Running head: Mind Perception in Digital Agents

Mind the robot! Variation in attributions of mind to a wide set of real and fictional robots

Oliver Jacobs^{1,2}, Kamel Gazzaz², & Alan Kingstone¹

¹ University of British Columbia

² Maslo Research

Author note

Correspondence concerning this article should be addressed to Oliver Jacobs, 2136 West Mall, V6T 1Z4, Vancouver, BC, Canada. E-mail: ojacobs@psych.ubc.ca

Mind Perception in Digital Agents

Abstract

The rapid rise of computing power over the last half-century has prompted the desire to understand and develop a paradigm of affective computing systems that can recognize, process, and simulate human features, including qualities like empathy and morality. Quantitatively comparing different computing systems in their abilities to simulate human qualities has been a major technical challenge. A recent framework put forth by Gray, Gray, and Wegner (2007) provides promise as a new means for comparing a wide landscape of different digital agents both real and fictional. Using this framework, we sought to investigate if attributions of mind towards robots suggests that people perceive robots as capable of emulating different degrees of mind. We asked participants to rate the agency (the ability "to do") and experience (the ability "to feel") of 24 characters made up of humans, robots, inanimate objects, and animals. Although robots were collectively rated much lower than humans on agency and experience, there was significant variation among robots (both real and well-known fictional robots). This implies that building digital agents to imitate aspects of experience is a fruitful avenue for future development. In addition, age was a critical factor in people's attributions of agency and experience indicating that there may be a generational shift towards greater acceptance of robots' ability to both do and feel.

Keywords: mind perception, artificial intelligence, humanoid robots, agency, experience,

Word count: 217

Mind Perception in Digital Agents

Mind Perception in Humans and Digital Agents

Introduction

In the 1950s, the cognitive revolution led to a convergence of psychology, neuroscience, anthropology, linguistics, and computer science in the study of human cognition and its processes (Miller, 2003). Cognitive scientists of the era began to view the human mind as an information-processor that takes in environmental input, processes the data, and generates an output (McCorduck & Cfe, 2004). Fueled by the rapid rise of computing power by way of Moore's law, researchers have since sought to develop intelligent computing systems that function in a similar manner to human beings (McCorduck & Cfe, 2004). Whereas most researchers targeted rational tasks such as logic and inference, a subset took an interest in developing affective computers that can recognize, process, and simulate human emotions (Picard, 2000). This quest was facilitated by the nascence of multi-signal processing of different contextual consequences that lead to learning and development (Cai, 2006). Despite these advances, the question of whether machines have developed any sort of authentic emotional capability remains unanswered.

Historically, many researchers have used the Turing test to determine whether machines have developed sufficient intelligence to be considered as having acquired thought (Saygin et al., 2003). If a human participant was unable to determine whether they were interacting with another human or a machine, then the machine was said to be intelligent (Saygin et al., 2003). However, criticisms of this approach quickly emerged after early chatbots like ELIZA (Weizenbaum, 1966) began to pass the Turing test by tricking humans in certain situations. Essentially, these programs were capable of manipulating symbols to appear as though they have understanding—without

Mind Perception in Digital Agents

actually having any underlying understanding. Illustrating this point, John Searle argued that a computing system could not be shown to have a mind even if it demonstrates seemingly intelligent or human-like behavior (Searle, 1980). In his now famous thought experiment, Searle proposed the idea of the Chinese room to highlight this point. In a simplified version, imagine you have an English-speaking person enclosed in a room with an input feeding Chinese characters through a slot in the wall. With sufficient tools in the room, such as an English to Chinese dictionary, paper, and a pen, the person in the room would be able to pass output through the slot in the room by using the dictionary to translate the incoming Chinese characters. To an observer outside the room, it may appear as though the person inside the room is fluent in both English and Chinese despite this not being the ground truth. This parallels the problem with interpreting AI systems as having a real understanding of whatever intellectual or human-like properties it is imitating when a more parsimonious account would simply be that they are only manipulating physical symbols.

While philosophical debates relating to AI continues and criticisms of John Searle's idea have emerged (e.g. Dennett, 1993; Chalmers, 1992), researchers have since agreed that people often perceive intelligent or human-like qualities in AI regardless of whether that computing system has any real understanding of its behavior (Epley et al., 2007). Despite individual differences and demographic variables like age and gender playing a role, anthropomorphism of non-humans, including machines, is relatively universal (Waytz et al., 2010). Indeed, one study examining the activation of the mirror neuron system found that the same neural circuitry underlies the perception of human-like behavior in an anthropomorphized robot as the perception of an actual human's behavior (Gazzola et al., 2007). One major advantage of studying these

Mind Perception in Digital Agents

perceptions of human-like behavior in computing systems is that unlike trying to quantify a computer's understanding of its behavior, we can much more readily assess people's perceptions of how a computer might imitate intelligent or human-like behavior. Additionally, researchers have been empirically investigating anthropomorphic attributions for almost 80 years (for example, see Heider and Simmel, 1944), leaving a broad and solid literature upon which to ground theory. As a result, several empirical paradigms have emerged dedicated to 'seeing human' in nonhumans, some with widespread support (e.g. Epley et al. 2007; H. Gray et al., 2007a).

One of the methods that has become popular for seeing human in nonhumans is called the mind perception approach (or framework). The mind perception approach allows researchers to directly probe people's self-reported perceptions. A major advantage of this approach over other paradigms is its focus on the hallmark of humanity—the mind (K. Gray & Wegner, 2012). This provides a promising means of investigating human qualities like emotion and empathy in digital agents because those traits are considered an essential part of the human mind. The mind perception framework was introduced in the seminal work of Gray and colleagues (H. Gray et al., 2007a; H. Gray et al., 2007b). They used pairwise comparisons and factor analysis to examine the critical components of mind perception in humans and non-humans. They found that there were two principal components: agency, and experience. Agency refers to a collection of different capacities including self-control, morality, memory, emotion recognition, planning, communication, and thought. Experience refers to a collection of other capacities including hunger, fear, pain, pleasure, rage, desire, personality, consciousness, pride, embarrassment, and joy. These two separable components of agency and experience map respectively onto Aristotle's

Mind Perception in Digital Agents

distinction between moral agents (agents that morally can do right or wrong) and moral patients (patients that can have moral right or wrong done to them). In the initial Gray, Gray, and Wegner (2007) study, they included a social humanoid robot and found that while it scored in the midrange for agency, people assigned it very little experience. While the social humanoid scored higher on agency than a dog, baby, and a dead person, the dead person and every other character other than God scored higher on ratings of experience. However, Gray et al. only included one type of social humanoid robot (and a rather crude one at that), limiting their ability to observe if the mind perception framework could be used to compare different types of computing systems or machines in terms of their abilities to display evidence of agency and experience.

Subsequent work has demonstrated that the mind perception framework can add to the growing body of research that suggests anthropomorphism can facilitate the development of human-computer interactions, including affective computing systems (Eyssel et al., 2011; K. Gray & Wegner, 2012; Moore, 2016). For example, the mind perception framework has demonstrated real-world value in predicting the success of implementing social humanoid robots into everyday life. Stafford et al. (2014) found positive attitudes towards robots and perceptions of agency predicted social humanoid robot use and acceptance in residents of a retirement home. Interestingly they also found that experience ratings were not predictive of acceptance. This could be indicative of an inherent problem with increasing people's perceptions of how capable robots are of experiencing, as this could lead to the uncanny valley, a general feeling of unease due to unnaturally realistic imitations of the real world (K. Gray & Wegner, 2012; Mori et al., 2012). Others have also criticized the entire concept of 'tricking' people into believing robots have feelings, pointing out that while it is easy to build robots capable of doing more things, it is

Mind Perception in Digital Agents

near impossible to build robots capable of feeling more things. These criticisms parallel Chalmer's (1995) distinction between the easy and hard problems of consciousness wherein he claimed we can make great advances in understanding how the mind does things, but the real problems are how the mind comes to experience – the so-called qualia associated with being. But, is there a hard problem in robotics? Can robots, of today or tomorrow, compare with living beings on both agency, and more doubtfully, experience? Is there enough perceived variation between robots in the mind perception framework to suggest that people may perceive a mind in a robot, both in terms of a robot's agency *and*, crucially, its experience?

Present Study

To address this issue, the present study applies the mind perception framework to compare a large group of affective computing systems ranging from fictional characters like Wall-E and R2D2, to a real-world social humanoid robot (Sophia), to other more mundane digital agents like Siri, Sphero, and a robotic vacuum. If we observe a significant variance amongst the robots on agency, and most critically, experience, then it would suggest that the mind perception framework can be used as a means of assessing and comparing different affective systems. Furthermore, if there is a significant distribution on agency and experience amongst robots that are not fictional, it would suggest that there already exist real-world robots that people perceive as having different qualities of mind. However, if as some previous research suggests, there is little variation in the ratings of the robots, then it would suggest that developing computing systems that will be perceived as being able to do *and* feel will be a most challenging (and dubious) exercise both presently and in the foreseeable future. As the key measurement is how others perceive robots that are drawn from the past, present, and future; and past work indicates

Mind Perception in Digital Agents

considerable individual differences in anthropomorphism (Waytz et al., 2010), we thought it prudent to examine if demographic variables including age and sex affect perceived differences in agency and experience.

Methods

Participants

There were 123 participants that took part in the study. Participants were recruited via Amazon Mechanical Turk (MTurk). MTurk is an Amazon hosted network for contracting short term work for tasks like surveys. MTurk's robustness has been extensively researched and results support that MTurk is a valid means of collecting online data for the behavioural sciences, and is sometimes even superior to in-person collection (Casler et al., 2013; Hauser & Schwarz, 2016). Still, as a precaution, only MTurk experts were recruited, those who have successfully completed 1000 other tasks with a success rate of 99%. Participants mean age was 41.47 ($SD = 11.29$). There were 55 females and 68 males. All participants completed the study in the United States.

Material

The study was formatted as a Qualtrics survey. Each survey question was comprised of an image, a one-line description of the image, and two rating scales. In total, there were 24 images of human and non-human characters, each paired with two 0-100 rating scales for agency and experience. The characters were selected from four different categories, as per the original methodology of Gray, Gray, and Wegner (2007). The four categories were humans, animals, inanimate objects, and computational beings. The human category included a woman, a dead person, a baby, a fetus, a child, and a person in persistent vegetative state (PVS). The animal category included a dog, a frog, a chimpanzee, and bacteria. The inanimate category included a rock, a teddy bear, and a shovel. The machine category included Sophia, Sphero, R2-D2, Siri, Roomba, Maslo, Atlas, Beam, future conscious robot (FCR), Alexa, and Wall-E.

Mind Perception in Digital Agents

Procedure

The survey was distributed through MTurk and Qualtrics. First, participants signed a consent form and filled in a questionnaire. It was a standard demographic questionnaire querying the participant's age and gender. Next, the participants were presented with the image of an adult woman along with a short description. The participants rated the woman on agency and experience on a 0-100 scale. The participants then repeated the process with an image of a rock. After these first two characters (the adult woman and the rock), the participants rated the rest of the characters, which were presented in a randomized order, on agency and experience in the same manner. The definitions of agency and experience and their component capacities were included alongside the rating scales during each character presentation should a participant need to remind themselves of one or both definitions.

Results

The analyses are split into 2 components with each comparing how agency and experience attributions vary across characters. The first section considers all the characters and pertains more to manipulation checks. The second section examines only robots and specifically investigates if there is significant variation in the attributions made to all the robots (real and fictitious), and to the real robots alone.

All Characters

The purpose of this initial set of analyses is to ground our results in the broader literature and provide manipulation checks. Two one-way Greenhouse Geisser corrected within-subjects ANOVAs revealed significant differences amongst the 24 characters on agency, $F(7.38, 900.96) = 163.22$, $MSE = 1,624.27$, $p < .001$, and experience, $F(6.92, 844.69) = 365.88$, $MSE = 1,317.36$, $p < .001$. We did not include age or sex in these analyses given that it would not add to our manipulation checks. In terms of these checks, it appeared that people accurately understood the purpose of the task given the consistently low scores of agency and experience for inanimate characters like the rock, shovel, and dead person. Furthermore, characters that would be expected to score high on agency and experience (e.g. an adult woman) did do so. See Figure 1 below for a visualization of mean agency and experience attributions by character.

Mind Perception in Digital Agents

Figure 1

Overall mean attributions of agency and experience by character

Note: Colour indicates character category.

Mind Perception in Digital Agents

Robots

The purpose of the second set of analyses is to examine the agency and experience attributions that were made for the robots, with and without including fictional robots. In order to more precisely measure the difference between robots, we include sex (self-reported) and age. For age, we perform a median split of the data with those 38 and under in one group and those over 39 in the other.

First, an 11 (all robots) x 2 (participant sex) x 2 (participant age split) Greenhouse Geisser corrected mixed model ANOVA was conducted for agency and experience attributions. For the main effect of robot, there were significant differences for attributions of agency, $F(4.64, 552.29) = 45.34$, $MSE = 818.79$, $p < .001$, and experience, $F(2.81, 334.70) = 63.86$, $MSE = 1,151.94$, $p < .001$. There was no main effect of sex for either agency or experience (p 's > 0.18) but there was a main effect of age for both agency, $F(1, 119) = 7.52$, $MSE = 6,157.81$, $p = .007$, and experience, $F(1, 119) = 4.75$, $MSE = 1,767.14$, $p = .031$. There were also significant interactions between robot and sex for agency, $F(4.64, 552.29) = 4.10$, $MSE = 818.79$, $p = .002$, and experience $F(2.81, 334.70) = 2.75$, $MSE = 1,151.94$, $p = .046$. No other significant interactions between any factors or the combination of the 3 factors were observed (p 's > 0.05 , see Table 3 in the appendix for all results). In sum, there was significant variance between robots both for agency and experience. Sex was not a significant factor by itself, though there were sex differences amongst attributions towards some robots. Age was a significant factor with the younger generation giving higher agency and experience scores for robots.

Mind Perception in Digital Agents

Next an 8 (real robots) x 2 (participant sex) x 2 (participant age split) Greenhouse Geisser adjusted mixed model ANOVA was conducted having omitted fictional robots (FCR, R2D2, Wall-E). For the main effect of robot, there were significant differences for attributions of agency, $F(5.50, 654.47) = 8.95$, $MSE = 255.32$, $p < .001$, and experience, $F(2.99, 356.12) = 12.41$, $MSE = 119.94$, $p < .001$. There was no main effect of gender for either agency or experience (p 's > 0.18) but there was a main effect of age for both agency, $F(1, 119) = 7.66$, $MSE = 5,152.79$, $p = .007$, and experience, $F(1, 119) = 4.69$, $MSE = 936.25$, $p = .032$. There were no significant interactions between factors or the combination of the 3 factors (p 's > 0.07 , see Table 3 in the appendix for all results). In sum, there was still significant variance amongst robots in terms of their ratings for agency and experience. Sex was not a significant factor, but age was, with the younger generation giving higher scores for robots again. The results of these ANOVAs are presented, collapsing across sex in Figures 3 and 4 below. See Figure 2 for a general visualization of mean scores.

Mind Perception in Digital Agents

Figure 2*Attributions for real robots*

Mind Perception in Digital Agents

Figure 3*Mean agency ratings by robot and age split*

Note: Error bars represent standard error.

Mind Perception in Digital Agents

Figure 4*Mean experience ratings by robot and age split*

Note: Error bars represent standard error.

Discussion

People have long been interested in developing machines that are capable of emulating key aspects of the human mind, including traits such as empathy and morality. A major technical challenge, however, has been quantitatively assessing different machines in their ability to imitate complex human phenomena. Indeed, it is a widely debated topic as to whether it is possible to design robots capable of having feelings. However, recent advances have introduced powerful new research tools for investigating how people view machines in terms their ability to think, do, and feel, thus opening to the door to answering outstanding questions. Can robots compare with living beings both on agency, and critically, experience? Is there enough variation in perceived agency and experience in robots of today or tomorrow to suggest that people can perceive robots as having aspects of the human mind? Here we use the mind perception framework to specifically address these questions by comparing a large group of affective computing systems ranging from fictional characters like Wall-E and R2D2, to a real-world social humanoid robot, to other more mundane digital agents like Siri, Sphero, and a robotic vacuum.

In general, the pattern of results was clear. Our manipulation checks revealed that people understood the purpose of the task given clear benchmarks between inanimate characters like the rock in comparison with humans including an adult and child. Indeed, many of the values were similar to Gray et al. (2006) original study. For instance, a rock scored near zero and an adult scored near perfect. Moreover, many predictable relationships also emerged (again replicating Gray et al.) amongst the characters. For instance, on both agency and experience an adult scored higher than a child who scored higher than a baby who scored higher than a fetus who scored higher than a dead person.

Mind Perception in Digital Agents

Notably, our results revealed that robots, both real and fictional, were perceived as having significantly weaker qualities of mind in comparison to humans, including a child or baby. As a group and individually (including fictional robots), robots scored consistently lower on agency, and particularly, on experience (see Figure 1).

Critically however, our results demonstrated that people perceive significant variation in aspects of mind in robots both in terms of agency and experience. In terms of agency, we found differences between digital agents, especially between more limited systems like Siri and Alexa compared to more complex robots like Sophia or Atlas. This makes sense given that the latter two can move around rooms while at best, in terms of embodiment, Siri and Alexa are confined to a speaker. Intuitively, Sophia and Atlas can *do* much more than Siri and Alexa. However, it is intriguing that a more complex robot like Sophia differed so profoundly in how much experience people attributed it to compared to more limited systems. After all, it seems less intuitive to think that Sophia or Atlas can *feel* much more than Siri or Alexa can. Furthermore, many are skeptical that people can ever design digital agents that can imitate human qualities, especially aspects of experience. Yet, our results reveal that people do reliably associate more experience capabilities with certain robots, like Sophia or Atlas, in comparison to more limited agents like Siri or Alexa.

The present study indicates that people are open to perceiving qualities of mind in robots - in both agency *and* experience - and that the mind perception framework can be used as a means of assessing and comparing different robots on both their agentic and affective capabilities. Furthermore, these results were robust to the omission of fictional robots, indicating that people currently perceive these capabilities in the real robots of today. This suggests that a landscape of robots differing in their agentic and affective capabilities is not a far-off science

Mind Perception in Digital Agents

fiction situation; rather it exists today. Indeed, if all the robots had scored low on experience, or if robots only differed in agency, it would have suggested that building robots capable of emulating key aspects of the human mind would have been a most challenging exercise. Instead, these results support the notion that building robots to emulate key hallmarks of the human mind, like empathy and morality, while still a difficult if not daunting task, is a fruitful avenue to pursue.

Our study also revealed that age may be a critical factor in mind perception in robots. While sex was not a significant factor, the older generation (> 38 years) gave lower scores for robots on both agency and experience compared to their younger (<39 years) counterparts. These data suggest that there are generational differences in how open people are to perceiving both agency and experience in robots. It is our speculation that these data may be emblematic of a gradual societal change in how people perceive digital agents as these agents become increasingly more sophisticated and intertwined with everyday living. This also supports the idea that people are becoming increasingly willing to perceive a range of capacities of agency and experience in robots.

In conclusion, the present investigation discovered that there is significant variation in the mind perception of robots, not only in terms of what they can do (agency), but crucially, with regard to what they can feel (experience). This challenges the idea that developing computing systems that will be perceived as being able to do *and* feel is a distant if not impossible reality. The present study reveals that significant progress has already been made with people perceiving a range of agency and experience in fictional and real robots. Indeed, the attribution of greater mind to fictional robots with more humanized traits (e.g. Wall-E, R2D2, FCR) relative to their existing counterparts (e.g. Siri, Sophia, Roomba) suggests that, although people are relatively

Mind Perception in Digital Agents

reluctant to ascribe conscious experience to the current robots, certain innovations related to affective technology may render greater mind perception to future digital beings. Moreover, the generational difference in how people attribute agency and experience to real robots suggests that people can learn to be more accepting of minds in digital agents, highlighting the potential influence of exposure and education. Collectively these data provide substantial support for the optimistic viewpoint that digital agents can be developed further to be perceived as possessing greater, closer approximations of the whole human mind.

References

- Cai, Y. (2006). Empathic computing. In *Ambient Intelligence in Everyday Life* (pp. 67-85). Springer, Berlin, Heidelberg.
- Casler, K., Bickel, L., & Hackett, E. (2013). Separate but equal? A comparison of participants and data gathered via Amazon's MTurk, social media, and face-to-face behavioral testing. *Computers in Human Behavior*, 29(6), 2156–2160.
<https://doi.org/10.1016/j.chb.2013.05.009>
- Chalmers, D. J. (1995). Facing Up To The Hard Problem of Consciousness. *Journal of Consciousness Studies*, 2(3).
- Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On Seeing Human: A Three-Factor Theory of Anthropomorphism. *Psychological Review*, 114(4), 864–886. <https://doi.org/10.1037/0033-295X.114.4.864>
- Eyssel, F., Kuchenbrandt, D., & Bobinger, S. (2011). Effects of anticipated human-robot interaction and predictability of robot behavior on perceptions of anthropomorphism. *HRI 2011 - Proceedings of the 6th ACM/IEEE International Conference on Human-Robot Interaction*, 61–67. <https://doi.org/10.1145/1957656.1957673>
- Gazzola, V., Rizzolatti, G., Wicker, B., & Keysers, C. (2007). The anthropomorphic brain: The mirror neuron system responds to human and robotic actions. *NeuroImage*, 35(4), 1674–1684. <https://doi.org/10.1016/j.neuroimage.2007.02.003>
- Gray, H., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*,

Mind Perception in Digital Agents

315(5812), 619. <https://doi.org/10.1126/science.1134475>

Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, 315(5812), 619. <https://doi.org/10.1126/science.1134475>

Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition*, 125(1), 125–130.
<https://doi.org/10.1016/j.cognition.2012.06.007>

Hauser, D. J., & Schwarz, N. (2016). Attentive Turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods*, 48(1), 400–407. <https://doi.org/10.3758/s13428-015-0578-z>

Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American journal of psychology*, 57(2), 243-259.

Miller, G. A. (2003). The cognitive revolution: A historical perspective. *Trends in Cognitive Sciences*, 7(3), 141–144. [https://doi.org/10.1016/S1364-6613\(03\)00029-9](https://doi.org/10.1016/S1364-6613(03)00029-9)

Moore, J. W. (2016). What is the sense of agency and why does it matter? *Frontiers in Psychology*, 7(AUG), 1–9. <https://doi.org/10.3389/fpsyg.2016.01272>

Mori, M., MacDorman, K. F., & Kageki, N. (2012). The uncanny valley. *IEEE Robotics and Automation Magazine*, 19(2), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>

Saygin, A. P., Cicekli, I., & Akman, V. (2003). *Turing Test: 50 Years Later*. 23–78.
https://doi.org/10.1007/978-94-010-0105-2_2

Mind Perception in Digital Agents

Searle, J. R. (1980). Minds and brains without programs. *Mindwaves*, 3, 1–19.

Stafford, R. Q., MacDonald, B. A., Jayawardena, C., Wegner, D. M., & Broadbent, E. (2014).

Does the Robot Have a Mind? Mind Perception and Attitudes Towards Robots Predict Use of an Eldercare Robot. *International Journal of Social Robotics*, 6(1), 17–32.

<https://doi.org/10.1007/s12369-013-0186-y>

Waytz, A., Cacioppo, J., & Epley, N. (2010). Who sees human? The stability and importance of individual differences in anthropomorphism. *Perspectives on Psychological Science*, 5(3), 219–232. <https://doi.org/10.1177/1745691610369336>

Weizenbaum, J. (1966). ELIZA-A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45. <https://doi.org/10.1145/365153.365168>