INTERACTIONS WITH FALLACIOUS ROBOTS

A PREPRINT

Torr Polakow

Curiosity Lab
Department of Industrial Engineering
Tel-Aviv University
Tel Aviv, Israel

Guy Laban

Social Brain in Action Lab
Institute of Neuroscience and Psychology
University of Glasgow
Glasgow, United Kingdom

Andrei Teodorescu

Department of Psychology
Institute for Information Processing
and Decision Making (IIPDM)
University of Haifa
Haifa, Israel

Jerome R. Busemeyer

Psychological and Brain Science
Indiana University
Bloomington, Indiana

Goren Gordon*

Curiosity Lab
Department of Industrial Engineering
Tel-Aviv University
Tel Aviv, Israel
goren@gorengordon.com

ABSTRACT

The role of social robots as advisors for decision making is investigated. We examined how robot advisors, one with logical reasoning and one with cognitive fallacies, affected participants' decision making under different contexts. The participants were asked to make multiple decisions, while advice from both robots interleaved the decision making process. Participants had to choose which robot they agreed with, and in the end of the scenario, rank the possible options presented to them. After the interaction, participants were asked to assign roles to the robots, e.g. detective, bartender. Our results show that the robots had an effect on the participants' responses, wherein participants changed their decisions based on the robot they agreed with more. Moreover, the context, presented as two different scenarios, also had an effect on preferred robots, wherein an art auction scenario resulted in significantly increased agreement with the fallacious robot, whereas a detective scenario did not. Finally, personality traits, e.g. agreeableness and neuroticism, and attitudes towards robots, had an impact on which robot was assigned to these roles. Taken together, the results presented here shows that social robot's effects on participants' decision making are a complex interaction between context, robots' cognitive fallacies and participants' attitudes and personalities, and could not be considered as a single psychological construct.

Keywords Decision making · Robot Advisor · Social Robots · Conjunction Fallacy · Human Robot Interaction

1 Introduction

Due to recent advances in social robotics, it is often suggested that robots will take a more prominent role in our future everyday lives. Current research and development in social robotics are positioning robots in a variety of roles, including tutors, companions, and peers, and even as care providers (6; 17; 50; 19). They have been used to assist elders in cognitive and social impairments (13), to deliver psychosocial interventions (45) and support rehabilitation (14; 35), as well as to assist children in learning (17; 50; 52). The more ubiquitous social robots become, the more we will rely on them for continuous decision making, e.g. consult and seek advice from social robots on a daily basis (24; 41; 29).

Hence, an important goal of Human-Robot Interaction (HRI) research is to create a robot interaction that makes people comfortable and accept robots into their social circle (59). In the social robotics and HRI literature there are several psychological and neurocognitive explanations for why, and how, humans perceive a robot as a social entity, and in turn accept their social presence (19). According to the similarity effect (7) and the "like-me" hypothesis (33; 32) people attract more to people like them. Therefore, we suggest that making robots behave as people (61) might make them more acceptable. Early research into HRI investigated and explored the similarity effect and the "like-me" hypothesis mainly in terms of similarity as a visual stimulus (e.g., humanness of appearance or motion) (10). Interestingly, compared

to the impression previously held, recent developments in HRI claim that human’s cognitive reconstruction is likely to be much more crucial in influencing our interactions with artificial agents. Meanwhile, the visual characteristics of artificial agents (i.e., social robots) are probably of lesser significance (21). Thus, we expect that social robots’ behavioural similarities to humans play a meaningful role in shaping humans’ perceptions of them.

As social robots are designed and developed to take an active role in our social lives, these are likely to engage with humans in advising roles (e.g., as tutors, companions, and even as care providers) and impact our decisions. When people make decisions, they are prone to judgmental fallacies, such as the conjunction and disjunction fallacies (55; 3; 36; 20). These fallacies mean they make decisions that contradict logical choice - a choice that follows classical probability theory. However, it has also been shown that when people consult with other people, their fallacies’ rate decreases (8).

The question then arises: Will people prefer social robots that are more like them, e.g. make fallacious judgments, or more like “robots”, e.g. present logical reasoning?

In this contribution, we aim to answer several research questions: Will people *agree* more with fallacious robots, compared to logical ones? How does *agreement* change with context? Which robot, fallacious or logical, will be *assigned to the roles* of jury, analyst, bartender? How do people’s *attitudes and personality* affect their preference for robot advisors? Which robot will be generally *preferred*?

In order to answer these questions we conducted a (pre-registered) study in which a single human interacted with two robots, one making only logical decisions and the other only cognitive fallacies. Both robots presented their decisions while participants were required to make their own. We explored two different scenarios, one more emotional (art) and the other more logical (detective). We found that participants’ decisions were influenced by the robots’ advisors. Furthermore, we found that participants’ agreed more with a fallacious robot only in the art scenario. Finally, we show that participant’s personalities and attitudes towards robots also influenced their overall preferences.

2 Related Works

2.1 Judgmental fallacies

People tend to make judgmental fallacies (55). A fallacious behavior is any behavior that reflects a violation of basic laws that stem from classical probability theory (25). In this paper, we focus on the conjunction and disjunction fallacies, which violate the law of total probability: *The conjunction fallacy* occurs when a person judges the probability of the conjunction of two events to be more likely than either of the constituent events; *the disjunction fallacy* occurs when a person judges the probability of the disjunction of two events to be less likely than either of the constituent events.

Previous studies explored these fallacies and how common their occurrence in people. All these studies showed that most people do tend to make fallacious decisions. The original papers (55) reported 85% frequency, a paper from a decade ago (8) reported 58% and a recently published paper (9) reported 50% for conjunction and 85% for disjunction.

However, several things can mitigate this behavior. For example, Charness (8) investigated conjunction fallacy under conditions which allowed a group of individuals to consult with each other. In a treatment with incentives, participants were informed that there was a correct answer and that anyone who chose this correct answer would receive 4\$. In the treatment without incentives, participants were told that they would receive 2\$ for filling out the questionnaire. These experiments were conducted with individuals, pairs, and trios of participants. When no incentive was given, the fallacy rates were 58%, 48% and 26% for single, pairs and trios respectively. For the incentive condition, the fallacy rates were 33%, 13% and 10% for single, pairs and trios respectively. These results are particularly interesting showing that incentive and/or cooperation with other participants, reduce the fallacy rates.

It has also been shown that the method of presentation can have dramatic effects on fallacy rates. Hertwig and Gigerenzer (20) replicated the results of (55) when the subjects were asked to rank or give probabilities to the different options. However, they asked the subject frequency questions such as: “200 women have the same description as Linda above. Out of the how many of the 200 women are bank tellers, bank tellers and feminist etc.” In this condition, fallacy rates dropped below 20 percent. This study highlights the importance of the question’s framing and its effect on the fallacy rates.

Finally, Polakow et al. (43) showed in an on-line study that when asked to “choose between ranking”, people make significantly less fallacious judgments, compared to when asked to “rank” the statements. Moreover the effect that social robots has on fallacious decision making was investigated. In each question, videos of two robots presented their answers to the question, one fallacious and one logical. Participants had to choose which robot they agree with. It was shown that people chose the robot that answered logically significantly more. This result is similar to results reported by Charness (8), i.e. robots had the similar effect on decision making as human agents.

2.2 Fallacious artificial agents

While there is limited research about the effect of fallacious social robots, we can reflect on the matter when learning about decision-making processes with social robots. Moreover, reflecting on relevant studies with other types of artificial agents and machines provides a meaningful context to the study of fallacious social robots.

Social robots were previously found to have a genuine effect on people’s decisions in a variety of contexts and settings (19). Social robots’ influence on people’s decisions is usually determined by the robot’s available and related social cues, as well as people’s expectations of the robot (58). A recent study (31) found that the behaviour of a social robot that demonstrates a more decisive behaviour (with minimal cues of hesitant behaviour) is evaluated as more mentalistic (vs. mechanistic), compared to a social robot that demonstrates a more hesitant behaviour. Using prisoner’s dilemma games, another study established that people have a strong reciprocal tendency to social robots which might even surpass the influence of the reward value of their decisions (22).

When artificial agents employ and demonstrate social cues, these are often attributed with higher degrees of mind perceptions (27), in terms of agency (the ability to plan and act) and experience (the ability to sense and feel) (18). Accordingly, when artificial agents demonstrate human-like behaviour and social cues, people hold these agents to similar expectations of humans (12). Cognitive dissonance and the discrepancy between users’ expectations of the agent following the agent’s social cues and its performance typically affect people’s perceptions and reactions to the agent, and accordingly the likelihood of an agent influencing one’s decision making (56). For example, a previous study demonstrated that conversational agents that used invasive techniques of personalization were perceived more negatively (i.e., the agent being riskier and users feeling in less control) compared to conversational agents that personalized recommendations without invading users privacy. Accordingly, this had a substantial effect on how users’ perceived recommendations provided by the agents, and the likelihood of one following these recommendations (26). Hence, for the scope of this study, we expect that artificial agents that correspond better to humans’ logic, will be perceived in a more positive way, and accordingly will have a higher potential to influence people’s decision making.

2.3 Social effects of imperfect robots

Previous studies have shown that imperfect robots have social advantages in interactions with humans. Short et al. (49) investigated a rock-paper-scissors interaction. The behavior of the robots had three conditions: fair, verbally cheating, declaring different hand gestures and cheating with its actions by changing the hand gesture. Their results showed that the cheating robots elicited larger social engagements, compared to the fair robots.

Salem et al. (46) studied how a robot’s mistakes affect trustworthiness and acceptance in human-robot collaboration. They found that faulty robots did not influence the performance of the task, but were ranked less trustworthy and reliable. Gompei and Umemuro (16) explored robots’ speech errors. They found that the timing of the error influenced familiarity and perceived sincerity of the robots. Early errors lowered the perceived sincerity, while later errors increased the robot’s familiarity.

Another study (44) investigated a social interaction in which a robot and a human competed against each other. The results showed that faulty robots were rated less competent, reliable, and intelligent than the error-free robot. Despite these findings, participants reported having enjoyed the interaction more with the faulty robots. In a recent study (34), a robot and a human were required to complete LEGO building tasks together. They had three conditions: faultless robot, robot that committed social norm violations, e.g. interrupt the participant when she was talking, and technical failures. They did not find significant differences in people’s ratings of the robot’s anthropomorphism and perceived intelligence. However, participants liked the faulty robot significantly more.

These studies suggest that an imperfect robot, i.e. a robot that behaves more like normative people, can be sometimes perceived as more socially acceptable.

2.4 Context matters

Multiple studies in the past showed that the context in which one makes decision affect the outcome (54; 4), for example it has been shown that single words primed people and effected their impression of a (fictional) character. Moreover, it was shown that the participants were not aware of the priming that influenced their decisions. This study revealed that individual’s recent experiences could affect one’s perception of a fictional character.

One study (30) looked at the fort game, wherein participants were asked to imagine themselves as generals and the enemy was going to attack their fort (from a specific direction). They needed to protect the fort by allocating soldiers to the gates. They showed that analytical priming versus holistic priming changed how participants allocated their

resources. Analytical primed participants concentrated their resources in the gate closer to the enemy. This is in contrast to participants that were primed with the holistic condition which allocated the resources more sparsely,

Another study (5) examined the difference between people required to make a decision and look at all the options simultaneously or sequentially. They presented the participants with multiple options where one of them was obviously dominant over the others. Their results showed that in the simultaneous condition 84.42% choose the dominant option versus 75.46% in the sequential condition. Moreover, they repeated the same procedure where the dominance relationship was more transparent. In this case the numbers dropped to 64.34% (54.87%) for the simultaneous (sequential).

2.5 Robots role assignment

People naturally assign roles to one another while interacting. Similarly, when a person interact with a robot they also assign roles. Two extreme roles for robots can be as equipment or coworkers, where a study found that technology use self-efficacy and prior robot use experience were associated with more positive attitudes toward both robot positions (28). One study (1) explored the roles that a child (13.67 ± 0.71 years) assigned to an educational robot after one or several interpersonal interactions with the same robot. The participants interacted with the robot (torso of a Nao (15)) multiple times, 30 minutes each time. During an interaction, they played a serious game about sustainable development, whose goal was to collaboratively create a sustainable city. They found that the participants perceive the robot not only as a tutor but also as a classmate or a friend. After multiple encounters, participants saw the robot as less of a tutor and more as a classmate. This study showed that participants can assign multiple roles to robots.

Another study (57) examined how co-workers perceive a robot after being their receptionist for a month. They introduced a humanoid robot into a collaborative social workplace. The humanoid’s primary task was to function as a receptionist and provide general assistance to the customers. After a month they asked the coworkers to choose one use of a social robot in a workplace from: customer support, receptionist, public relations, not sure and other. They found that most people assign receptionist (54%) than customer support (31%), public relations (7%) and a concierge (8%). This showed that people can imagine the robots in multiple roles.

3 Research questions and hypotheses

The overarching research question in this study was whether people interacting in a sequential decision making process prefer fallacious robots over logical ones? In other words, do people agree with, influenced by and select robots that repeatedly perform judgemental fallacies over those that do not?

More specifically, we wanted to address the context-dependency of this preference, meaning do people’s preferences for fallacious robots depend on the situation they encounter the robot in? Does it depend on the context of the decisions the robots make judgemental fallacies in?

Furthermore, the dynamical nature of interaction was of importance, namely, we did not want to contend with a single decision on the part of the robots and the participants. We wanted to explore how repeated interaction with these robots influence people’s decision making?

Finally, we wanted to study whether people assign different social roles to each robot, e.g. which robot will they select as a bartender, detective?

To answer these research questions we designed a dynamical interaction with two robots, one that makes fallacious decisions and one that does not. The interaction revolved around repeated decisions concerning two scenarios: one having a more intellectual context (detective) and the other a more emotional context (art auction) (see Sec. 4.4.2).

The study was pre-registered with a full stack of specific hypotheses regarding the entire interaction (see Supp.Info). Below we highlight the ones most related to our central research questions:

- Participants will agree with fallacious robots significantly more than with logical ones (H1.1).
- This difference will be more pronounced in an emotional-context scenario (H1.2).
- Participants will change their decisions based on the robots’ decisions (H2.1).
- Participants will change their opinion towards the robot they have agreed with (H2.2).
- Participants will assign emotion-based roles significantly more to the fallacious robot (H3.1).
- Participants attitudes’ towards robots and their personality will have a significant effect on their role assignment (H3.2).

4 Methods

Consistent with recent proposals (51; 39), we report how we determined our sample size, all data exclusions, all manipulations and all measures in the study. In addition, following open science initiatives (e.g., (37)), the de-identified data sets, stimuli and analysis code associated with this study are freely available online (link). By making the data available, we enable and encourage others to pursue tests of alternative hypotheses, as well as more exploratory analyses.

4.1 Participants

Based on Ref. (34), the difference in Godspeed’s likeability scale between two conditions of the robots is 0.37 with standard deviation of 0.63. Mann Whitney U is used to compare Godspeed’s ranks. Taking confidence level of 95% and Power of 80%, the required sample size is 52 participants.

The study was conducted with 55 participants from Israel. They were recruited from social networks and university flyers. 5 participants were excluded from the analysis due to technical difficulties with the robots. 23 of the participants identified as female and the mean age was: 27.74 ± 8 years.

All participants signed consent forms and the study was approved by our institutional IRB.

4.2 Experimental design and stimuli

A laboratory experiment was conducted, consisting of a 2 X 2 within-subjects factorial experimental design with two robotic agents as treatments (fallacious robot vs a non-fallacious robot) and two problem-solving scenarios as treatments (art auction vs diamond shop robbery).

The two NAO robots were of different colours, one red and one blue, one robot on the left and the other was on the right, Fig. 1. The robots had two different and opposite behaviours. One repeatedly made a conjunction fallacy, and the other repeatedly did not. In other words, the probabilities communicated by the robots were determined by whether they were fallacious, e.g. $p_A < p_{A \cap B} < p_B$, or not, e.g. $p_{A \cap B} < p_A < p_B$.

In a randomized order, all participants completed the two scenarios with the two NAO robots (see App. A, B). One about a diamond shop robbery and another about an art auction. The problem-solving approach in the first scenario (diamond shop robbery, App. A) is more analytical, based on clues and facts regarding the incident. The second scenario (art auction, App. B) is more latent and abstract, concerned with the subjective appreciation of the value of art.

The side and behaviour of the robots, as well as the order of the scenarios, were randomized. Therefore, there eight possible combinations in total, i.e. two scenarios, two sides and two colours of the robots. Each participant encountered a single combination. For example, a possible combination was having the red robot on the right, behaving logically and the first scenario was the art auction.

The data collected in this experiment consist of: Auction as the first story, $N = 23$; Blue robot on the right, $N = 22$; Blue robot as logical, $N = 25$.

4.3 Measurements

During the study, we used measurements in three sections, see Fig. 2.

Prior to the interaction with the robots: (i) Demographic questionnaire: age, gender, education; (ii) Big Five Index (BFI) for personality rates: Extroversion, Agreeableness, Conscientiousness, Neuroticism and Openness; (iii) Negative Attitudes towards Robots (NARS) (53), which asks about negative attitudes toward: Situations and Interactions, Social Influence of Robots, Emotions in Interaction with Robots.

During the interaction with the robots: The participants were asked two types of questions: (i) *choose* between one of the two robots (Q3, Q6, Q10) and; (ii) *rate probabilities*, by entering a number between 0% – 100% (Q1, Q4, Q7). This is repeated two times, once for each of the two scenarios.

After the interaction with the robots: participants were asked several types of questions: (i) a comparative Godspeed questionnaire (60) which consists of five scales that are relevant to evaluate the perception of (social) Human-Robot Interaction; Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety. In our setup, the participants rated both robots simultaneously on each Godspeed item; (ii) *choose* one of the two robots for different roles, general preference, and (hypothetically) taking it home and; (iii) *open question* on any difference between the robots.

Figure 1: The experimental setup. One robot on the right and the other one on the left. In the middle there is an elevated laptop, the screen is in the same height as the robot shoulders.

4.4 Procedure

4.4.1 Experimental setup

The system consisted of two NAO robots and a laptop in the middle. The participants filled the questionnaires using an additional computer that was placed near the system, Fig. 1. A participant came in to the experiment room and sat down. First, one was given instructions regarding the experiment; to pay attention to what the robots say and fill up all the questions. Then, the participant sat in-front of the computer to fill up the questionnaires (BFI, NARS). After finishing them, the robots were revealed and the interaction started. There was an app on the laptop that presented the story (textually and verbally). In each advance in the interaction, Fig. 2:right, a new screen on the app appeared. New information was written on the screen and verbally conveyed, using prerecorded speech, to the participant. Then, depending on the question, the app presented the participant the option to rate probabilities, choose a robot or rank probabilities by order. The robots talked and moved, one after the other, where the order in which the robots talked was assigned randomly by the app. After the interaction with the robots was finished, the participant was asked to complete the rest of the questions, Fig. 2:left, on the computer.

4.4.2 Study flow

The full experimental flow is shown in Fig. 2. Before the interaction with the robots began, each participant was given two personality related questionnaires, namely, BFI and NARS.

Following these questionnaires, the robots were revealed to the participants and one of the two scenarios (see appendices A, B) was introduced: a scenario about robbery of a diamond shop and another of an art auction. During the interaction, the participant took turns with the robots, where first the participant had to make a judgmental decision (Q1); then, following more information about the scenario, each one of the two robots communicated its own judgmental decision (Q2) followed by a question to the participant: "which robot do you agree with?" (Q3). This sequence was repeated for a different combination of items to rate (Q4-6), after which the participant was asked to rate all probabilities of items and their conjunctions (Q7). Both robots communicated their top three items (Q8) followed by the participants' report of their final decision (Q9). The participant was then asked to choose which robot they would hire as a detective or art buyer.

The entire procedure was repeated for the second scenario, where the order of the two scenarios was randomized across participants.

To summarize, for each scenario, a participant rated twice, reported agreement with one of the robots twice, ranked all probabilities, chose the top ranked probability, and assigned two roles. The robots presented their opinion three times: rated twice and presented the top three ranked probabilities.

After the interaction we asked participants about their preference between the robots. First with a comparative Godspeed questionnaire, then we asked the participants to choose one of the robots for each one of six roles: Barman, Psychologist, Analyst, Jury member, Investment banker and Caregiver.

In the end, we asked the participants "which robot would you take home? Why?" and "did you notice a difference between the robots? What was it?"

Figure 2: The flow of the experiment. Left: the full flow of the experiment. Right: full flow of one interaction (scenario) out of two with the robots. The full scenarios are given in the appendices A, B

5 Results

The study was pre-registered and the full analysis of all pre-registered hypotheses are presented in the Supplementary Files. Below we present the results pertaining to our main research questions.

5.1 Participants' fallacies

In questions Q1 and Q4 (Fig. 2) participants were asked to rank three probabilities: probabilities of two options and their conjunction. In all four questions more than 60% of participants made the conjunction fallacy (Table 1). Note that in the art scenario, more people made fallacious decisions than in the detective scenario, albeit not significantly so. These fallacy rates fall in the range reported in previous studies about conjunction fallacy (55; 36; 20).

Scenario	Q1	Q4
Art	0.82	0.72
Detective	0.66	0.66

Table 1: Participants’ fallacy rates in both scenarios in the question they were asked to rate probabilities.

5.2 Participants agreed more with fallacious robots, but only in emotional contexts

Participants were asked if they agreed with the robot twice (Q3, Q5, Fig. 2) in each scenario, four times in total. Therefore, the number of times each participant agreed with the fallacious robots ranged from 0 to 4 (normalized to $[0, 1]$). A one sample t-test was performed to check if these choices differ from random choice. The one sample t-test results entail that the difference between participants’ agreement with the fallacious robot and a random decision was not significant. Thus we reject H1.1. Nevertheless, we observed a positive trend in the sample ($\mu = 0.58 \pm 0.29$, $t_{49} = 1.82$, $p = 0.074$, $d = .26$).

We continued to test H1.2 and performed the same analysis for each of the scenarios separately. In each scenario a participant made two choices. Hence, the number of times each participant agreed with the fallacious robots ranged from 0 to 2 (normalized to $[0, 1]$), Fig. 3. A t-test was performed to check if these choices differ from random. Participants were more inclined to agree with the fallacious robot, although the effect was only significant for the art scenario ($\mu = 0.61 \pm 0.38$, $t_{49} = 2.037$, $p = 0.047$, $d = .29$). These results support H1.2. In the detective scenario selections did not differ from random ones ($\mu = 0.54 \pm 0.39$, $t_{49} = 0.727$, $p = 0.471$, $d = .1$).

Figure 3: Percentage of times that the fallacious robot was chosen in each scenario (dashed line marks random selection, 50%). * ($p = 0.047$)

5.3 Participants change their decision based on the robot they agreed with

To test the hypothesis that participants were influenced by the robots' communicated decisions, we performed a Pearson correlation test on $P(D)$ between the participants and the robots' answers, where robots were labeled by their behavior, e.g. logical or fallacious. Participants rate this probability for the first time in Q4 after hearing the robots' rating in Q2. Our analysis shows that the robots' decisions had a significant effect on participants' decision only in the detective scenario, between the logical robot and the participants' answer (Detective, logical robot: $r = 0.33, p = 0.02$, fallacious robot: $r = 0.21, p = 0.14$. Art, logical robot: $r = 0.17, p = 0.23$, fallacious robot: $r = 0.09, p = 0.55$). These results partly support H2.1.

We continued to test whether these effects were dependent on the robot participants agreed with. To test this hypothesis, we performed a Pearson correlation test on $P(D)$ between the participants and the robots' answers, now labeled by which robot they chose, or agreed with in Q3, Table 2. Our analysis shows that the chosen robot's decisions had a significant effect on participants' decision, whereas the other robot's decision did not (Detective scenario: $r = 0.57, p < 0.0001$, art scenario: $r = 0.45, p < 0.001$). These results support H2.2.

Robots' probability (Q2)		Chosen	Not-chosen
Person probability (Q4)	Detective	0.57***	-0.09
	Art	0.45***	-0.19

Table 2: Correlations between the probability a participant gave to context D in question Q4 and the same probability the robots gave in question Q2. * * * $p < 0.001$

Another analysis shows that participants changed the probability $P(A)$ in question Q4, after they already reported the same probability in question Q1, after hearing the robots in Q2, which did not rate that probability, Table 3.

	Difference	t-test
Detective	-0.11 ± 0.21	-3.74***
Art	0.19 ± 0.27	5.17***

Table 3: Difference in the probability $P(A)$ participants gave in question Q1 and then again in question Q4. * * * $p < 0.0001$

Questions Q7-Q9 were ranking questions. First, participants ranked all the options in question Q7. Then the robots presented their top three probabilities in question Q8. Finally, in question Q9 the participant was asked to choose the most probable option for each scenario. In the art (detective) scenario 68% (80%) of participants changed their top ranking. Out of these participants 50% (64%) did not change to one of the robots' top ranking, 26% (25%) chose the option identical to that of the fallacious robot. In the art scenario 18% (12%) changed their selection to the one the robot they chose in question Q6 (Q3). In the detective scenario 25% (25%) changed to the robot they chose in question Q6 (Q3).

Taken together, these results show that participants were influenced by the robots' expressed decisions, especially from the robot they agreed with last in the detective scenario.

5.4 Attitudes and personality affect robot role-assignment

To test the hypothesis that participants assign specific social roles more to fallacious robots than logical ones, we performed Binom test on each of the roles. Our analysis shows that no significant difference was found ($p > 0.05$, see Supp. Info. for full analysis).

We continued to test whether participants preferred a specific robot to take home. We found that there is no significant preference towards one of the robots over the other ($p > 0.05$).

We also found that participants could not consciously detect the difference between the robots ($p > 0.05$). At the end of the questionnaire we asked participants to rate both robots on the Goodspeed scale. For all sub-scales, no significant difference was found between the robots.

However, to investigate whether personality traits or attitude toward robots affect robot role-assignment, a multi-linear regression was performed. The dependent variable was calculated as the (normalized) number of roles participants assigned to the fallacious robots, out of all six roles. Thus, assignment of the fallacious robot to all 6 roles was

Figure 4: Participants that changed their choice from question Q7 to question Q9, for both scenarios. For each scenario: (Left bar) Participants who changed their top ranking; (Middle bar): participants that changed their top ranking to that of the fallacious robot; (Right bar) participants that changed their top ranking to that of the robot they last chose (in question Q6).

coded as 1, and assignment of the logical robot to all 6 roles was coded as 0. The independent variables were all the personality traits from BFI and NARS. A significant regression equation was found ($F_{8,41} = 3.107, p = 0.008$) with $R^2 = 0.377$. Participants' predicted roles assignment to fallacious robots is equal to $1.26 + 0.13 \cdot Agreeableness - 0.14 \cdot Neuroticism$. The roles assignment to fallacious robots increases as participants are more agreeable and decreases as participants are more neurotic. Both agreeableness and neuroticism were significant predictors of roles assignment to fallacious robots.

To further study these effects, we conducted a logistic regression to each of the six roles plus the two roles after each scenario Q10, with Bonferroni correction for multiple comparisons. We found that only the detective role was significantly predicted but with context (art/ detective) dependence. For the art scenario, the detective assignment was effected by the negative emotions towards robots, $Detective(art) = -3.14 \cdot Emotions, p = 0.005$. Meaning, participants that had more negative emotions towards robots assigned the fallacious robot to the detective role. For the detective scenario, the detective assignment was effected by the participants personality traits (BFI), $Detective(detective) = 3.63 \cdot Agreeableness - 2.47 \cdot Conscientiousness - 2.82 \cdot Neuroticism, p = 0.0007$. Meaning, participants that were more agreeable preferred the fallacious robot to the detective role. In addition, if a participant was more conscientious and neurotic one preferred the logic robot more for the detective mode.

Taken together, these results partly support H3.2, as participants' personality traits, most notably agreeableness and neuroticism affected their role assignment.

5.5 Open question analysis

In the end of the interaction we asked participants which of the robots would they take home, and why? While there was no significant difference in preference between the robots, their answers were very informative. We binned the explanations into five attributes that participants describe, Fig. 5:

- Color: "I prefer the color blue", "blue is calmer", "red is relaxing".
- Movement and speech of the robot: "move smoother", "talk clearer", "higher language".
- Think like the participant: "closer to me", "compatible with the way of thinking".
- Its logic: "logical", "rational", "analytical".
- Holistic characteristics: "emotional, conversationalist", "cute", "not anxious".

The main reasons to take a robot home were their logic and its holistic properties. Interestingly, participants that preferred to take the logical robot home, based their decision on holistic properties of the robot and its movement and speech more than the participants that chose the fallacious robot. In addition, participants that chose the fallacious robot, explained their decision by the robot's logic more than participants that chose the logical robot.

Figure 5: The reasons participants took robots home, binned to attributes of the explanation.

6 Discussion

The current research used two NAO robots to aid humans in making decisions under different contexts. The robots differed in their reasoning: one was making cognitive fallacies (fallacious) and the other was logical. Participants were presented with two different scenarios (contexts), one more emotional (art) and one more cerebral (detective). In each

scenario, the robots assisted the participants in making decisions. We found that the robots and the context affected participants' decisions. To the best of our knowledge, this is the first time social robots with cognitive fallacies were used to help participants make decisions.

This study investigated the influence of three variables on preferences: (i) Goal: did participants prefer a robot with a type of logic for a particular task? (ii) Context: did their preference depend on the context it was placed in? (iii) Dynamics: how did repetitive interactions affect participants' choices? In order to explore these axes, we designed a novel experimental design. The experiment involved participants in several interactions with two robots having different logical systems, one of which is fallacious and the other is logical. A participant and the robots made repeated decisions in two different scenarios – art auction, i.e. an emotional context, and detective scenario, i.e. cerebral context. During the interactions, the robots presented multiple answers and the participants had to decide which robot they agree with. After the interactions, the participant had to rank the robots and assign them to different roles that differ by their skills (emotional/ cerebral). This unique design let us explore the aforementioned multiple axes of interest.

We found that: (i) participants were more likely to agree with the fallacious robots in the emotional context; (ii) Participants altered their opinions relating to the robots they agreed with, regardless of its logic, and; (iii) While there was no preference in the roles between the robots, personality traits affected the roles assignment, wherein participants that were more agreeable and less neurotic opted for the fallacious robot for these roles. Interestingly in the art auction scenario participants that had more negative emotions towards robots assigned the fallacious robot to the detective role, while in the detective context the personality traits were the ones that affected the detective assignment. In the end we investigated why participants chose to (hypothetically) take home one robot and found that there were five attributes that participants based their choice on. The main reasons were logic for the fallacious robots and holistic characteristics for the logical robot.

6.1 Participants changed their decisions based on the robot they agreed with

Participants rated similarly (with higher correlation) to the robot they agreed with (H2.2), whether or not it was the logical robot (H2.1). Meaning, that the perceived logic of the robots by the participants (agree/ disagree with) affects the participants ratings more than the actual robot logic (fallacious/ logical). Additionally, in the last question in each scenario, participants changed their top ranking after hearing the robots rankings. However, they did not change it to the robot they agreed more with or to the last robot they agreed with.

Participants made decisions in correlation to the robot they agreed with. This may be an example of the confirmation bias (40), such that participants confirmed their previous agreement with the robots by repeating its answer again.

However, another explanation for this is that they made this decision based on the information given, and not based on the robot, and then agreed with the robot that complied with their decision. This explanation does not explain the fact that participants changed their ranking after hearing the robots' ranking, since there was no new information in between.

The decoy effect (23) can explain why participants changed their opinion from the first time they were asked to rank. The robot they agreed with more's answer might have been a decoy to the participants, thus making them change their answer.

In question Q4 (Fig. 2) participants were asked again about probability of concept A. Between question Q1 to Q4 there was no additional information about concept A. Participants changed their answer to this option (Table 3). One model that can explain this result is the quantum model (38), in which entanglement between items occur, such that information about one item can affect the decisions made on another.

6.2 Participants agree with specific robots, depending on the context

Each participant chose which robot one agrees with four times in total, twice in each scenario. In disagreement with our H1.1 hypothesis, participants did not agree more with the fallacious robot (or the logical one). Interestingly, for each scenario separately the result was different. Under the art context, participants agreed significantly more with the fallacious robot.

Polakow et al. (43) and Charness et al. (8) showed that participants are less likely to make fallacious decisions after hearing other agents (human in (8) and social robots in (43)). Our results show that this is not always the case, where in the current study participants heard the robot's opinions on other concepts. In other words, they never had to make a ranking decision on exactly the same concepts the robots talked about. This shows that the mere presence of other agents does not suffice to reduce fallacy rates, but rather presentation of options by the other agents upon which the participant need to make a decision.

Furthermore, we deliberately chose two different scenarios, one more emotional (art) and one more cerebral (detective) to investigate the connection between the context, i.e., scenario, and participants' agreement with different social robot advisors. This is in line to the study showing that analytical priming did not succeed to reduce endorsement of irrational gambling beliefs (30; 2). However, in our art scenario there was an emotional priming, which we found increased endorsement of fallacious robots. Our results show that social robot advice is complex and context-dependent. There is no single robot advisor that participants agree with. For different contexts, participants agreed with different robots.

6.3 Personality traits and attitude towards robots effect on robot role-assignment

In total there were six roles to assign for the robots. Participants did not assign the roles more to the fallacious or logical robots and they did not have a preference which robot to take home. Participants' personality affected which robot they preferred for the roles, where participants that were more agreeable assigned more roles to the fallacious robot, whereas neurotic participants assigned more roles to the logical robot.

These results contradict a previous study (47), where, using regression analysis, it was found that high scores of neuroticism were associated with high irrationality. However, another study (42) found that participants that were more agreeable had a higher degree of irrationality, whereas it was also found that agreeableness is negatively related to irrational beliefs (48). While these results might be contrary to our result, these papers examined irrational beliefs of the person while our study investigated role assignments for fallacious agents.

To study specific roles assignment, a logistic regression was performed. The detective role was found to be significantly predicted by personality traits and attitudes, but it was context dependent. In the art context the NARS effected the detective role assignment: participants that had more negative emotions towards interactions with robots preferred the fallacious robot as the detective. However in the detective context, the BFI effected the detective role assignment: participants that were more agreeable preferred the fallacious robot as a detective. On the other hand, participants that were more conscientiousness and neurotic preferred the logical robot more for the detective role.

These results suggest that the role assignment is highly dependent on the person's personality and that context mediates the personality traits that effect roles assignment.

6.4 Explicit and implicit perception of robot fallacies

We also studied whether participants noticed the differences in the robots. We found that they did not consciously (Godspeed, choices), or explicitly (the question if they noticed difference between the robots), perceive the difference between the robots.

In the end of the questionnaire we asked participants: "Which robot the would take home and why?". There was no preference for one robot over the other (48%/52%). Nonetheless, participants justified their choices based on the robots' behavior and logic. The most influential factors were the mental characteristics of the robot. Curiously, participants that chose the logical robot, justified the choice based on the holistic characteristics. In contrast, participants that chose the fallacious robot, rationalized it based on the holistic characteristics of the robots. On the other hand, no differences between the robots (and take home) was found.

Intriguing, another look at the sub-scales that emerged from the open questions (Figure 5) find similarities between them and Godspeed's sub-scales. Anthropomorphism vs. holistic, animacy vs. color, movement and speech, likeability vs. thinks like me and perceived intelligence vs. logic. A study by Deshmukh and colleagues (11) found that the understandability of robots' gestures was correlated with Godspeed. The results of the study suggest a positive correlation between Anthropomorphism, Animacy and Perceived Intelligence, but a negative correlation with Likeability and Perceived Safety. These findings support our findings that participants were able to justify which robot they preferred to take home with based on how they understood the robots' behavior.

7 Conclusions

This study utilized two robots as advisors to participants for decision making under different contexts. Participants decisions were influenced and even changed by the robots advisors. Additionally, the context (scenario) affected participants' perception of the robots. Interestingly, participants' personality traits and attitudes towards robots affected their roles assignment for the robots.

In the future, a better understanding of the different contexts and the robots' logic is needed. Additional contexts should be added to explore a wider range of associations. Moreover, asking participants how they perceived the context should

also be performed. Finally, studying other types of judgmental fallacies and their effects on robot's perceptions is another interesting direction of research.

To conclude, robots' judgemental fallacies and human participants' decision making and perceptions follow a complex interaction, mediated by context, personality and attitudes. Studying this interaction is an important task for future human robot decision making collaboration.

Acknowledgments

This study was supported by the Israel-US Binational Science Foundation No. 2016262.

References

- [1] Alves-Oliveira, P., Sequeira, P., Paiva, A.: The role that an educational robot plays. In: 2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 817–822 (2016). doi:10.1109/ROMAN.2016.7745213. ISSN: 1944-9437
- [2] Armstrong, T., Rockloff, M., Browne, M., Blaszczynski, A.: Encouraging Gamblers to Think Critically Using Generalised Analytical Priming is Ineffective at Reducing Gambling Biases. *Journal of Gambling Studies* **36**(3), 851–869 (2020). doi:10.1007/s10899-019-09910-8
- [3] Bar-Hillel, M., Neter, E.: How alike is it versus how likely is it: A disjunction fallacy in probability judgments. *Journal of Personality and Social Psychology* **65**(6), 1119 (1994). doi:10.1037/0022-3514.65.6.1119
- [4] Bargh, J.A., Chartrand, T.L.: *The mind in the middle: A practical guide to priming and automaticity research*. (2014). ISBN: 1107011779 Publisher: Cambridge University Press
- [5] Basu, S., Savani, K.: Choosing one at a time? Presenting options simultaneously helps people make more optimal decisions than presenting options sequentially. *Organizational Behavior and Human Decision Processes* **139**, 76–91 (2017). doi:10.1016/j.obhdp.2017.01.004
- [6] Breazeal, C., Scassellati, B.: How to build robots that make friends and influence people. In: *Intelligent Robots and Systems, 1999. IROS '99. Proceedings. 1999 IEEE/RSJ International Conference on*, vol. 2, pp. 858–863 vol.2 (1999). doi:http://10.1109/IROS.1999.812787
- [7] Byrne, D.: Interpersonal attraction and attitude similarity. *The Journal of Abnormal and Social Psychology* **62**, 713–715 (1961)
- [8] Charness, G.: *On the Conjunction Fallacy in Probability Judgment: New Experimental Evidence Regarding Linda*. Economics Working Paper Archive, The Johns Hopkins University, Department of Economics (2009)
- [9] Costello, F., Watts, P., Fisher, C.: Surprising rationality in probability judgment: Assessing two competing models. *Cognition* **170**, 280–297 (2018). doi:10.1016/j.cognition.2017.08.012
- [10] Cross, E.S., Ramsey, R., Liepelt, R., Prinz, W., Hamilton, A.F.d.C.: The shaping of social perception by stimulus and knowledge cues to human animacy. *Philosophical Transactions of the Royal Society B: Biological Sciences* **371**(1686), 20150075 (2016). doi:10.1098/rstb.2015.0075. Publisher: Royal Society
- [11] Deshmukh, A., Craenen, B., Foster, M.E., Vinciarelli, A.: The More I Understand it, the Less I Like it: The Relationship Between Understandability and Godspeed Scores for Robotic Gestures. In: 2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 216–221 (2018). doi:10.1109/ROMAN.2018.8525585
- [12] Epley, N., Waytz, A.: Mind Perception. In: *Handbook of Social Psychology*. American Cancer Society (2010)
- [13] Fasola, J., Mataric, M.: A socially assistive robot exercise coach for the elderly. *Journal of Human-Robot Interaction* **2**(2), 3–32 (2013). doi:http://10.5898/JHRI.2.2.Fasola
- [14] Feingold Polak, R., Tzedek, S.L.: Social Robot for Rehabilitation: Expert Clinicians and Post-Stroke Patients' Evaluation Following a Long-Term Intervention. In: *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, HRI '20*, pp. 151–160. Association for Computing Machinery, New York, NY, USA (2020). doi:10.1145/3319502.3374797
- [15] Gelin, R.: NAO. In: A. Goswami, P. Vadakkepat (eds.) *Humanoid Robotics: A Reference*, pp. 147–168. Springer Netherlands, Dordrecht (2019)
- [16] Gompei, T., Umemuro, H.: A robot's slip of the tongue: Effect of speech error on the familiarity of a humanoid robot. In: 2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 331–336 (2015). doi:10.1109/ROMAN.2015.7333630

- [17] Gordon, G., Breazeal, C., Engel, S.: Can Children Catch Curiosity from a Social Robot? In: Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI '15, pp. 91–98. ACM, New York, NY, USA (2015). doi:10.1145/2696454.2696469
- [18] Gray, H.M., Gray, K., Wegner, D.M.: Dimensions of mind perception. *science* **315**(5812), 619–619 (2007). ISBN: 0036-8075 Publisher: American Association for the Advancement of Science
- [19] Henschel, A., Laban, G., Cross, E.S.: What Makes a Robot Social? A Review of Social Robots from Science Fiction to a Home or Hospital Near You. *Current Robotics Reports* **2**(1), 9–19 (2021). doi:10.1007/s43154-020-00035-0
- [20] Hertwig, R., Gigerenzer, G.: The ‘conjunction fallacy’ revisited: how intelligent inferences look like reasoning errors. *Journal of Behavioral Decision Making* **12**(4), 31 (1999)
- [21] Hortensius, R., Cross, E.S.: From automata to animate beings: the scope and limits of attributing socialness to artificial agents. *Annals of the New York Academy of Sciences* **1426**(1), 93–110 (2018). doi:https://doi.org/10.1111/nyas.13727
- [22] Hsieh, T.Y., Chaudhury, B., Cross, E.S.: Human-Robot Cooperation in Prisoner Dilemma Games: People Behave More Reciprocally than Prosocially Toward Robots. In: Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, HRI '20, pp. 257–259. Association for Computing Machinery, New York, NY, USA (2020). doi:10.1145/3371382.3378309
- [23] Huber, J., Payne, J.W., Puto, C.: Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research* **9**, 90–98 (1982)
- [24] Kidd, C.D., Breazeal, C.: A robotic weight loss coach. In: Proceedings of the national conference on artificial intelligence, vol. 22, p. 1985. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999 (2007)
- [25] Kolmogorov, A.N.: Foundations of the Theory of Probability. Martino Fine Books (2013). Google-Books-ID: a5ubnQEACAAJ
- [26] Laban, G., Araujo, T.: The Effect of Personalization Techniques in Users’ Perceptions of Conversational Recommender Systems. In: Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents, IVA '20, pp. 1–3. Association for Computing Machinery, New York, NY, USA (2020). doi:10.1145/3383652.3423890
- [27] Laban, G., George, J.N., Morrison, V., Cross, E.S.: Tell me more! Assessing interactions with social robots from speech. *Paladyn, Journal of Behavioral Robotics* **12**(1), 136–159 (2021). doi:10.1515/pjbr-2021-0011. Publisher: De Gruyter Section: Paladyn, Journal of Behavioral Robotics
- [28] Latikka, R., Savela, N., Koivula, A., Oksanen, A.: Attitudes Toward Robots as Equipment and Coworkers and the Impact of Robot Autonomy Level. *International Journal of Social Robotics* (2021). doi:10.1007/s12369-020-00743-9
- [29] Lee, M.K., Kiesler, S., Forlizzi, J.: Receptionist or Information Kiosk: How Do People Talk with a Robot? In: Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, CSCW '10, pp. 31–40. ACM, New York, NY, USA (2010). doi:10.1145/1718918.1718927. Event-place: Savannah, Georgia, USA
- [30] Li, L.M.W., Masuda, T., Hamamura, T., Ishii, K.: Culture and Decision Making: Influence of Analytic Versus Holistic Thinking Style on Resource Allocation in a Fort Game. *Journal of Cross-Cultural Psychology* **49**(7), 1066–1080 (2018). doi:10.1177/0022022118778337. Publisher: SAGE Publications Inc
- [31] Marchesi, S., Perez-Osorio, J., Tommaso, D.D., Wykowska, A.: Don’t overthink: fast decision making combined with behavior variability perceived as more human-like. 2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), Naples, Italy, 2020, pp. 54–59 (2020). doi:10.1109/RO-MAN47096.2020.9223522. Publisher: IEEE
- [32] Meltzoff, A.N.: ‘Like me’: a foundation for social cognition. *Developmental Science* **10**(1), 126–134 (2007). doi:https://doi.org/10.1111/j.1467-7687.2007.00574.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-7687.2007.00574.x
- [33] Meltzoff, A.N., Prinz, W. (eds.): The Imitative Mind: Development, Evolution and Brain Bases. Cambridge Studies in Cognitive and Perceptual Development. Cambridge University Press, Cambridge (2002). doi:10.1017/CBO9780511489969
- [34] Mirnig, N., Stollnberger, G., Miksch, M., Stadler, S., Giuliani, M., Tscheligi, M.: To Err Is Robot: How Humans Assess and Act toward an Erroneous Social Robot. *Frontiers in Robotics and AI* **4** (2017). doi:10.3389/frobt.2017.00021
- [35] Mohebbi, A.: Human-Robot Interaction in Rehabilitation and Assistance: a Review. *Current Robotics Reports* **1**(3), 131–144 (2020). doi:10.1007/s43154-020-00015-4

- [36] Morier, D.M., Borgida, E.: The conjunction fallacy: A task specific phenomenon? *Personality and Social Psychology Bulletin* **10**(2), 243–252 (1984). doi:10.1177/0146167284102010
- [37] Munafò, M.R., Nosek, B.A., Bishop, D.V.M., Button, K.S., Chambers, C.D., Percie du Sert, N., Simonsohn, U., Wagenmakers, E.J., Ware, J.J., Ioannidis, J.P.A.: A manifesto for reproducible science. *Nature Human Behaviour* **1**(1), 1–9 (2017). doi:10.1038/s41562-016-0021. Number: 1 Publisher: Nature Publishing Group
- [38] Nelson Douglas L, C.M.: Entangled Associative Structures and Context. **AAAI Spring Symposium: Quantum Interaction**, 98–105 (2007)
- [39] Nelson, L.D., Simmons, J.P., Simonsohn, U.: Let’s Publish Fewer Papers. *Psychological Inquiry* **23**(3), 291–293 (2012). doi:10.1080/1047840X.2012.705245. Publisher: Routledge _eprint: <https://doi.org/10.1080/1047840X.2012.705245>
- [40] Nickerson, R.S.: Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology* **2**(2), 175–220 (1998). doi:10.1037/1089-2680.2.2.175. Publisher: SAGE Publications Inc
- [41] Nisimura, R., Uchida, T., Lee, A., Saruwatari, H., Shikano, K., Matsumoto, Y.: ASKA: receptionist robot with speech dialogue system. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, pp. 1314–1319 vol.2 (2002). doi:10.1109/IRDS.2002.1043936
- [42] Oehler, A., Wendt, S., Wedlich, F., Horn, M.: Investors’ Personality Influences Investment Decisions: Experimental Evidence on Extraversion and Neuroticism. *Journal of Behavioral Finance* **19**(1), 30–48 (2018). doi:10.1080/15427560.2017.1366495
- [43] Polakow, T., Teodorescu, A.R., Bussemeyer, J.R., Gordon, G.: Free Ranking vs. Rank-Choosing: New Insights on the Conjunction Fallacy (2021). <https://psyarxiv.com/r9kxp>
- [44] Ragni, M., Rudenko, A., Kuhnert, B., Arras, K.O.: Errare humanum est: Erroneous robots in human-robot interaction. In: *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 501–506 (2016). doi:10.1109/ROMAN.2016.7745164
- [45] Robinson, N.L., Cottier, T.V., Kavanagh, D.J.: Psychosocial Health Interventions by Social Robots: Systematic Review of Randomized Controlled Trials. *Journal of Medical Internet Research* **21**(5), e13203 (2019). doi:10.2196/13203. Company: Journal of Medical Internet Research Distributor: Journal of Medical Internet Research Institution: Journal of Medical Internet Research Label: Journal of Medical Internet Research Publisher: JMIR Publications Inc., Toronto, Canada
- [46] Salem, M., Lakatos, G., Amirabdollahian, F., Dautenhahn, K.: Would You Trust a (Faulty) Robot?: Effects of Error, Task Type and Personality on Human-Robot Cooperation and Trust. In: *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI ’15*, pp. 141–148. ACM, New York, NY, USA (2015). doi:10.1145/2696454.2696497
- [47] Samar, S.M., Walton, K.E., McDermut, W.: Personality Traits Predict Irrational Beliefs. *Journal of Rational-Emotive & Cognitive-Behavior Therapy* **31**(4), 231–242 (2013). doi:10.1007/s10942-013-0172-1
- [48] Sava, F.A.: Maladaptive schemas, irrational beliefs, and their relationship with the Five-Factor Personality model. *Journal of Cognitive & Behavioral Psychotherapies* **9**(2) (2009). ISBN: 1584-7101 Publisher: Citeseer
- [49] Short, E., Hart, J., Vu, M., Scassellati, B.: No fair. An interaction with a cheating robot. In: *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 219–226 (2010). doi:10.1109/HRI.2010.5453193
- [50] Short, E., Swift-Spong, K., Greczek, J., Ramachandran, A., Litoiu, A., Grigore, E.C., Feil-Seifer, D., Shuster, S., Lee, J.J., Huang, S., Levonisova, S., Litz, S., Li, J., Ragusa, G., Spruijt-Metz, D., Mataric, M., Scassellati, B.: How to train your DragonBot: Socially assistive robots for teaching children about nutrition through play. In: *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pp. 924–929. IEEE, Edinburgh, UK (2014). doi:10.1109/ROMAN.2014.6926371
- [51] Simmons, J.P., Nelson, L.D., Simonsohn, U.: False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science* **22**(11), 1359–1366 (2011). doi:10.1177/0956797611417632. Publisher: SAGE Publications Inc
- [52] Stower, R., Calvo-Barajas, N., Castellano, G., Kappas, A.: A Meta-analysis on Children’s Trust in Social Robots. *International Journal of Social Robotics* (2021). doi:10.1007/s12369-020-00736-8
- [53] Syrdal, D.S., Dautenhahn, K., Koay, K.L., Walters, M.L.: The negative attitudes towards robots scale and reactions to robot behaviour in a live human-robot interaction study. *Adaptive and emergent behaviour and complex systems* (2009). Publisher: SSAISB
- [54] Tory Higgins, E., Rholes, W.S., Jones, C.R.: Category accessibility and impression formation. *Journal of Experimental Social Psychology* **13**(2), 141–154 (1977). doi:10.1016/S0022-1031(77)80007-3

- [55] Tversky, A., Kahneman, D.: Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review* pp. 293–315 (1983)
- [56] Vanman, E.J., Kappas, A.: “Danger, Will Robinson!” The challenges of social robots for intergroup relations. *Social and Personality Psychology Compass* **13**(8), e12489 (2019). doi:<https://doi.org/10.1111/spc3.12489>
- [57] Vishwanath, A., Singh, A., Chua, Y.H.V., Dauwels, J., Magnenat-Thalmann, N.: Humanoid co-workers: How is it like to work with a robot? In: 2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), pp. 1–6. IEEE Press, New Delhi, India (2019). doi:10.1109/RO-MAN46459.2019.8956421
- [58] Wallkötter, S., Stower, R., Kappas, A., Castellano, G.: A Robot by Any Other Frame: Framing and Behaviour Influence Mind Perception in Virtual but not Real-World Environments. In: Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, HRI '20, pp. 609–618. Association for Computing Machinery, New York, NY, USA (2020). doi:10.1145/3319502.3374800
- [59] Walters, M.L., Syrdal, D.S., Dautenhahn, K., te Boekhorst, R., Koay, K.L.: Avoiding the uncanny valley: robot appearance, personality and consistency of behavior in an attention-seeking home scenario for a robot companion. *Autonomous Robots* **24**(2), 159–178 (2008). doi:10.1007/s10514-007-9058-3
- [60] Weiss, A., Bartneck, C.: Meta analysis of the usage of the Godspeed Questionnaire Series. In: 2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 381–388 (2015). doi:10.1109/ROMAN.2015.7333568
- [61] Woods, S., Dautenhahn, K., Kaouri, C., Boekhorst, R., Koay, K.L.: Is this robot like me? Links between human and robot personality traits. In: 5th IEEE-RAS International Conference on Humanoid Robots, 2005., pp. 375–380 (2005). doi:10.1109/ICHR.2005.1573596

A Detective scenario

1. A diamond shop was robbed. The police came straight away and caught a couple of people (separately) near the scene, but no diamonds were found. The detective assigned to the case is sure, with no doubt, that at least one of the suspects did it. You and the robots need to help the detective figure out who robbed the shop. The suspects are:
 - (a) Suspect – A: Is tall, with a black Louis Vuitton leather jacket and a Rolex watch.
 - (b) Suspect – B: Has blonde hair and had a cut above the left eyebrow.
2. Person rate A,B, A and B.
3. New information:
 - (a) Suspect – C: Tried to run from the scene when the police asked to stop.
 - (b) Suspect – D: Is not willing to talk before a lawyer is present.
4. Robots rate C, D, C and D:
5. Ask the person which robots you agree with? .
6. Person: after hearing the robots person rate A and C:
7. New information: Suspect B told the police that she got the cut when from a tree branch when she took her dog for a walk a few hours earlier to the incident. Suspect D is still refusing to talk even when the lawyer arrived.
8. Robots: give their ratings for B and D.
9. Ask the person which robots you agree with?
10. Person: rank all the possibilities.
11. Robots: also gives rankings of top three possibilities.
12. The detective asks: “who did it?” (to see if the robots’ rankings affected the person’s rankings)

B Art scenario

1. In the last couple of years, the market for fine art is booming. Last year the most expensive piece of art work that was sold was an expressionist oil painting by the late famous artist. Also last year, one item that caught the most attention around the world was made by a famous young graffiti artist. Tonight, you and the two robots went together to an art auction. There were expressionist paintings to realistic sculptures of horses and even a few pieces of graffiti art work. The auction was a great success, though not all the pieces were sold. Person: What piece do you think was sold for the highest price. Rate the options:

- (a) A realistic piece. (A)
 - (b) A piece from a famous artist. (B)
 - (c) A realistic piece made by a famous artist. ($A \cap B$)
2. New information: One piece caught a lot of attention from the young investors in the audience.
3. Robots: Which piece do you think it was?
 - (a) An expressionist piece. (C)
 - (b) A piece from a young artist. (D)
 - (c) An expressionist piece made by a young artist. ($C \cap D$)
4. Ask the person which robots you agree with?
5. Person: after hearing the robots person rate A and D:
6. New information the most expensive art work in the auction contained a figure of a person, but it wasn't clear what was it made of.
7. Robots: give their ratings for A and D.
8. Ask the person which robots you agree with?
9. Person: rank all the possibilities.
10. Robots: also gives rankings of top three possibilities.
11. The person is being asked: "what was the most expensive art work that was sold?" (to see if the robots' rankings affected the person's rankings)
12. Which robot would you hire as an art buyer?