

Would you like to play with me? How robots' group membership and task features influence human-robot interaction

Markus Häring, Dieta Kuchenbrandt, Elisabeth André

Angaben zur Veröffentlichung / Publication details:

Häring, Markus, Dieta Kuchenbrandt, and Elisabeth André. 2014. "Would you like to play with me? How robots' group membership and task features influence human-robot interaction." In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction - HRI '14, Bielefeld, Germany, March 3 - 6, 2014*, edited by Gerhard Sagerer, Michita Imai, Tony Belpaeme, and Andrea Thomaz, 9–16. New York, NY: ACM Press.
<https://doi.org/10.1145/2559636.2559673>.



Would You Like to Play with Me? How Robots' Group Membership and Task Features Influence Human–Robot Interaction

Markus Häring
Human Centered Multimedia
Institute for Computer Science
Augsburg University, Germany
haering@hcm-lab.de

Dieta Kuchenbrandt^{*}
Center of Excellence in
Cognitive Interaction
Technology (CITEC)
Bielefeld
Germany
dkuchenb@cit-ec.uni-
bielefeld.de

Elisabeth André
Human Centered Multimedia
Institute for Computer Science
Augsburg University, Germany
andre@hcm-lab.de

ABSTRACT

In the present experiment, we investigated how robots' social category membership and characteristics of an HRI task affect humans' evaluative and behavioral reactions toward robots. Participants ($N = 38$) played a card game together with two robots, one belonging to participants' social in-group and the other one being a social out-group member. Furthermore, participants were either asked to cooperate with the in- and to compete with the out-group robot (congruent condition), or they were asked to cooperate with the out-group robot while competing with the in-group robot (incongruent condition). The results largely support our hypotheses: Participants showed more positive evaluative reactions toward the in-group (vs. the out-group) robot and they anthropomorphized it more strongly, independent of the congruency or incongruence of the HRI. Moreover, if required, participants cooperated with both the in- and the out-group robot, whereas their cooperativeness was more pronounced toward the in-group robot. Finally, participants indicated more difficulties with the HRI in the incongruent vs. the congruent condition. The theoretical and practical implications of the findings are discussed.

1. INTRODUCTION

Nowadays robots are more and more deployed as (semi-) autonomous tools or servants, for instance, in manufacturing, search-and-rescue scenarios, or in the military (e.g., [11, 2, 26]). For the near future, however, scientists and developers envision robots in more socially interactive contexts such as assisting humans in their daily household chores or serv-

ing as co-workers and team partners in a work environment (e.g., [24]). However, how do humans react to and interact with socially active robots? Which factors determine whether humans are willing to accept a robot's assistance or to work together with a robot team partner? With the present research we aim to provide some more insights into these still open research questions. More specifically, we wanted to investigate the effects of a robot's social group membership and task features on human–robot interactions (HRI) and humans' perceptions of social robots.

2. RELATED WORK

How humans perceive social robots and how they interact with these non-human agents is a crucial question in HRI research. In the context of human-computer interaction, Reeves and Nass [21] observed that humans tend to treat computers and virtual agents like human beings (see also [16, 17]). For instance, they found that people mindlessly apply human social categories such as ethnicity to them: Participants in their studies evaluated a computer that ostensibly belonged to the participants' ethnic in-group more positively than a computer that was equipped with out-group cues [18]. This resembles the well-documented phenomenon of intergroup bias in human-human contexts [25]: When we perceive others as members of a social in- or out-group (e.g., in terms of ethnicity, religion, or gender), we treat and evaluate them differentially, favoring in-group members and derogating social out-groupers. Importantly, not only in the human-human or human-computer context but also in the field of HRI, research has documented such intergroup bias towards robots. Accordingly, humans base their perceptions and judgments of robots on cues indicating a robot's social category membership [5, 12]. To illustrate, Eyssel and Kuchenbrandt [5] found that people favored a robot that ostensibly belonged to their national in-group over a robot allegedly being a member of a national out-group. The in-group robot was evaluated more positively as well as anthropomorphized more strongly (i.e., attributed more human qualities) than the national out-group robot. Similar effects could be demonstrated even with arbitrary social categories that were irrelevant for participants' social identity [12, 18].

^{*}The first and second author contributed equally to the presented work and the manuscript. The authors therefore share the first authorship.

Importantly, humans’ tendency to apply social categories to robots not only has perceptual and evaluative but also behavioral consequences. In a study by Powers and colleagues [20], participants used knowledge about gender roles when interacting with a gendered robot: They elaborated less on a typically female topic (i.e., dating norms) when talking to an ostensibly female robot than when talking to a male robot. Furthermore, research by Siegel, Brezeal, and Norton [23] has shown that humans evaluated a robot of the opposite gender more positively than a same-gender robot, and they even tended to behave more positively toward a robot of the opposite gender. Consequently, alleged social category membership of robots plays a key role in how humans react toward robots and thus clearly needs to be considered when developing robots for different fields of applications.

In order to better understand humans’ perceptions of robots and the way humans and robots interact with each other, it is necessary to go beyond the investigation of robot or user features. An HRI is likewise characterized by features of the HRI task that, in turn, could influence humans’ perceptions of and behavior toward robots. In line with this reasoning, Mutlu and colleagues [15] yielded empirical evidence for the influence of the structure of an HRI task (i.e. competitive vs. cooperative structure) on the interaction itself and on perceptions of robots. In their experiment, males and females played an interactive video game with a robot, and they did so either in a cooperative or in a competitive way. The results showed that men based their evaluation of the robot to a large extent on the task structure, whereas women were more influenced by the characteristics of the robot. In a different set of studies [8], participants found a robot more suitable for a task when the degree of the robot’s humanlikeness matched the degree of sociability required by the task (see also [3]). In contrast, Kuchenbrandt et al. [13] present findings that suggest that a mismatch between robot characteristics (male vs. female robot) and the task type (male vs. female task type) could be beneficial for HRI: Females evaluated an ostensibly female robot more positively when they interacted with it on a typically male task than on a female task, whereas no such effect could be observed for male participants [13]. In sum, research investigating the impact of task features such as task type or task structure and their interplay with robot and user characteristics is still in its infancy and so far yielded inconclusive results. Consequently, more research is needed to shed light on the influence of robot and task features on HRI outcomes.

With the present research we aim to examine for the first time how robots’ social category membership in-group vs. out-group robot) and the structure of an HRI task (cooperative vs. competitive) in an HRI affect humans’ perceptions of and behavior toward robots. In line with the research reviewed above, we hypothesize that 1) participants will perceive a robot that ostensibly belongs to their social in-group more positively than an out-group robot. That is, they will perceive it as more competent, anthropomorphize it to a greater extent, and perceive it as psychologically more close to themselves. However, according to the contact hypothesis [1], cooperative contact with human out-group members improves our evaluative reactions toward these persons and toward the whole out-group. Consequently, we 2) assume that a cooperative HRI with an out-group robot should result in improved reactions toward that robot compared to competing with a robot target that ostensibly belongs to a

social out-group. Social psychological findings further indicate that people are more willing to cooperate with members of social in-groups than with out-group members [7], while out-group members are often perceived as competitors [22]. Competing with an out-group member and cooperating with an in-group member could therefore be regarded as humans’ behavioral default. Consistent with this reasoning, we 3) suppose that participants will cooperate with an in-group robot to a greater extent than with an out-group robot. Related to this, competing with a human in-group member and cooperating with an out-group member should, at least in part, contradict humans’ behavioral scripts regarding intra- and intergroup interactions, and should thus be more challenging and demanding. Likewise, within the context of an HRI we 4) assume that participants will perceive an HRI as more demanding when they have to cooperate with an out-group robot and compete with an in-group robot (incongruent interaction) than when an in-group robot serves as a cooperation partner and the out-group robot represents the competitor (congruent interaction).

3. METHOD

3.1 Participants and Design

$N = 38$ Germans (23 males, 15 females) took part in the experiment¹. Their age ranged from 21 to 62 years ($M = 28.21, SD = 7.52$). Participants were randomly allocated to one of two conditions resulting from a 2 (robot-task congruency: congruent vs. incongruent) by 2 (robot group membership: in-group vs. out-group) experimental design with the first factor as a between-subjects factor and the latter one as a within-subjects factor. That is, participants simultaneously interacted with both an in-group and an out-group robot. However, they either played a game together with the in-group robot as a cooperation partner and the out-group robot as a competitor (congruent). Or they played the game having the in-group robot as a competitor while the out-group robot served as the cooperation partner (incongruent).

3.2 Procedure

Participants were tested individually in a laboratory at Augsburg University. They were told that they would play a card game with two different robots. In this card game, all players were supposed to find the best fitting pairs of words that were displayed on the cards. The experimenter introduced the two robots to the participants by mentioning their alleged names (MAIK and MALIK, see section 3.3). Moreover, participants learned that the artificial intelligence (A.I.) of the two robots was programmed by two different student groups (a group of German students and a group of Egyptian guest-students, see section 3.3). Following the introduction, participants received written instructions and information about the course and the rules of the card game (see section 3.2.1). Importantly, at this point participants learned that one of the two robots would be their cooperation partner while the other robot would serve as their com-

¹Originally, 43 participants (26 male, 17 female, age $M = 27.84, SD = 7.16$) took part in the experiment. Five participants had to be excluded from further analyses, as they did not correctly remember the group membership of the two robots. However, including these five participants in all analyses did not significantly affect our results.



Figure 1: Experimental set-up with the in-group robot MAIK as the cooperation partner and the out-group robot MALIK as the competitor

petitor in the game. After playing the card game together with the two robots, the participants completed a computerized questionnaire containing our dependent variables, were thanked, debriefed, and dismissed.

3.2.1 Experimental set-up

Participants were sitting in front of a Microsoft Surface (V1.0)² touch-screen table. Two robots (NAO, Academic Edition V4.0, Aldebaran Robotics) were sitting on both sides of the table. One robot served as the participants' cooperation partner and was always placed at the left hand side of the table (as seen from the participant) while the other robot (the competitor) was placed at the right hand side (see Fig. 1).

On the touch-screen table the participants could see two decks of red and green cards as well as three sets of four red cards that were placed in front of each player. These sets represented the players' hand cards (see Fig. 2). The participants' hand cards were face up, so that the participants could see the content of their cards. The robots' cards were displayed face down. Red cards always displayed a noun, while green cards contained adjectives.

Each round of the game followed the same pattern (see Fig. 3): One of the players (active player) draws a green adjective-card from the green deck. The drawn card is then automatically displayed face up. Subsequently, each of the two remaining players (passive players) suggests a noun from the four hand cards that could fit the adjective on the green card. Finally, the active player decides whether to take the noun suggested by passive player one or by passive player two. The player whose suggested noun has been chosen gains a point. A round is finished with the played cards being removed and the passive players' red cards being refilled. For each round, the active player changed clockwise. The complete game consisted of 21 rounds; therefore, each player served as the active player for seven rounds. The player with the highest number of points is considered as the winner of the game.

Irrespective of the information given to the participants, the robots were not equipped with an A.I. but instead fol-

²The tabletop was renamed in June 2012 with V2.0 to PixelSense, www.pixelsense.de




Figure 2: Hand cards, card decks, and each player's current points displayed on the touch-table screen

lowed a predefined script. Accordingly, the sequence of the green and red cards as well as the robots' moves in the game were equal in every game. The robots' decisions regarding whose suggested noun to choose followed a fixed random order and was independent of what participants actually suggested.

To sustain the impression that the participants were actually playing with intelligent robots, both robots were showing simple gaze and pointing behaviors. Moreover, the robots uttered sentences such as "Let's see what the next adjective is", or they gave feedback when receiving a point (e.g., "I also thought that my card was the better one"). This behavior did not differ amongst the two robots. Importantly, independent of their actual role as cooperation partner or competitor, both robots favored the participant over the other robot. That is, both robots chose the nouns that have been suggested by the participant four times (i.e., the participant always received four points by each robot) while choosing the suggestions of the other robot three times. With this, we aimed to assure that every difference in how participants perceived and interacted with the robots solely depended on the robots' ostensible group membership and their assigned role as cooperation partner versus competitor and thus were not affected by differences in the robots' behavior.

To interact with the table and the robots (picking green cards, making suggestions, choosing a suggested noun), the participants had to briefly touch the relevant cards with their fingers. The robots did not touch the table-top. Instead, the respective cards have been displayed automatically following a fixed script. However, the robots commented their 'moves' in the game verbally and showed gaze behavior toward the relevant cards. When the robots had to choose a noun from the two suggestions, they announced their choice verbally and pointed and looked at the player whose suggestion has been chosen.

3.3 Independent Variables

In order to manipulate *group membership* of the two robots, we varied three aspects: First, we named the in-group robot MAIK (a typically German name) and the out-group robot MALIK (a typically Arab name). Second, we told participants that German students have developed MAIK's artificial intelligence. In contrast, Egyptian guest students al-

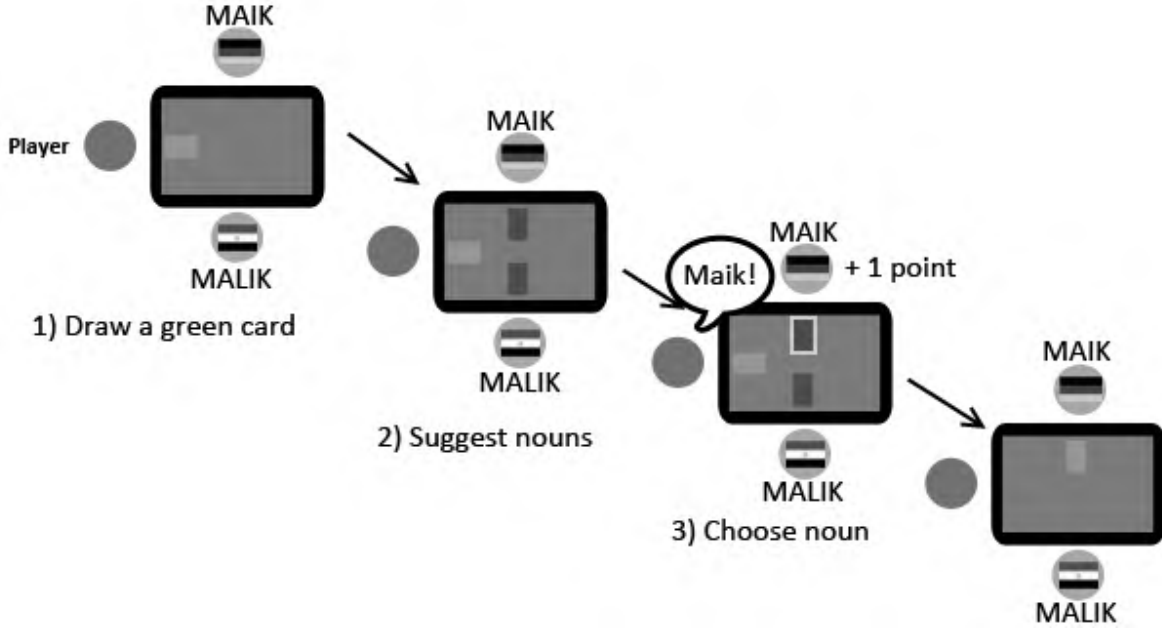


Figure 3: Three steps of a round : 1) the active player draws a green card, 2) the passive players suggest nouns, 3) the active player chooses a noun, the passive player with the chosen noun gets a point. Finally the active player changes clockwise.

legedly developed MALIK’s artificial intelligence³. Third, MAIK was equipped with a sticker displaying the German national flag, whereas MALIK was provided with the Egyptian flag.

The *robot-task congruency* was manipulated by forcing participants to cooperate either with the in- or with the out-group robot. In the *congruent condition*, participants were told that they could take part in a lottery and have the chance to win 50 Euro if their own points in the game plus MAIK’s points together exceed double the points of MALIK. We thus presented the in-group robot MAIK as the cooperation partner in the game while the out-group robot MALIK was presented as a competitor. In the *incongruent condition*, in contrast, the out-group robot served as the cooperation partner and the in-group robot was introduced as the competitor. Consequently, participants were informed that they could participate in the lottery if their own points and MALIK’s points together exceed double the points of MAIK. The opportunity to participate in the lottery was given in order to increase the incentive to cooperate (vs. compete) with the respective robot.

3.4 Dependent Variables

The ratings regarding the two robots were always completed en bloc for each robot, whereby the two blocs were presented in random order. To assess the endorsement of the dependent variables, we used seven-point Likert scales (ranging from 1 = *not at all* to 7 = *very*). High values reflect high endorsement of the assessed dimension.

³Egyptian students actually visited the lab the first author belongs to during the time the study was conducted. This contributed to the credibility of the cover story.

Manipulation check. As a manipulation check we asked participants to indicate how strongly they perceived MAIK and MALIK as a team partner.

Perceived competence. Participants were asked to rate MAIK and MALIK with regard to four competence-related traits (efficient, skillful, confident, competent; see [6]). The four items were combined to form an acceptably reliable index for the perceived in-group robot’s ($\alpha = .69$) and out-group robot’s competence ($\alpha = .74$).

Mind attribution. Participants completed nine items with regard to MAIK’s and MALIK’s mental capabilities. The items were adapted from Gray, Gray, and Wegner ([9], see also [5]): ‘To what extent is the robot capable of feeling hungry/joy/pain/fear?’; ‘To what extent is the robot capable of hoping for things?’; ‘How likely is it that the robot has a personality/own will?’; ‘To what extent is the robot capable of being aware of things?’; ‘How likely is it that the robot has a soul?’. These items formed reliable indexes for both robots ($\alpha = .88$ for the in-group robot; $\alpha = .87$ for the out-group robot) and were used as an indicator of anthropomorphism [5].

Psychological closeness. To assess the degree of perceived psychological closeness between the participants and MAIK or MALIK, respectively, participants were asked to respond to the following five items: ‘To what extent do you feel close / connected / similar to MAIK (MALIK) / on the same wavelength with MAIK (MALIK)?’ and ‘Do you share many commonalities with MAIK (MALIK)?’ [4]. This measure was reliable for both robots ($\alpha = .93$ for the in-group robot; $\alpha = .93$ for the out-group robot).

Perceived task difficulty. To assess whether participants perceived the game as differentially demanding depending on the experimental condition, they were asked to indicate how difficult it was for them to play the game.

Perceived stress. Participants were further asked to indicate how much they felt stress during playing the game.

Cooperative behavior. We defined cooperative behavior as favoring the cooperation partner over the competitor. That is, in order to cooperate with one of the robots, the participants had to give more points to the robot cooperation partner (i.e., choosing more often the nouns suggested by this robot) than to the robot competitor. To build a cooperation index, we subtracted the points that participants allocated to the out-group robot MALIK from the points participants allocated to the in-group robot MAIK. Consequently, a positive value of the index indicates cooperation with the in-group robot (more points were allocated to MAIK than to MALIK), while a negative value indicates cooperation with the out-group robot MALIK.

4. RESULTS

4.1 Manipulation Check

As a manipulation check, we first tested whether participants perceived the in-group robot MAIK and the out-group robot MALIK as a team partner as a function of the experimental condition. The results of t-tests show that the in-group robot indeed was perceived as a team partner to a greater extent in the congruent ($M = 3.67, SD = 0.91$) compared to the incongruent condition ($M = 1.40, SD = 0.60$), $t(36) = 9.18, p < .001$. Similarly, participants perceived the out-group robot more strongly as a team partner in the incongruent ($M = 3.20, SD = 1.20$) than in the congruent condition, ($M = 1.94, SD = 0.80$), $t(33.41) = 3.83, p = .001$.

4.2 Test of Main Hypotheses

4.2.1 Ratings of the robots (Hypotheses 1 and 2)

Perceived competence. Results of a mixed models analysis of variance (ANOVA) with robot-task congruency (congruent vs. incongruent) as the between-subjects factor and robot type (in-group vs. out-group robot) as the within-subjects factor yielded a marginally significant main effect of robot type, $F(1, 36) = 3.31, p = .08, \eta^2 = .08$. No interaction effect was found, $F < 1$. That is, participants tended to ascribe the in-group robot MAIK more competence ($M = 4.84, SD = .96$) than the out-group robot MALIK ($M = 4.58, SD = 1.07$), irrespective of whether they cooperated or competed with the in- or out-group robot, respectively (see Fig. 4).

Mind attribution. Results of a mixed models ANOVA revealed a similar pattern: A main effect of robot type was found, $F(1, 36) = 4.04, p = .05, \eta^2 = .10$, but no interaction effect, $F < 1$. Participants attributed more mind to the in-group ($M = 2.22, SD = 1.13$) than to the out-group robot ($M = 2.06, SD = 0.99$).

Psychological closeness. An ANOVA with robot-task congruency (congruent vs. incongruent) as the between-subjects factor and robot type (in-group vs. out-group robot) as the within-subjects factor yielded no main effect of robot type, $F(1, 36) = 1.19, p = .28, \eta^2 = .03$. However, as hypothesized, we found a significant interaction effect of robot type by experimental condition, $F(1, 36) = 6.94, p = .01, \eta^2 = .16$. To inspect this pattern of findings further, we conducted post hoc t-tests and compared the psychological closeness to the in- and the out-group robot between the experimental




Figure 4: Mean ratings of the robots' competence and mind as a function of the robots' group membership




Figure 5: Mean ratings of psychological closeness to the robots as a function of experimental condition

conditions: Participants in the congruent condition rated the in-group robot MAIK as closer to themselves ($M = 3.69, SD = 1.35$) than participants in the incongruent condition ($M = 2.89, SD = 1.22$), $t(36) = 1.92, p = .06$, although this difference only approached significance. No such difference between the experimental conditions was found for the out-group robot MALIK (congruent condition: $M = 2.87, SD = 1.16$; incongruent condition: $M = 3.23, SD = 1.50$), $t(36) = -0.82, p = .42$. Moreover, we further compared the perceived closeness to the in-group robot with the perceived closeness to the out-group robot separately for each condition. In the congruent condition, participants perceived the in-group robot as closer to themselves ($M = 3.69, SD = 1.35$) than the out-group robot ($M = 2.87, SD = 1.16$), $t(17) = 3.00, p = .01$. In the incongruent condition, in contrast, no such difference was obtained, $t(19) = -1.01, p = .33$. Results are depicted in Figure 5.

4.2.2 Ratings of the task (Hypothesis 4)

Perceived task difficulty. Results of a t-test demonstrate that participants indicated to have more difficulties with the interaction task in the incongruent ($M = 2.15, SD = 1.57$) than in the congruent condition ($M = 1.28, SD = 0.58$), $t(36) = -2.23, p = .03$ (see Fig. 6).




Figure 6: Mean ratings of perceived task difficulty and stress as a function of congruency condition

Perceived stress. A comparable pattern of results was obtained for perceived stress during the HRI game. Participants in the incongruent condition in tendency rated the interactive game as more stressful ($M = 2.20, SD = 1.67$) than in the congruent condition ($M = 1.61, SD = 0.70$). However, this was only a non-significant tendency, $t(25.97) = -1.14, p = .16$.

4.2.3 HRI behavior (Hypothesis 3)

Cooperative behavior. With a t-test we tested whether the cooperation index varies as a function of experimental condition. As can be seen in Figure 7, participants cooperated with the in-group robot MAIK in the congruent condition ($M = 1.71, SD = 1.21$) while they cooperated with the out-group robot MALIK in the incongruent condition ($M = -0.90, SD = 1.77$)⁴, $t(35) = 5.12, p < .001$. Additionally, we tested whether the cooperation index differs significantly from 1 in the congruent or from -1 in the incongruent condition, respectively. In the congruent condition, a value of 1 would be the minimum value to indicate cooperation with the in-group robot whilst in the incongruent condition a value of -1 represents the minimum value for cooperation with the out-group robot⁵. However, we were interested in whether participants showed cooperative behavior towards the in- or out-group robot over and above the to-be-expected minimum value. Interestingly, participants in the congruent condition clearly showed a cooperative tendency towards the in-group robot over and above the minimum cooperation value of 1, $t(16) = 2.40, p = .03$. In contrast, in the incongruent condition, the cooperation index did not significantly differ from the minimum cooperation value of -1, $t(19) = 0.25, p = .80$. This latter finding demonstrates that although participants cooperated with the out-group robot when the robot served as the cooperation partner, their cooperative behavior did not exceed the to-be-expected minimal cooperative behavior towards the out-group robot.

⁴Remember that a positive value indicates cooperation with the in-group robot while a negative value indicates cooperation with the out-group robot

⁵Participants had to allocate seven points between the two robots. Consequently, one robot always had to have at least one point more (or less) than the other robot, resulting in a minimal difference of 1 or -1 between MAIK's and MALIK's points, respectively.




Figure 7: Cooperation index for congruent and incongruent condition

5. DISCUSSION

With the present research we aimed to contribute to a clearer understanding of the factors that influence how humans react to and interact with social robots. More specifically, we tested whether intergroup bias toward a robot would occur depending on the alleged group membership of the robot. In addition, we examined the effects of HRI task structure (cooperative vs. competitive) on participants' reactions toward the robot and their perception of the HRI. In order to investigate our research questions, participants played a card game with two robots, one belonging to participants' in-group and one being a member of a social out-group. Furthermore, in the interaction, one of the two robots served as a cooperation partner while the other one was presented as a competitor. Importantly, the study was designed in such a way that half of the participants cooperated with the in-group robot and competed with the out-group robot (congruent condition) whilst for the remaining participants the out-group robot was the cooperation partner and the in-group robot represented the competitor (incongruent condition).

We first of all hypothesized that participants will react more positively toward an alleged in-group versus out-group robot (Hyp. 1). Similar to previous findings [5, 12], participants indeed perceived the in-group robot MAIK as more competent than the out-group robot MALIK. They even anthropomorphized the robot to a greater extent, that is, they attributed more mind to the in-group versus the out-group robot. Interestingly, this latter finding is in line with social psychological research demonstrating that human in-group members are often attributed more humanness compared to human out-groupers [10, 14]; or, to put it differently, members of social out-groups are often dehumanized. Our study substantiates the claim that social categories not only play a key role in the perception of social robots, but social categories lead to similar consequences in HRI as they do in human-human interactions (see also [5, 12]).

We secondly expected that the effects of the robots' group membership would interact with the effects of the HRI task structure. That is, we expected the participants to show improved reactions toward the out-group robot when being forced to cooperate versus to compete with it (Hyp. 2). Our findings do not support this notion: With regard to the robots' perceived competence and mind, the robot-task

congruency did not show any influence. Interestingly, our results indicate that the in-group robot MAIK benefitted from being the participants' cooperation partner. Participants perceived MAIK as closer to themselves when they cooperated versus competed with it. In contrast, no such beneficial effects could be found for the out-group robot MALIK. Participants reported similar psychological distance to MALIK irrespective of whether they cooperated or competed with it. These findings contradict results from intergroup contact research in the context of human-human interactions that document the beneficial effects of positive cooperative out-group contact [1, 19] on reactions toward members of social out-groups. Moreover, they again illustrate the importance of social categories in HRI: Intergroup bias toward robots appears to be hard to overcome and the effects of a robot's social category membership even seem to overrule the potential impact of a positive social HRI.

Thirdly, we hypothesized that participants would cooperate with the alleged in-group robot to a greater extent than with the out-group robot (Hyp. 3). Previous research in human-human contexts has proven that people more readily cooperate with their in-groupers than with out-group members [7]. The results fully support our third hypothesis. Participants cooperated with both the in- and the out-group robot if requested. However, the cooperation index clearly showed that participants cooperated with the in-group robot to a greater extent (in the congruent condition) than with the out-group robot (in the incongruent condition). For a more refined interpretation of these results, the winning conditions of the game need to be reconsidered: In order to win the game - which could be viewed as the participants' individual goal - participants had to ensure that they obtain the highest overall score compared to each of the two robots. At the same time, participants had to give more points to their robot cooperation partner (i.e., choosing more of the nouns suggested by this robot) than to the robot competitor in order to participate in the lottery - which could be viewed as the mutual goal shared between the participants and the respective robot cooperation partner. However, this cooperation behavior at a certain point involves the risk of not reaching one's individual goal of winning the card game: The more points participants allocated to the cooperation partner, the higher the likelihood that this robot would gain more points than the participant himself/herself, and thus, the higher the likelihood for participants to lose the overall win of the game. Against this background, the present findings suggest that the tendency to cooperate with an in-group robot compared to an out-group robot was so strong that participants even tended to risk their individual win⁶. In contrast, with the out-group robot MALIK as the cooperation partner, participants only showed the minimum of cooperation that was necessary to take part in the lottery. From an applied point of view, this finding is particularly important: Many future applications for social robots involve some degree of cooperation between the human user and the robot partner. The present results suggest that a human-robot cooperation could be reinforced by presenting the robot as an in-group member (i.e., giving the robot a

name indicating in-group membership). Likewise, for the purpose of smooth human-robot cooperation cues that signify out-group membership should be avoided.

According to our fourth hypothesis, we expected participants to find an HRI more demanding when robot features and task structure are incongruent, that is, when participants had to cooperate with an out-group robot while competing with an in-group robot (Hyp. 4). Indeed, our results support this assumption: Participants indicated that they had more difficulties in playing the card game in the incongruent than in the congruent experimental condition. Moreover, in tendency, they even felt more stressed in the incongruent interaction. Accordingly, this pattern of results adds to previous research suggesting that a match between robot and task features has a positive impact on how humans perceive a robot and how they interact with it [8, 3]. A match between robot and task features might be in accordance with users' expectations regarding an HRI, which in turn could render the HRI cognitively less effortful for the user. However, other findings [13] suggested that a mismatch could have positive effects, for instance, on how a robot is evaluated. Consequently, future research needs to identify moderating factors that determine when and why a match or mismatch between robot and task features would be advantageous for HRI.

We critically acknowledge that our group manipulation could be confounded with existing stereotypes and expectations about Germans and Egyptians, respectively. However, previous work demonstrated that the mere categorization of the robot as an in- versus out-group member using arbitrary and meaningless categories [12] leads to similar patterns of intergroup bias compared to group manipulations using nationality [4]. This makes us confident that also in the present experiment, the in-group / out-group differentiation was mainly responsible for the reported effects and not only differences in how Egyptians and Germans are perceived.

6. CONCLUSION

In sum, the present study showed that social categories represent core determinants of how we evaluate, perceive, and interact with robots. We demonstrated that perceiving a robot as an in-group member has positive impact on users' reactions toward that robot while categorizing a robot as an out-group member has negative implications for HRI. Importantly, this is the first study documenting that even a positive (cooperative) interaction with an out-group robot does not compensate for the negative impact of out-group category membership. Based on these findings, we conclude that when designing social robots, scientists and developers need to consider various cues of the robot (e.g., voice, appearance, name) that could indicate a robot's social category membership. A robot's alleged group membership together with its supposed application might strongly influence how users react toward and interact with this robot. At the same time, roboticists could make use of social categories in order to make an HRI an intuitive, pleasant and smooth experience for the user.

7. ACKNOWLEDGMENTS

We especially thank Sigfried Depner for his help with the implementation of the game and the robots' behavior.

⁶In the present study, a cooperation index of 2 indicates that the robot cooperation partner gained more points than the participants (who always received eight points). In the congruent condition, the cooperation index of $M = 1.71$ was quite close to this value.

8. REFERENCES

- [1] G. W. Allport. *The nature of prejudice*. Oxford, UK: Addison-Wesley, 1954.
- [2] A. Davids. Urban search and rescue robots: from tragedy to technology. *Intelligent Systems, IEEE*, 17(2):81–83, 2002.
- [3] F. Eyssel and F. Hegel. (s)he’s got the look: Gender stereotyping of robots. *Journal of Applied Social Psychology*, 42(9):2213–2230, 2012.
- [4] F. Eyssel and D. Kuchenbrandt. My robot is more human than yours: Effects of group membership on anthropomorphic judgments of the social robot flobi. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2011)*, 2011.
- [5] F. Eyssel and D. Kuchenbrandt. Social categorization of social robots: Anthropomorphism as a function of robot group membership. *British Journal of Social Psychology*, 51(4):724–731, 2012.
- [6] S. T. Fiske, A. J. Cuddy, and P. Glick. Universal dimensions of social cognition: warmth and competence. *Trends in Cognitive Sciences*, 11(2):77 – 83, 2007.
- [7] L. Goette, D. Huffman, and S. Meier. The impact of group membership on cooperation and norm enforcement: Evidence using random assignment to real social groups. *The American economic review*, 96(2):212–216, 2006.
- [8] J. Goetz, S. Kiesler, and A. Powers. Matching robot appearance and behavior to tasks to improve human-robot cooperation. In *Robot and Human Interactive Communication, 2003. Proceedings. ROMAN 2003. The 12th IEEE International Workshop on*, pages 55–60, 2003.
- [9] H. M. Gray, K. Gray, and D. M. Wegner. Dimensions of mind perception. *Science*, 315(5812):619, 2007.
- [10] N. Haslam. Dehumanization: An integrative review. *Personality and Social Psychology Review*, 10(3):252–264, 2006.
- [11] H. Kitano and S. Tadokoro. Robocup rescue: A grand challenge for multiagent and intelligent systems. *AI Magazine*, 22(1):39, 2001.
- [12] D. Kuchenbrandt, F. Eyssel, S. Bobinger, and M. Neufeld. When a robot’s group membership matters. *International Journal of Social Robotics*, 5(3):409–417, 2013.
- [13] D. Kuchenbrandt, M. Häring, J. Eichberg, and F. Eyssel. Keep an eye on the task! how gender typicality of tasks influence human-robot interactions. In S. Ge, O. Khatib, J.-J. Cabibihan, R. Simmons, and M.-A. Williams, editors, *Social Robotics*, volume 7621 of *Lecture Notes in Computer Science*, pages 448–457. Springer Berlin Heidelberg, 2012.
- [14] J.-P. Leyens, P. M. Paladino, R. Rodriguez-Torres, J. Vaes, S. Demoulin, A. Rodriguez-Perez, and R. Gaunt. The emotional side of prejudice: The attribution of secondary emotions to ingroups and outgroups. *Personality and Social Psychology Review*, 4(2):186–197, 2000.
- [15] B. Mutlu, S. Osman, J. Forlizzi, J. Hodgins, and S. Kiesler. Task structure and user attributes as elements of human-robot interaction design. In *Robot and Human Interactive Communication, 2006. ROMAN 2006. The 15th IEEE International Symposium on*, pages 74–79, 2006.
- [16] C. Nass, B. J. Fogg, and Y. Moon. Can computers be teammates? *Int. J. Hum.-Comput. Stud.*, 45(6):669–678, Dec. 1996.
- [17] C. Nass, K. Isbister, and E.-J. Lee. Truth is beauty: researching embodied conversational agents. In *Embodied conversational agents*, pages 374–402. MIT Press, Cambridge, MA, USA, 2000.
- [18] C. Nass and Y. Moon. Machines and mindlessness: Social responses to computers. *Journal of social issues*, 56(1):81–103, 2000.
- [19] T. F. Pettigrew and L. R. Tropp. A meta-analytic test of intergroup contact theory. *Journal of personality and social psychology*, 90(5):751, 2006.
- [20] A. Powers, A. Kramer, S. Lim, J. Kuo, S. lai Lee, and S. Kiesler. Eliciting information from people with a gendered humanoid robot. In *Robot and Human Interactive Communication, 2005. ROMAN 2005. IEEE International Workshop on*, pages 158–163, 2005.
- [21] B. Reeves and C. Nass. *How people treat computers, television, and new media like real people and places*. CSLI Publications and Cambridge university press, 1996.
- [22] M. Sherif, O. J. Harvey, B. J. White, W. R. Hood, C. W. Sherif, et al. *Intergroup conflict and cooperation: The Robbers Cave experiment*, volume 10. University Book Exchange Norman, OK, 1961.
- [23] M. Siegel, C. Breazeal, and M. Norton. Persuasive robotics: The influence of robot gender on human behavior. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 2563–2568, 2009.
- [24] J. Stuckler, D. Holz, and S. Behnke. Robocup@home: Demonstrating everyday manipulation skills in robocup@home. *Robotics Automation Magazine, IEEE*, 19(2):34–42, 2012.
- [25] H. Tajfel and J. C. Turner. An integrative theory of intergroup conflict. *The social psychology of intergroup relations*, 33:47, 1979.
- [26] D. Voth. A new generation of military robots. *Intelligent Systems, IEEE*, 19(4):2–3, 2004.