Better know who you are starving with

Judging humanness in a multiplayer videogame

Stephanie Tulk

George Mason University
Fairfax, Virginia
stulk@gmu.edu

Taha Zafar

George Mason University
Fairfax, Virginia

Ryon Cumings

George Mason University
Fairfax, Virginia

Eva Wiese

George Mason University
Fairfax, Virginia

ABSTRACT

Perceived humanness affects how others behave towards artificial agents, and can be examined with imitation games, where participants interact with either a human or an artificial agent and then have to guess who they interacted with. The current experiment uses the multiplayer, socially interactive videogame *Don't Starve Together* to examine whether it is possible for participants (Player 1) to distinguish a human from an artificial agent (Player 2) based on how they make decisions in complex environments. We explore how sensitive participants are at discriminating human from AI behavior, and whether Player 2's actual and/or perceived humanness has an impact on the participants' willingness to cooperate with them. The results show that while participants were good at correctly identifying AI behavior, they misidentified the human player with a high likelihood. Cooperation during the game was higher for participants who correctly identified Player 2 (as either human or AI), and participants always preferred to cooperate with Player 2 when they thought they were human. The willingness to cooperate with Player 2 in future interactions was higher when participants misidentified Player 2, and was not impacted by the actual or perceived humanness of Player 2.

CCS CONCEPTS

• **Human-centered computing** → **Collaborative interaction**;
• **Applied computing** → **Psychology**; • **Computer systems organization** → **Robotic autonomy**;

KEYWORDS

Imitation Game, Behavioral Economics, Videogames

ACM Reference Format:

Stephanie Tulk, Ryon Cumings, Taha Zafar, and Eva Wiese. 2018. Better know who you are starving with: Judging humanness in a multiplayer videogame. In *Proceedings of ACM Woodstock conference (TechMindSociety '18)*. ACM, New York, NY, USA, Article ..., 6 pages. <https://doi.org/10.1145/3183654.3183710>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

TechMindSociety '18, April 2018, Washington, DC, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5420-2/18/04...\$15.00

<https://doi.org/10.1145/3183654.3183710>

1 INTRODUCTION

While it is easy to imagine ways in which artificial agents, such as social robots and virtual avatars, can be incorporated into our world, many have urged caution in further developing Artificial Intelligence (AI) that could someday autonomously control these agents [19]. While current technologies are not at the level of sophistication where people need to fear that their actions can threaten the whole of humanity, experts like Ray Kurzweil expect artificial agents to reach that critical threshold soon. Detecting the working of artificial agents or bots can particularly be an issue if their identity is hidden (they do not look like artificial agents), if we are only presented with the outcome of their doing (bots producing fake likes or generating false news on social media platforms), or if their identity needs to be inferred from complex behaviors (intelligent cars making decisions about driving). Furthermore, while interactions with other humans usually follow commonly accepted moral standards (lying is bad, friends need to be protected, selfishness needs to be prevented, etc.) and have naturally emerged in a cooperative fashion over the course of evolution [8], AIs are not necessarily expected to "play by the same rules" as humans or benefit from cooperative behavior. The goal of the current study is to explore whether participants can distinguish human from artificial agents merely based on their decision making behavior and whether the actual or perceived humanness of an interaction partner influences whether we choose to cooperate or compete with them in an interactive social environment. Modern multiplayer videogames are an ideal platform to examine these questions, since they allow players a wide range of different behaviors, require them to act while monitoring the other player's decisions, and allow to hide the actual identity of the other player behind a virtual character. Here, we use the game *Don't Starve Together*, where participants maneuver a complex virtual environment together with another player in order to survive, to examine the aforementioned questions in an imitation game.

1.1 Imitation Game: Who am I interacting with?

When people interact in complex decisionmaking tasks with agents of unknown identity (human or artificial), the assessment of the agent's identity needs to be coordinated with strategic decisions on how to best accomplish a shared or separate goal and social

decisions regarding how to keep the other person engaged and entertained. Whether we interact with a human or an AI potentially has a tremendous impact on whether we focus on the strategic or social component of the interaction, and whether we choose to cooperate or compete in order to accomplish the desired outcome. Decisions regarding our interaction partners' humanness also affect moral judgments, with immoral acts being more likely executed towards nonhuman agents (or dehumanized humans), and protective and prosocial acts more likely being devoted to human agents (or anthropomorphized nonhuman agents); see [2]. But how easy is it to detect humanness in the complex behaviors of others? The question of whether artificial agents can be designed in a way that makes them indistinguishable from humans was first raised by Alan Turing in the 1950s. To investigate this question, he developed the so-called Imitation Game, where humans interact with a machine that is designed to convince them that it is 'human'. Participants can ask as many questions as they like in a communicative setting, but ultimately need to make the decision whether they think that they are interacting with a human or a machine. The machine traditionally passes the Turing Test, if 30% of the participants or more misidentify it as "human" in 5 minutes of unrestricted conversation [17]. To date, only one machine is said to have passed the Turing test [18], although the victory is not universally recognized [5]. One major criticism of the traditional Turing Test is that it is more a test of human intelligence rather than machine intelligence [1], [3], and that the interaction is too unilateral, with the machine being "questioned" by the human without actual mutual interactions. To overcome these limitations, modern versions of the Imitation Game require both the human and the machine to demonstrate their intelligence on a task that poses a similar challenge for each [16]. This approach allows for greater freedom in how to conduct the test and permits that humanness can be judged based on complex behavior instead of verbal conversations [10]. But what are behavioral signs of humanness and how can we identify humans based on how they make decisions?

1.2 How do humans make decisions?

One feature associated with human decisionmaking is rational thought, or the ability to make pragmatic decisions based on available information and facts [4]. Herbert Simon proposed that people have *bounded rationality*, which in the psychological sense means that people make the most rational choice they can given a complex "inner world", including variable internal states and cognitive limitations which may interfere with their ability to make sense of the external environment [15]. When an individual's outcomes interdepend on the decisions of others, social factors like trust, fairness, or affection for others will influence the determination of the best decision given these additional constraints [4]. In order to investigate how people make complex decisions under conflict (i.e., where social and personal pressures are at odds), behavioral economics research uses twoperson interaction models adapted from Game Theory to explore how participants maximize their benefits [7]. For instance, the Prisoner's Dilemma is a commonly used social dilemma where strictly rational choices should always lead both parties to try to cheat (i.e., competing for a better outcome), which would lead to worse outcomes than mutual cooperation. However,

when people actually play the game, they are about as likely to cooperate, as they are to compete, with the presence of social constraints being the number one factor that biases people towards cooperation [12]. Adding a social component to the game has also lead to more cooperative behavior in the Ultimatum Game and even the Dictator Game, although participants have no economic incentive to cooperate since the second player has no control over the outcome and can therefore not minimize the outcome for the participant [11]. The potential of social information in biasing human decision making away from rational choices seems to depend on who we are interacting with, and in particular, whether we are interacting with a human or an artificial agent. For instance, people who were treated unfairly by a human in the Ultimatum Game were more likely to punish the other player compared to when the unfair player is a computer; they also exhibit stronger activation in brain regions associated with disgust when treated unfairly by a human versus a computer [13]. Since perceiving nonhuman agents as humanlike (i.e., anthropomorphism) can cause people to see them as entities who deserve to be treated respectfully and fairly [2], it is possible that the less humanlike people perceive their interaction partners during economic games, the less likely they are to have concern for their outcomes in these games. In line with this assumption, participants assign less emotional valence and likability to artificial agents than to humans when playing a competitive video game against them, which suggests that they were less able or willing to empathize with nonhuman agents [6]. Another possibility is that humans draw from their own experience to understand and predict the behaviors of others, and a lack of experiences with artificial agents could cause uncertainty in how to treat them and how to interact with them. In support of this notion, it has been shown that observing cheating behavior from a robot can actually cause people to like the robot (and praise it for being clever) [14], which indicates that participants originally did not think that the robot was able to show such "human" behavior and were therefore pleasantly surprised when it did. When given the opportunity to interact with artificial agents, most people use it to increase their knowledge about these agents in order to better understand them and to strategically benefit from these insights [9]. In the current study, we expand upon these findings from previous research by investigating whether (a) participants can distinguish human from artificial agents based on their decision making behavior in video games, and whether (b) the actual and/or perceived humanness of the other player affects the decision whether to cooperate or compete in current or future interactions.

2 METHODS AND MATERIALS

2.1 Participants

20 undergraduate students (mean age: 21; 9 female) participated in this study and were recruited through a psychology research participation system. Students who had prior experience with the video game were excluded from participation. Participants reported that they spend an average of 6.4 hours (SE = 1.15) on a computer per day and an average of 6.2 hours (SE = 2.02) playing video games per week. Participants were randomly assigned to either the human player condition (10 participants; 4 female) or the AI condition (10

Figure 1: A Typical Game Scene: Both players (Wilson and Woodie) are in view. The game statistics (hunger, health, sanity) are displayed in the upper right corner of the screen and the player's inventory is displayed on the bottom of the screen.

participants, 5 female). The experiment took about 45 minutes to complete.

2.2 Game

All participants (Player 1) were assigned the same the virtual character during the game (Wilson). Player 2 (human or AI) was also assigned a fixed character (Woody), which did not differ between the human player and AI player condition. The participants started the game with some items in their inventory (i.e., basic food items, an axe) and a tiki torch at the spawn point that was always burning, providing a safe place for participants to return to at nightfall (which is a potential source of harm in the game). This home base also constituted a place to which Player 2 would occasionally return to create more opportunities for interactions. In order to keep Two copies of DST were purchased from the developer via the digital distribution platform, Steam, and were modified for this experiment in order to be able to conduct the imitation game and record player data (outlined in further detail below). DST is played in an open world, where players collect resources (e.g., food items or wood for fire) and face environmental threats (e.g., animal attacks, weather changes, or physical battles) over the course of the day. Their task is to stay alive by building a fire the night, cooking food, exchanging items with other characters or collecting resources to craft tools and clothing. Three statistics (i.e., health, hunger and sanity), displayed in the right upper corner of the screen, inform the player at any time how well they are doing. Low hunger and sanity negatively affect the player's health, and if health is depleted, their character will die and the game ends. If the player manages to stay alive, the game gets more challenging over time.

All participants (Player 1) were assigned the same the virtual character during the game (Wilson). Player 2 (human or AI) was also assigned a fixed character (Woody), which did not differ between the human player and AI player condition. The participants started the game with some items in their inventory (i.e., basic food items, an axe) and a tiki torch at the spawn point that was always burning, providing a safe place for participants to return to at nightfall (which is a potential source of harm in the game). This home base also constituted a place to which Player 2 would

Figure 2: Game Characters: Wilson (left) was played by the participant; Woodie (right) was played either by another human or an AI.

occasionally return to create more opportunities for interactions. In order to keep the playtime stable across participants, the health value for their character was set to not fall below a certain minimal threshold so that their character could never die (participants were not informed about this). **Figure 1** shows a representative scene from the game containing the home base, the player stats and the two players. Since the goal of this study is to examine whether participants can infer the humanness of their interaction partner from complex decisionmaking behaviors, we disabled the chat function during the experiment. Participants could view Player 2's behavior when they were close enough for both to be displayed on the same screen, and could interact with them by engaging in physical battles or donating food or resource items to them. In situations where the participant could not view Player 2 on their screen because they were too far away, participants had access to a minimap showing Player 2's position, which allowed them to approach them at any time during the game.

2.3 Characters

Two main characters were used during the game: Wilson, who was always played by the participant (Player 1) and Woodie, who was always played by the interaction partner (human or AI; Player 2); see **Figure 2**. The human interaction partner was played by another student who was trained to suppress certain behaviors (e.g., attacking the participant without provocation), but otherwise had a lot of freedom in terms of how to play, how often to interact with the participant and which goals to accomplish. The AI interaction partner was developed to look and behave like another human player, and operated based on a behavior tree, which was meant to emulate human behavior. Some of the behavior was adapted from that of other agents already in the game and some was independently developed based on observing how other people play the game. The AI would collect various resources, fight other entities and occasionally follow the participant. When close enough to the participant, the AI would help fight any entities that were attacking the participant. However, the AI was designed to never fight with the participant and to not get uncomfortably close or collect items that were close to the participant's character to avoid appearing impolite or off putting. Both the human and the AI were modified to

be invincible and were given an axe and some food at the beginning of the game.

2.4 Apparatus

The game was played on a computer through the Steam platform, and participants were given the option of using an Xbox One controller or mouse and keyboard. A demographic, presence and enjoyment questionnaires were administered through Google Forms. The interviews were conducted verbally and transcribed by the experimenter.

2.5 Measurements

At the beginning of the experiment, participants filled out a basic demographic questionnaire, which also asked them about how many hours per day they spent on the computer and how many hours per week they spent playing video games. During the game, the distance between Player 1 and 2 was recorded whenever the participant performed some game action, and used as a measure for the participant's willingness to cooperate with Player 2 (*Present Cooperation*). After the video game, participants were asked whether they thought they interacted with a human or an AI, how confident they were in their decision (Imitation Game; data reported), what factors influenced the decision, and how skilled they thought the other entity was at playing the game (data not reported here). Participants rated how much they would like to play cooperatively or competitively with the same agent in the future (*Future Cooperation*), and were also interviewed on which strategy they would pick (i.e., cooperate or compete) and why (*Anticipated Strategy Use*).

2.6 Procedure

At the beginning of the experiment, participants signed the informed consent form, filled out the demographic questionnaire and read the instructions. Afterwards, the experimenter initialized the game, introduced the participant to their virtual character (Wilson) and allowed them to familiarize themselves with the video game environment until they reported being comfortable enough with the game to begin with the actual experiment (about 5 minutes). During the familiarization phase, participants were made aware of the most important basics of the game, such as needing to pay attention to their stats (health, hunger, sanity), how to collect and craft items, that living in this virtual world follows a day cycle and that night fall is associated with more danger. Once participants stated that they understood the instructions and the basics of the game, the actual experiment started. To begin the actual experiment, participants were instructed to turn around while the experimenter invited Player 2 to join the game. In the human condition, Player 2 was a confederate who was running the game from a different location, and was informed by the experimenter via text message to join the game. In the AI condition, a new instance of the AI was created (see Characters). Participants played the game for about 15 minutes together with Player 2, and were allowed to play however they wanted, but were aware that the main goal was to decide whether Player 2 was a human or artificial agent based on their behavior throughout the game.

2.7 Analysis

Performance during the imitation game was quantified using d' , which is a measure from signal detection theory used to assess one's sensitivity in distinguishing between two conditions; it was calculated based on hit rates (e.g., AI identified as AI) and false alarm rates (e.g., human identified as AI). The distance data (i.e., distance between the two characters whenever the participant performed any game action), and explicit ratings of how interested players were in cooperating with Player 2 during future interactions were subjected to two 2 (actual humanness: human vs. AI) $\times$ 2 (perceived humanness: human vs. AI) ANOVAs. Responses to the interview question (If you were to play the game again with this agent, would you choose to cooperate or compete, and why?) were conceptually aligned with classical approaches within the Game Theory framework. All data was analyzed using Rstudio.

3 RESULTS

3.1 Imitation game

7 out of 10 participants successfully detected the AI (hit rate: 70%, miss: 30%), and 3 out of 10 participants successfully detected the human (correct rejection (hit human): 30%, false alarm (miss human): 70%), with a d' of 0.0 (i.e., hit rate = false alarm rate). The overall accuracy was 50%. Participants' overall confidence in their decision was 61%, with no differences between correct and incorrect identifications ($p = .218$) or human versus AI identifications ($p = .863$).

3.2 Present cooperation

The descriptive statistics for the distance measures obtained during the game are shown in **Figure 3**. Participants were significantly more likely to judge the identity of Player 2 correctly if they stayed closer to them during the game ($t(2822) = 24.80, p < .001$). Distance between Player 1 and 2 was significantly related to the actual humanness ($F(1,2940) = 108.35, p < .001$), as well as perceived humanness ($F(1, 2940) = 31.64, p < .001$). The interaction between actual and perceived humanness was also significant ($F(1, 2940) = 7.248, p = .007$), with the distance between Player 1 and 2 being significantly larger when the human was misidentified as AI than in any other condition.

3.3 Future cooperation

The descriptive statistics for the interest in future cooperation measures are shown in **Figure 4**. Overall, participants who accurately detected Player 2's identity had less interest in playing cooperatively in the future ($t(17.56) = 2.89, p < .01$), and less overall interest in playing with the agent again in the future ($t(16.78) = 2.51, p = 0.023$). Willingness to cooperate in the future was not affected by actual humanness ($F(1,16) = .07, p = .790$), or perceived humanness ($F(1,16) = .01, p = .923$). The interaction between the two variables was significant ($F(1,16) = 8.60, p < .01$). The same pattern was seen for interest in playing with the agent again in the future, with no main effects, but a significant interaction ($F(1,16) = 5.45, p = .033$).

Figure 3: Distance Measures: A) Participants who correctly identified Player 2 stayed closer to them during the experiment. B) Participants stayed closer to Player 2 when they believed they were human and stayed significantly further away when they misjudged the human player as AI.

Figure 4: Figure 4. Interest in Future Cooperation: A) Participants who correctly identified Player 2 were less willing to cooperate with them in the future. B) Perceived and actual humanness did not significantly affect participants willingness to cooperate in the future, but the interaction effect was significant.

3.4 Anticipated Strategy Use

9 participants stated that they would compete (4 with the human as Player 2, 5 with the AI), and 11 participants that they would cooperate (6 with the human, 5 with the AI), if they were to interact again with Player 2 in the future. Of those 20 participants, 13 indicated that they would choose this strategy because they believed it would help them win (5 compete, 8 cooperate), 6 believed that using this strategy would be more fun (4 compete, 3 cooperate) and

1 participant did not provide a reason (1 cooperate). The provided explanations for cooperation versus competition conceptually align well with classical interactions within the Game Theory framework: the motivation to compete to win fits the Deadlock Game (i.e., goal is to win, no coordination necessary; 2 human, 5 AI), the motivation to compete for fun fits the Coordination Game (i.e., goal is to have fun, winning is not necessary; 2 human, 2 AI), and the motivation to cooperate for utility fits the Stag Hunt (i.e., goal is to win via mutual cooperation, only compete if the other does not cooperate; 7 human, 5 AI).

4 DISCUSSION

The goal of the current experiment was to examine whether complex behavior, such as decisionmaking in video games, can be used to distinguish human from artificial agent players during an imitation game. We also explored whether the actual and/or perceived humanness of the other player had an impact on the how much participants cooperated with them during the game and how much they wanted to cooperate or compete with them again in the future. In order to understand the reasons why participants prefer a competitive to a cooperative strategy for future interactions, we reclassified their named reasons within a game theoretical framework. The imitation game data shows that the AI was correctly identified as nonhuman 70% of the time (i.e., passing a typical Turing test), while the human was correctly identified only 30% of the time. The overall accuracy was no better than chance with no sensitivity (i.e., participants had no ability to distinguish signal (AI performance) from noise (human performance)); regardless of true identity, the agent is identified as AI 70% of the time) during the imitation game indicates that it was difficult for participants to reliably distinguish behaviors during the video game. The distance measures obtained during the game showed that participants who stayed closer to Player 2 during the game were better at correctly judging its identity (human vs. AI) compared to participants who stayed further away, which indicates that focused attention facilitated accurate identification. Participants also always stayed closer to Player 2 when they thought it was played by another human (i.e., perceived humanness), independent of whether it was actually the human or the AI player (i.e., actual humanness). Participants stayed farthest away from the actual human player when they judged it to be the AI, which could potentially be due to a conflict between perceived humanness and actual humanness. The willingness to stay close to the other player could indicate an interest in cooperation or at least in learning more about the player in order to understand how to treat them and how to interact with them. In terms of anticipated future interactions, the data shows that participants who correctly identified their interaction partners as either human or AI were less willing to cooperate than participants who incorrectly identified them, and less interested in playing with the agent again in the future. In contrast to their behavior during the game, neither the actual nor perceived humanness of the interaction partner predicted the use of cooperative strategies in the future. The reported results raise at least two questions that the current experiment cannot answer: First, why is the distance between players during the game greatest, when the human was incorrectly identified as AI? One

possible explanation is that participants make decisions regarding the humanness of Player 2 early on in the game, then adjust their behavior to this early prediction (i.e., stay away from the AI, and get closer to the human), and might miss critical information that would lead to an updated model. That is, participants who initially judged the human player as AI and thus stayed away from it might have missed behaviors from this agent that would have lead them to correctly change their model to "human". Participants who initially judged Player 2 as human stayed closer to them throughout the game and were less likely switch their model to 'AI'. In order to validate this assumption, followup studies need to determine if and when initial models are formed, how they change over time and whether the dynamics differ for "human" versus "AI" predictions. Second, why are participants with better identification performance less willing to cooperate with Player 2 again in the future, and why is there no preference for human versus AI partners? One possible explanation is that the participants who correctly identified the agents paid closer attention to the task at hand and were less interested in the social interaction, while participants who focused more on the social component of the game might not have had enough resources for the actual identification task. This interpretation would be in line with the observation that overall identification performance was low (i.e., task was difficult and required attentional resources), and that "cooperation" during the game was necessary for successful performance on the identification task, but was not associated with anticipated utility in future interactions. In order to validate this assumption, future studies need to examine participants' motivations when engaging in the game in more detail (i.e., succeeding at the imitation game vs. having fun playing DST vs. succeeding at DST). Future studies will address these open questions and a more complex qualitative analysis will be conducted on the remaining behavioral and interview data. We will also collect more data in order to increase the statistical power of our analyses, and incorporate neurophysiological measures in our procedure to gain a better understanding of the underlying cognitive processes. We will also develop the artificial agent further by incorporating a learning algorithm so that the AI can develop its own interaction strategies and adjust to Player 1's behavior.

REFERENCES

- [1] Daniel C. Dennett. 2004. Can machines think? In *Alan Turing: Life and Legacy of a Great Thinker*. Springer, 295–316.
- [2] Nicholas Epley, Adam Waytz, and John T. Cacioppo. 2007. On seeing human: a three-factor theory of anthropomorphism. *Psychological review* 114, 4 (2007), 864.
- [3] Robert M. French. 1990. Subcognition and the Limits of the Turing Test. *Mind* 99, 393 (1990), 53–65.
- [4] Gerd Gigerenzer and Reinhard Selten. 2002. *Bounded rationality: The adaptive toolbox*. MIT press.
- [5] Thomas Hornigold. 2017. How Close Is Turing's Dream of Conversational Machines? *Singularity Hub* (2017).
- [6] Sohye Lim and Byron Reeves. 2010. Computer agents versus avatars: Responses to interactive game characters controlled by a computer or other player. *International Journal of Human-Computer Studies* 68, 1-2 (2010), 57–68.
- [7] R. Duncan Luce and Howard Raiffa. 2012. *Games and decisions: Introduction and critical survey*. Courier Corporation.
- [8] John M. McNamara, Zoltan Barta, Lutz Fromhage, and Alasdair I. Houston. 2008. The coevolution of choosiness and cooperation. *Nature* 451, 7175 (2008), 189.
- [9] Roberta C. Ramos Mota, Daniel J. Rea, Anna Le Tran, James E. Young, Ehud Sharlin, and Mario C Sousa. 2016. Playing the 'trust game' with robots: Social strategies and experiences. In *Robot and Human Interactive Communication (RO-MAN), 2016 25th IEEE International Symposium on*. IEEE, 519–524.
- [10] Hirotaka Osawa, Kunitoshi Tobita, Yuki Kuwayama, Michita Imai, and Seiji Yamada. 2012. Behavioral turing test using two-axis actuators. In *RO-MAN, 2012 IEEE*. IEEE, 328–333.
- [11] Robert D. Rogers, Andrew P. Bayliss, Anna Szepletowska, Laura Dale, Lydia Reeder, Gloria Pizzamiglio, Karolina Czarna, Judi Wakeley, Phillip J. Cowen, and Steven P. Tipper. 2014. I want to help you, but I am not sure why: Gaze-cuing induces altruistic giving. *Journal of Experimental Psychology: General* 143, 2 (2014), 763.
- [12] David Sally. 1995. Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992. *Rationality and society* 7, 1 (1995), 58–92.
- [13] Alan G. Sanfey, James K. Rilling, Jessica A. Aronson, Leigh E. Nystrom, and Jonathan D. Cohen. 2003. The neural basis of economic decision-making in the ultimatum game. *Science* 300, 5626 (2003), 1755–1758.
- [14] Elaine Short, Justin Hart, Michelle Vu, and Brian Scassellati. 2010. No fair!! an interaction with a cheating robot. In *Human-Robot Interaction (HRI), 2010 5th ACM/IEEE International Conference on*. IEEE, 219–226.
- [15] Herbert A. Simon. 2000. Bounded rationality in social science: Today and tomorrow. *Mind & Society* 1, 1 (2000), 25–39.
- [16] Susan G. Sterrett. 2003. Turing's two tests for intelligence. In *The turing test*. Springer, 79–97.
- [17] Alan M. Turing. 2009. Computing machinery and intelligence. In *Parsing the Turing Test*. Springer, 23–65.
- [18] Kevin Warwick and Huma Shah. 2016. Can machines think? A report on Turing test experiments at the Royal Society. *Journal of experimental & Theoretical artificial Intelligence* 28, 6 (2016), 989–1007.
- [19] Ellie Zolfgharifard. 2015. Don't let AI take our jobs (or kill us): Stephen Hawking and Elon Musk sign open letter warning of a robot uprising. *Daily Mail* 1 (2015).