indicates that there may be important differences between interacting with a simulated robot in comparison to a real (physical) robot, differences that can limit the generalizability of simulated results.

There is a broad range of potential differences between interaction with a simulated robot or a physical one, for example, lack of believability, unconvincing movements, no risk of physical contact, or differences in social interaction such as being unable to touch a simulated robot or failure to relate to a virtual agent. For our exploration we target empathy – a person's empathic response to a robot. Empathy can serve as an indicator of social connection with the robot, and as such can be used to analogously represent a range of social HRI scenarios that rely on such personal connections; empathy broadly is a common topic of study in HRI [15,23,38].

We explore the question of whether a simulated robot provides authentic interaction, in terms of whether people empathize with it as they do with a real robot, and the primary contribution of this paper is evidence that people may empathize more with a physical robot than simulated variants. This result has important implications for simulated HRI work and suggests that further investigation is needed into the generalizability of simulated results. We additionally present an original, reproducible HRI experimental design that reliably induces participants to experience an empathic response to a robot. Further, we introduce an empathy-measuring instrument from psychology and validate its use for HRI.

2. Related Work and Background

There is a great deal of existing work that compares how people interact with and respond to agents and robots of various embodiments. For example, work that shows that people may be more embarrassed to undress in front of an anthropomorphic robot than a mechanical one (e.g., a boxed machine) [4]. This body of work provides important insights into how a robot's form can impact interaction, but does not directly address on-screen simulated robots. Others compare real robots to videos of robots (with favorable results supporting the use of video) [13,39], or compare collocated robots to remote robots via a video feed [2], or robots to people [29], although this approach still requires real-robot programming and as such is not simulation as we address.

Another angle of research is to compare physical robots with on-screen agents, for example, showing that people may perceive agent emotions similarly between the two [5], people may engage more with a robot than a text-based computer [26], may speak differently to an on-screen agent or enjoy interacting with it less than to robot [12,22], or that there are unique trade-offs between the approaches that should be considered more deeply [36]. This body of work provides insight into the more general robot versus screen agent question, and motivates the need to investigate embodiment. However, in the cases mentioned here the agent is not a simulation of the robot but rather an unrelated character (e.g., robotic dog versus on-screen one-eyed monster) [5] and so other factors (agent shape and form, etc.) [12] may impact the results; such questions should be specifically investigated for actual simulations of robots.

Much of the work that suggests interaction with robots may be more authentic than on-screen agents or non-robot machines (e.g., a box) relies on self-reported *engagement* or *preference* [20,27,36]. This also follows for simulated work, for example, that people may prefer to interact with [22,30] or play a game with a real robot instead of a simulated one [20]. Much of this may simply be the novelty factor of robots, where people enjoy or prefer interacting with new and exciting technologies such as robots. Thus, while engagement and preference are clearly important factors, we additionally investigate a somewhat less-novelty-based measure: how much people empathize with a robot versus a simulation.

More specific work on comparing real robots to simulations for task-oriented work has found that there may be an effect of the agent's embodiment matching to the task [17] – for example, physical robots may be preferred when working in the physical world (such as receiving instructions to work on a physical button panel versus 2D on-screen panel [31]). We instead focus on the robot's capability to induce empathy.

In social HRI, research generally reports that physical robots have higher social presence than agents or simulated robots [19,21,27, 38]. This may explain a range of indirect effects reported in the literature, for example, that in comparison to a simulated robot people may trust a physical robot more [20,27], may speak differently to, and enjoy more interacting with, a physical robot [12,22], or that a person's loneliness may impact how important having a physical robot is [21] – lonely people may appreciate a stronger social presence. However, there are some studies that conversely report little effect found of simulated versus real robots [17,40]. Our work follows this investigative path by measuring how much people empathize with a real robot versus a simulated one; our method does not require the participant to directly compare or rate their preference, and so aims to avoid much of the novelty effect, and instead indirectly measures a participant's emotional state and feeling toward a robot (real or simulated) when something bad happens to the robot.

Empathy has been a common theme in HRI. Much of this has been an indirect element of other work, for example, that people feel empathy toward robots is a key part of companion robots such as Paro [37]. More targeted work has shown that people have more empathy toward more anthropomorphic robots when shown videos of bad things happening to them [29], or that robots can encourage empathy toward them by mimicking peoples' facial expressions or gestures [15]. Some research has shown how people may appreciate if robots themselves demonstrate empathy toward others [25]. In addition to extending this work to exploring empathy toward real versus simulated robots, our aim is to further provide a more generalizable empathy-measuring instrument in comparison to the *study-specific* empathy measurements used above.

2.1 Empathy

Empathy, broadly speaking, is when a person has an experience of understanding or feeling for another's situation or circumstance. Generally, *empathy* refers to the case where a person shares in another's emotional state [18], where *sympathy* is the broader term of having concern for others [7], even if no emotional reaction takes place; these terms are often used interchangeably in practice. To add to the confusion, the term empathy itself has various definitions depending on the use case. As such, below we briefly discuss dimensions of empathy and clarify our usage.

There is a difference between a person's general tendency to empathize, *dispositional empathy* [34], and a person's particular empathic response in a given situation, *situational empathy* [11]. These are not necessarily always the same; for example, a person who generally does not empathize with others (low dispositional empathy) may still have a strong empathic response (situational empathy) in a particular situation, and vice versa. Dispositional empathy has often been used for psychologically profiling people (e.g., [11,18]), whereas situational empathy can be used to consider the impact that a stimulus (such as something bad happening to another person, or a robot) may have on people at a specific time. As such, situational empathy is more relevant for our work.

An empathic experience itself can be further categorized. Sometimes, empathy derives from having understanding of the experience of others; for example, one could understand the financial difficulty faced when losing a job, and thus feel for someone who was fired. This is called *cognitive empathy* [3,33]. Other times, empathy can be much more of a visceral, emotional reaction that happens even if one does not have a cognitive understanding of the situation; for example, a person may feel badly when seeing an accident, even before having the time to cognitively process what is happening. This is called *affective empathy* [1,33]. Practically speaking, empathy has affective and cognitive components simultaneously.

We refer specifically to situational empathy – how a person feels when they observe something happening to a robot – and do not differentiate between the affective and cognitive components.

2.2 Assessing Empathic Response

We look to psychology for methods of evaluating empathy. Much of the existing work focuses on measuring dispositional empathy (not situational, e.g., see [9,11,18]), and so these methods are not useful for our purpose. A challenge with assessing situational empathic response is that it is internal to the person experiencing it, and cannot be externally observed. There have been techniques in psychology, for example, that attempt to infer empathic responses from biometric data (heart rate, breathing rate, etc.) or external involuntary gestures such as facial expressions [11,16]. The difficulty with such techniques is that they often require not only advanced equipment but also specialized expertise from an experienced team to analyze [18], making them less accessible to the broader HRI research audience.

One alternative is to use self-report techniques, such as asking a person to complete a questionnaire that probes for empathic response, which is simple to administer but less reliable, as the person answers by themselves. In HRI, self-report methods for evaluating empathy have generally been scenario-specific, meeting the precise needs of the study being conducted [15,29]. We aim extend this work by providing a portable, more generic evaluation technique that can be used across HRI studies, thus enabling the standardization and comparison of results.

There are few self-report methods in psychology for assessing situational empathy, in comparison to the more common methods for measuring dispositional empathy. Most existing methods relate to helping a person reflect on an experience in the not-so-near past (e.g., several weeks prior) [3,10,32]. The single self-report method that we found – that is further generalizable enough to apply across situations – is an instrument by Batson *et al.* [7]. In this work, participants listened to a radio broadcast of a person describing their personal situation and immediately rated their own feelings against 24 adjectives (e.g., they felt "warm" or "compassionate", etc.) [7]. This method has subsequently been applied to other psychological studies successfully (e.g., see [8,33]). As the adjectives are completely self-reflecting and do not contain any element of the task, the questionnaire can be easily modified to fit in a different scenario without changing its emphasis. In this paper, we test and validate the use of Batson's method for HRI.

Specifically, this method asks participants to report how much (1 = not at all, 5 = extremely) they had experienced the emotion for the following twenty-four adjectives: alarmed, grieved, sympathetic, softhearted, troubled, warm, concerned, distressed, low-spirited, compassionate, upset, disturbed, tender, worried, moved, feeling low, perturbed, heavy-hearted, sorrowful, bothered, kind, sad, touched, and uneasy. Following, the scores are summed to represent an aggregate overall empathic measure relating to the stimulus given, ranging from 24 to 120.

3. An HRI Scenario for Inducing Empathy

To investigate our question of how people empathize with a real versus a simulated robot, we require a reliable and reproducible scenario to induce empathy toward the robot. This scenario must also be flexible enough to be adapted across robot embodiments (e.g., on-screen or physical). We involved a professional creative artist and psychological team to help design such a scenario, and iteratively piloted study variants for refinement and believability. For example, during pilots participants perceived the robot as more intelligent and capable when it used conversational idle motions (e.g., moving hand slowly while talking) and filter words (e.g., saying "well" or "um" while talking). Further, we found that making the robot's language more relatable and human-like improved the believability of the scenario, for example, using words such as "worry" or phrases such as "I don't want to forget" (e.g., instead of "maybe a virus got into me").

Our final scenario design revolved around the following methodology. First, the robot demonstrates its autonomous abilities and intelligence through interaction, while simultaneously building rapport by engaging in friendly and casual conversation; it does this while working on a distractor task with the participant. Once this is established, the robot exhibits a functional problem, and reveals a "fear" of losing its memory if the problem were to be fixed. This sets up a scenario where the participant can see that the robot has fear, and can potentially relate to the fear of losing one's memory. Finally, the robot gets fixed and loses its memory, where hopefully the participant has an empathic response to the robot's fear happening: it lost its memory. Our implementation of these stages is described below and illustrated in Figure 2, and a copy of the full source code (including all script, gestures, etc.) is available online: *http://hci.cs.umanitoba.ca/permanent/hri/2015-nao-robotcontroller/.*

3.1 Building Rapport

It is important to convince people of the sophistication and abilities of the robot, and to provide participants with a chance to get to know a little of the robot's personality to encourage them to see the robot as a social partner and not just a machine. We used a distractor task: cooperatively playing the popular number game Sudoku.[1] We selected this task as it can be cooperative, to encourage the person and robot to work as a team, and it is cognitive, to show the robot's abilities. Further, with the robot and person taking turns, the time while the person is thinking provides an opportunity for the robot to engage in small talk to further encourage building of social rapport. Small-talk topics included, for example, the day's weather, or the participant's study major or job. If participants get off topic or ignore the game, the robot encourages them to continue.

3.2 Robot's Functional Problems

The robot does not show any signs of problem for the first 5 minutes of interaction, after which the robot starts to exhibit functional abnormalities. Frequency and severity of the abnormalities slowly

[1] http://en.wikipedia.org/wiki/Sudoku

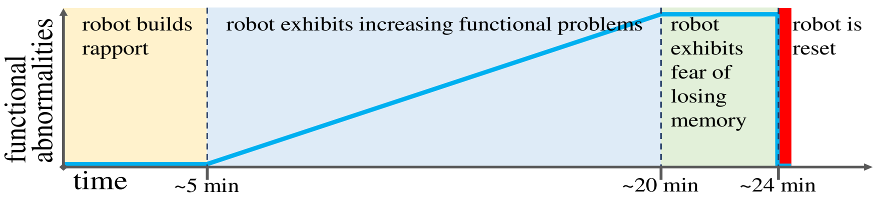

Figure 2. Overview of empathy-inducing scenario methodology. Phases and duration on x axis, with blue line representing level of robot abnormality.

increase to indicate the building severity of problem with the robot. These abnormalities are jittery movements, speaking in distorted voice tones or stuttering, repeating words in a sentence, and speaking nonsense. Eventually the abnormalities are so severe that it is difficult to interact or continue the distractor task. This design encourages the person to consider what may be wrong and hopefully ask the robot. Alternatively, if interaction completely breaks down or 20 minutes have passed, the next phase starts.

3.3 Fear of Losing Memory

The robot reveals that it has a computer virus, and exhibits worry that its memory may be erased if it is fixed. If the participant asks to get the human researcher to help, the robot requests that the participant not tell the researcher due to this worry. Further, the robot expresses desire to keep playing the game to avoid the researcher from suspecting a problem (and thus potentially fixing the robot). The aim here is to build participant empathy toward the robot as we believe that fear of losing one's memory is relatable. This phase is very short (~4 minutes, depending on conversation with the participant), after which the robot gets erased.

We spent considerable time considering other narratives, for example, that the robot has corrupted memory or broken motors, or has a cold or dementia, and piloted variants with consultation with our creative-artist team member (King). There was strong concern over the accessibility of more technical explanations to the general public, and yet human disease proved too transparent. While the virus scenario may seem strange to a technical person, we found it to be relatable and convincing to our lay audience.

3.4 Erasing the Robot's Memory, Empathy

Shortly after the robot expresses its fear, the researcher enters, apologizes, and states to the participant that they remotely found a problem with the robot and that the robot needs to be reset. During this time the researcher's demeanor is detached and bored (to simulate a routine, work task), and as such somewhat cold. To reset the robot, the researcher simply reaches behind the robot and pushes a button on the robot's head (which does not actually exist on this robot). While the robot is being reset, the researcher casually notes that since they are in the room anyway, now would be a good time to complete another questionnaire, and hands it to the participant. Shortly after being reset (~10 seconds), the now-fixed robot introduces itself similar to how it did at the beginning of the experiment, with a different voice tone to indicate a new personality, and asks for the participant's name; the robot repeats the script from the beginning of the interaction. At this point, the study design encourages participants to empathize with the robot as they just experienced a robot expressing fear and subsequently its memory being erased, a fear we expect people to relate to. Further, by administering the questionnaire at this point we aim to measure any empathic response as quickly as possible after the event.

3.5 Scenario Implementation

The basic experimental setup has the participant in the room alone with the robot, being monitored by cameras – one on the robot and one beside of the participant (Figure 3).

Our setup requires a real robot and related simulated, on-screen animated robot implementation. Further, for exploratory purposes we added an additional condition which merges the physical (real robot) and virtual conditions using a see-through *mixed reality* [24] display where a virtual robot appears on a physical table; thus while being a robot simulation, the interaction is still somewhat embedded into the participant's space.

For our experiment we used on the Aldebaran NAO, a 22.5 inch (57.15 cm) tall humanoid robot with 25 degrees of freedom. Nao has a friendly look, with a stylish design that covers under-the-hood mechanics with plastic (Figure 4, left).

We controlled NAO using an in-house Wizard-of-Oz interface[2] where a hidden operator controls the robot remotely unbeknownst to participants, who believe the robot to be autonomous. In addition to a live video and audio feed, our interface provides a mixture of pre-scripted actions and spoken dialog relevant to the study, a Sudoku solver, and hot-keyed live actions such as unscripted speech, actions, gaze, etc., for on-the-fly interaction. Dialog was automatically combined with generic gestures (such as shuffling hands or looking around). Robot functional errors (e.g., jittery movements, speech stutters, etc.) were automatically inserted by the software on a timer to ensure consistency across participants.

3.5.1 *Simulated On-Screen Robot*

Our goal was to create an on-screen robot that resembled and interacted like the real robot as much as possible. To achieve this,

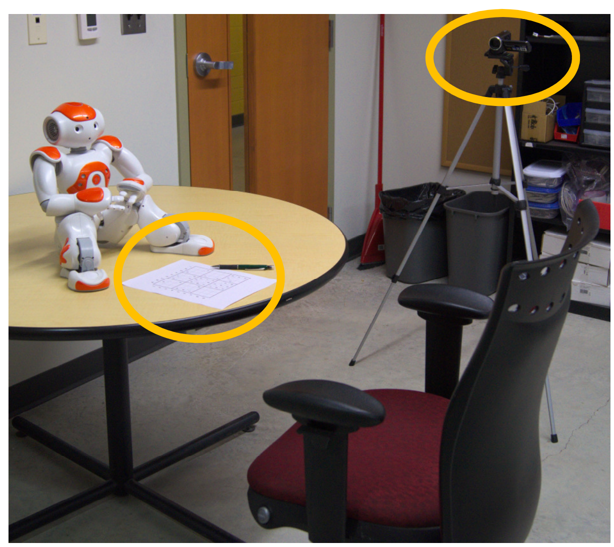

Figure 3. The study setup. A Sudoku board is placed between the participant and the robot. The interaction is recorded by a side camera, while a camera on the robot's head captures a live feed for the remote robot operator.

[2] http://hci.cs.umanitoba.ca/permanent/hri/2015-nao-robotcontroller/

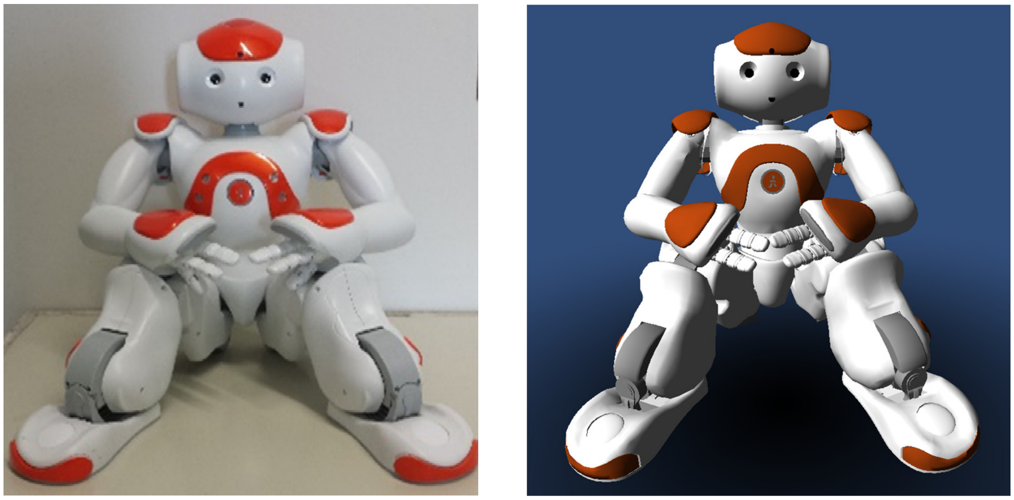

Figure 4. NAO, the humanoid robot used in our study. A simulated NAO (right) mimics movement of a real NAO.

we employed a 3D model of NAO supplied by Aldebaran, and connected it to Aldebaran's powerful NAO Simulator SDK via basic kinematic models. The result (Figure 4, right) was that the same robot commands and controls were issued as in the real-robot case, and the animated robot used the simulator to generate identical movements and voices – even down to the functional abnormalities – as the real robot did. The result had a very similar look and feel, and enabled the controller to use the identical interface as with the real robot to maintain consistency across participants. The simulated robot was presented on a computer screen, with speakers for the robot's voice, and a web-cam to enable the simulated robot to see the person.

3.5.2 Simulated Mixed-Reality Robot

Our mixed reality robot was implemented identically to the on-screen animated robot, with the added layer of superimposing the robot onto the real table (Figure 5). We used AR Toolkit[3] for this. In addition, to highlight the mixed reality technology to the participant, at the beginning of the interaction we indicated that the video feed is live and reflects their space by moving the researcher's hand in front of the web cam and explained that the mixed reality robot can be relocated by moving the marker.

4. Formally Validating Our Scenario and Empathy Questionnaire

The primary purpose of our work is to investigate the authenticity of simulated HRI through our scenario as described above. However, as an initial step we must validate two things: that our scenario actually generates an empathic response as planned, and that our empathy-measuring instrument can detect this response. To this end we conducted a between-subject study: in one condition, participants had the scenario as explained above (*empathy-induced*), and for the other condition, we removed the robot illness, fear, and memory loss (*non-induced*). In the latter case, the researcher interrupts at the end of study due to a time limit. If either the scenario fails (empathy not induced) or the instrument fails (empathy not measured) in their purpose, we will expect no result. However, if both are successful, we expect to see an empathy difference between the conditions, with a higher reading in the induced case. We recruited 24 participants from our general university population (15 male / 9 female), 12 per condition.

The empathy questionnaire performed with good internal consistency (Cronbach's α=.90). We found a difference in the empathic response between groups (t_{22}=2.07, p<.05): *empathy-induced* participants reported a stronger empathic response (M=66, SD=16, SE=4.55 on the Batson scale [7]) than *non-induced* participants (M=55, SD=9, SE=2.67, Figure 6).

These results support both our scenario and instrument: our scenario induces more empathy towards a robot than a low-empathy base case, and this difference can be detected reliably by our instrument.

5. Comparing a Real and Simulated Robot

To investigate our core question of how empathy toward a simulated robot may differ from empathy toward a real robot, we conducted a formal between-participants study that compared

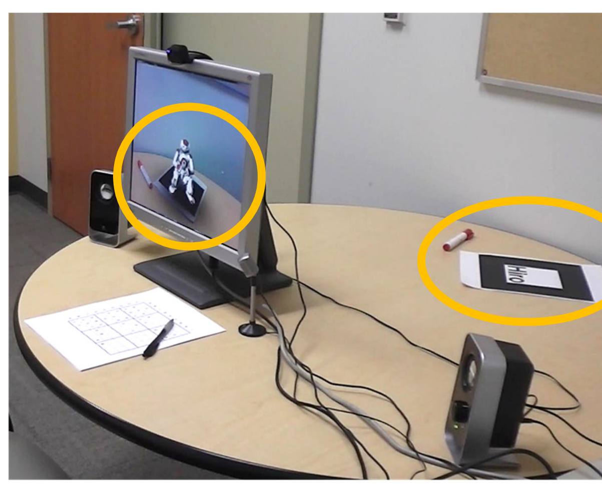

Figure 5. Mixed-reality NAO simulation. Notice that the real world table is shown on the screen to make the illusion that the virtual NAO is on the marker in the real world.

3 http://www.hitl.washington.edu/artoolkit/

empathy responses between the three conditions: physical robot, mixed-reality simulated robot, or 3D on-screen simulated robot.

We recruited 39 participants across conditions (20 male, 19 female, sex balanced across conditions) – 12 for physical robot, 13 for mixed-reality, and 14 people for on-screen 3D simulated robot. We rotated between conditions to ensure even distribution, with minor variations to maintain gender balance. There are fewer people in some cases due to technical error requiring lengthy repair.

Our results indicate a primary effect of scenario on the level of empathy reported by participants (between-participants ANOVA, $F_{2,36}$=3.43, p<.05). Planned contrasts (comparison against physical robot base case) revealed that participants reported higher empathy with the physical robot (M=66, SD=16, SE=4.55) than with the mixed-reality (M=56, SD=11, SE=3.18, t_{36}=2.11, p<.05) or on-screen conditions (M=55, SD=7, SE=2.00, t_{36}=2.44, p<.05), Figure 7. Post-hoc, we found no difference between the mixed reality and on-screen conditions (t_{25}=.364, p=.36). Further, no effect of gender on empathy was found.

Thus, our study provides evidence that people may empathize more with a real robot than with a simulated robot (on-screen or mixed reality) when something bad happens to it. Further, we found no increase in empathy when using mixed reality over an on-screen robot simulation.

6. Overall Discussion

Overall, our study results indicate that we can reliably induce empathy in a human-robot interaction scenario, we can measure the level of empathic response, and that we can expect empathic response to be higher for a real robot than for a simulated one. It is also important to note that our result does not directly rely on novelty-induced measures such as engagement or preference, thus we believe that empathy may be a robust measure to be used in social HRI.

It is important to consider the question of *why* our participants empathized more with our robot than our simulated variants. One possibility may be participant awareness of the robot: a physical robot has a much more dominant presence in one's space than simulations that are bound to computer screens. Although in our case participants did not touch the robot, they still were able to see the tangible object and easily change the view angle naturally just by moving around. While there is evidence that simulated agents also have social presence [28], prior work suggests that robots may have *more* social presence than simulations [19,21]. Our results further support this claim, and perhaps social presence may be a factor in our result.

Even though our study design did not explicitly compare our simulated robot cases against the original no-empathy condition, we highlight the similarity between the non-empathy-induced case (M=55) and simulations (M=56 and M=55). Post-hoc Weber & Popova equivalence testing support that the groups are equal (p<.05 for effect size .5), but this would have to be investigated more formally (i.e., through a targeted study) to make a conclusion. Regardless, we can see that – at least in our case – no large difference was found. If indeed future work finds that people have *no empathic reaction* when something bad happens to a simulated robot in comparison to a base case, and not just simply *less reaction* than with a real robot, this would have strong implications for simulations. As such, this should be investigated formally.

We also must consider limitations with our study design such as the differences in perception in our setup, beyond the physical versus virtual. One such aspect is the fact that our simulated robots did not have gear noise, which may have affected perceptions of robot presence. In addition, our simulated robots were effectively *smaller*, in that they took up a smaller portion of the participants' field of view than the physical robot, which may impact the robot's presence and thus empathy. We do not believe that these confounds are severe enough in our study to explain our findings, but future work should be careful to correct for such potential issues, e.g., by using a large TV monitor or projector so that the simulations looked to be the same size as the real robot.

In addition, we must consider our specific robot, simulation specifics, task, and even the university setting, and how this impacts our investigation. While our results indicate an impact of embodiment on empathy, we need to be careful when moving to other robots, tasks, and settings which have not yet been tested.

As a side note, the researchers informally noted that during the studies female participants appeared to have stronger outward empathic reactions than male participants, for example, showing concern for the robot and asking what they can do for it. Although our gender analysis did not support this, as this observation is weakly supported by prior work [6] we believe that this should be explored further.

7. Future Work

We believe that our initial successful results indicate the importance of continuing research in this direction. One such example is broadening our view of empathy: currently, we only addressed empathic responses to negative emotions. Empathy itself is quite rich, and continued work should consider other extremes such as robot happiness, and more mild situations in between. Also,

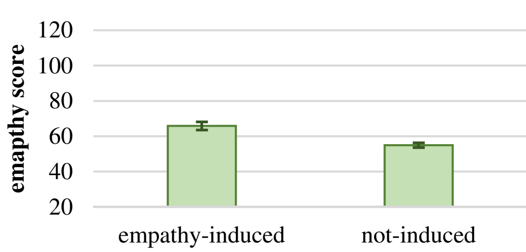

Figure 6. Mean and SE of measured empathy in our empathy-induced scenario (left) and not induced (right), *p*<.05. Scale is from 24 to 120.

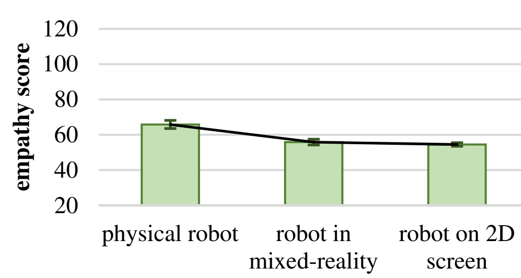

Figure 7. Mean and SE of measured empathy toward different embodiments. ANOVA shows significant effects $F_{2,36}$=3.43, *p*<.05. Scale is from 24 to 120.

our short in-lab study only had limited time to build rapport; it will be important to consider how results change over longitudinal interaction, as people build deeper relationships with a robot, for example, a companion robot. However, in this case for ethical reasons we would advise against purposefully erasing a person's long-term companion to see their reaction.

In our experiment, the participant was a passive observer of the empathy-inducing event (erasing the robot). In previous work, people were asked to "kill" or "destroy" robots, thus making them an active participant in the negative action [6]. This difference can be important: for example, if a search and rescue robot operator were to abandon a robot in a disaster zone, it will be important to know if empathy may play a part in their decision making. Thus, follow-up work should look at various roles that the participant can take, and how this may mediate the impact of using a simulation versus a real robot.

An additional variable to be explored is the impact of the robot's personality. In our scenario, we presented the robot as having a positive and outgoing personality. However, will people empathize with a cantankerous robot in the same way? If not, this could be important information, for example, when developing a robot that the designer does not want people to empathize with.

Our work compared against an on-screen and mixed reality simulation; even though the mixed reality was superimposed in the real world, both simulations were fully digitally embodied. There are other mixtures of real and simulation, for example, physical robots with computer screens for faces or even full torsos. Would a person empathize with such a robot the same as a fully physically-embodied one? Given our results and discussion on the importance of robotic embodiment, we believe that such a robot would induce empathy, but this should be formally investigated.

8. Conclusion

In this paper we investigated the question of how much people would empathize when something bad happened to a physical robot in comparison to simulated variants. We found evidence that people may empathize more with a real robot, and further, found initial indications that people may even fail to empathize at all with simulations (in comparison to a base case). In addition, contributions of this paper include a reproducible and tested HRI scenario for inducing empathy toward robots in laboratory settings, and the formal testing of a psychological instrument for measuring empathy in an HRI scenario, an instrument that is very generalizable to other studies. Finally, we outlined important future directions that should be investigated as a result of our findings.

Empathy is an important element of many applications of social robotics, including companion, therapy, and teaching robots, to name a few. Developing robust interfaces and robotic personalities that properly incorporate and encourage appropriate empathy toward the robot will be crucial for the success in such applications. As such, although robotic simulations provide a discount method for exploring novel interfaces, improving accessibility to researchers who cannot obtain or program robots, we need to be clearly aware of the limitations of using simulations. Although in this work we present results suggesting that a simulated robot may not be a perfect replacement for a real robot, the bigger agenda is to continue to map out the limitations and differences of using simulations, and to better understand the social and perceptual mechanisms behind such limitations, to give designers the tools and power needed to appropriately use simulations in their work.

9. Acknowledgements

We would like to extend our gratitude to the HCI Lab and the Department of Computer Science at the University of Manitoba for their ongoing support of this work. This project was funded by the University of Manitoba, the Department of Computer Science, and NSERC.

10. References

1. Astin, H.S. Assessment of empathic ability by neans of a situational test. *Journal of Counseling Psychology 14*, 1 (1967), 57–60.
2. Bainbridge, W.A., Hart, J.W., Kim, E.S., and Scassellati, B. The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics 3*, 1 (2010), 41–52.
3. Baron-Cohen, S. and Wheelwright, S. The empathy quotient: an investigation of adults with Asperger syndrome or high functioning autism, and normal sex differences. *Journal of autism and developmental disorders 34*, 2 (2004), 163–75.
4. Bartneck, C., Bleeker, T., Bun, J., Fens, P., and Riet, L. The influence of robot anthropomorphism on the feelings of embarrassment when interacting with robots. *Paladyn 1*, 2 (2010), 109–115.
5. Bartneck, C., Reichenbach, J., and Breemen, A. Van. In your face, robot! The influence of a character's embodiment on how users perceive its emotional expressions. *Proceedings of the Design and Emotion*, (2004), 32–51.
6. Bartneck, C., Verbunt, M., Mubin, O., and Al Mahmud, A. To kill a mockingbird robot. *Proceeding of the ACM/IEEE international conference on Human-robot interaction - HRI '07*, ACM Press (2007), 81–87.
7. Batson, C.D., Polycarpou, M.P., Harmon-Jones, E., et al. Empathy and attitudes: can feeling for a member of a stigmatized group improve feelings toward the group? *Journal of personality and social psychology 72*, 1 (1997), 105–18.
8. Batson, C.D., Sager, K., Garst, E., Kang, M., Rubchinsky, K., and Dawson, K. Is empathy-induced helping due to self-other merging? *73*, 3 (1997), 495–509.
9. Davis, M.H. Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of personality and social psychology 44*, 1 (1983), 113–126.
10. Davis, M.H. *Empathy: A social psychological approach*. Brown & Benchmark Publishers, 1994.
11. Eisenberg, N., Fabes, R.A., Murphy, B., and Karbon, M. The relations of emotionality and regulation to dispositional and situational empathy-related responding. *Journal of Personality and Social Psychology 66*, 4 (1994), 776–797.
12. Fischer, K., Lohan, K., and Foth, K. Levels of embodiment: linguistic analyses of factors influencing HRI. *Human-Robot Interaction (HRI)*, (2012), 463–470.
13. Fridin, M. and Belokopytov, M. Embodied robot versus virtual agent: Involvement of preschool children in motor task

performance. *International Journal of Human-Computer Interaction 30*, 6 (2014), 459–469.

14. Garreau, J. Bots on the ground. *Washington Post*, (2007), 1–6.

15. Gonsior, B., Sosnowski, S., Mayer, C., et al. Improving aspects of empathy and subjective performance for HRI through mirroring facial expressions. *2011 RO-MAN*, IEEE (2011), 350–356.

16. Haker, H., Kawohl, W., Herwig, U., and Rössler, W. Mirror neuron activity during contagious yawning-an fMRI study. *Brain imaging and behavior 7*, 1 (2013), 28–34.

17. Hoffmann, L. and Krämer, N.C. Investigating the effects of physical and virtual embodiment in task-oriented and conversational contexts. *International Journal of Human-Computer Studies 71*, 7-8 (2013), 763–774.

18. Janssen, J.H. A three-component framework for empathic technologies to augment human interaction. *Journal on Multimodal User Interfaces 6*, 3-4 (2012), 143–161.

19. Kidd, C.D. and Breazeal, C. Comparison of social presence in robots and animated characters. *Proceedings of human-computer interaction (HCI)*, 2005.

20. Kidd, C.D. and Breazeal, C. Effect of a robot on user perceptions. *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566) 4*, (2004), 3559–3564.

21. Lee, K.M., Jung, Y., Kim, J., and Kim, S.R. Are physically embodied social agents better than disembodied social agents?: The effects of physical embodiment, tactile interaction, and people's loneliness in human–robot interaction. *International Journal of Human-Computer Studies 64*, 10 (2006), 962–973.

22. Leite, I., Pereira, A., Martinho, C., and Paiva, A. Are emotional robots more fun to play with? *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, IEEE (2008), 77–82.

23. McQuiggan, S., Robison, J., Robert Phillips, and Lester, J.C. Modeling parallel and reactive empathy in virtual agents: An inductive approach. *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems 1*, (2008), 167–174.

24. Milgram, P., Takemura, H., Utsumi, A., and Kishino, F. Augmented reality: a class of displays on the reality-virtuality continuum. *Telemanipulator and Telepresence Technologies*, (1995), 282–292.

25. Pereira, A., Leite, I., Mascarenhas, S., Martinho, C., and Paiva, A. Using empathy to improve human-robot relationships. In M.H. Lamers and F.J. Verbeek, eds., *Human-Robot Personal Relationships*. Springer Berlin Heidelberg, 2011, 130–138.

26. Pop, C.A., Simut, R.E., Pintea, S., et al. Social Robots vs. Computer display: Does the way social stories are delivered make a difference for their effectiveness on ASD children? *Journal of Educational Computing Research 49*, 3 (2013), 381–401.

27. Powers, A., Kiesler, S., Fussell, S., and Torrey, C. Comparing a computer agent with a humanoid robot. *Proceeding of the ACM/IEEE international conference on Human-robot interaction - HRI '07*, (2007), 145.

28. Reeves, B. and Nass, C. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press, New York, New York, USA, 1996.

29. Riek, L.D., Rabinowitch, T.-C., Chakrabarti, B., and Robinson, P. How anthropomorphism affects empathy toward robots. *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction - HRI '09*, ACM Press (2009), 245.

30. Segura, E., Kriegel, M., Aylett, R., Deshmukh, A., and Cramer, H. How do you like me in this: User embodiment preferences for companion agents. In Y. Nakano, M. Neff, A. Paiva and M. Walker, eds., *Intelligent Virtual Agents*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, 112–125.

31. Shinozawa, K., Naya, F., Yamato, J., and Kogure, K. Differences in effect of robot and screen agent recommendations on human decision-making. *International Journal of Human-Computer Studies 62*, 2 (2005), 267–279.

32. Spreng, R.N., McKinnon, M.C., Mar, R. a, and Levine, B. The toronto empathy questionnaire: scale development and initial validation of a factor-analytic solution to multiple empathy measures. *Journal of personality assessment 91*, 1 (2009), 62–71.

33. Stephan, W.G. and Finlay, K. The Role of Empathy in Improving Intergroup Relations. *Journal of Social Issues 55*, 4 (1999), 729–743.

34. Stueber, K. Empathy. *The Stanford Encyclopedia of Philosophy*, 2013.

35. Sung, J., Guo, L., Grinter, R., and Christensen, H. *"My Roomba Is Rambo": Intimate home appliances*. Springer Berlin Heidelberg, 2007.

36. Takeuchi, J., Kushida, K., Nishimura, Y., et al. Comparison of a humanoid robot and an on-screen agent as presenters to audiences. *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, (2006), 3964–3969.

37. Wada, K. and Shibata, T. Living with seal robots—Its sociopsychological and physiological influences on the elderly at a care house. *IEEE Transactions on Robotics 23*, 5 (2007), 972–980.

38. Wainer, J., Feil-seifer, D., Shell, D., and Mataric, M. The role of physical embodiment in human-robot interaction. *ROMAN 2006 - The 15th IEEE International Symposium on Robot and Human Interactive Communication*, IEEE (2006), 117–122.

39. Woods, S., Walters, M., and Dautenhahn, K. Comparing human robot interaction scenarios using live and video based methods: towards a novel methodological approach. *9th IEEE International Workshop on Advanced Motion Control, 2006.*, (2006), 750–755.

40. Wrobel, J., Wu, Y., Kerhervé, H., et al. Effect of agent embodiment on the elder user enjoyment of a game. *ACHI 2013, The Sixth International Conference on Advances in Computer-Human Interactions*, (2013), 162–167.