

The Dark Side of Embodiment

Teaming Up With Robots VS Disembodied Agents

Filipa Correia, Samuel Gomes, Samuel Mascarenhas, Francisco S. Melo, Ana Paiva
INESC-ID and Instituto Superior Técnico, University of Lisbon
Portugal, Porto Salvo 2744-016
Email: filipacorreia@tecnico.ulisboa.pt

Abstract—In the past years, research on the embodiment of interactive social agents has been focused on comparisons between robots and virtually-displayed agents. Our work contributes to this line of research by providing a comparison between social robots and disembodied agents exploring the role of embodiment within group interactions. We conducted a user study where participants formed a team with two agents to play a Collective Risk Dilemma (CRD). Besides having two levels of embodiment as between-subjects —physically-embodied and disembodied—, we also manipulated the agents’ degree of cooperation as a within-subjects variable —one of the agents used a prosocial strategy and the other used selfish strategy. Our results show that while trust levels were similar between the two conditions of embodiment, participants identified more with the team of embodied agents. Surprisingly, when the agents were disembodied, the prosocial agent was rated more positively and the selfish agent was rated more negatively, compared to when they were embodied. The obtained results support that embodied interactions might improve how humans relate with agents in team settings. However, if the social aspects can positively mask selfish behaviours, as our results suggest, a dark side of embodiment may emerge.

I. INTRODUCTION

Over the past two decades, researchers have tried to understand how the embodiment of interactive social agents affects the way humans perceive them, and the impact it has on the interaction itself. In particular, the efforts in the field of Human-Robot Interaction have been focused on comparisons between robots and virtual or virtually-displayed versions of themselves. Generally, most findings support that robots can positively impact the interaction when compared to their virtual counterparts [26, 10]. Naturally, one can argue that robots have interaction affordances that are harder or even impossible for virtual agents to have, e.g. touch or manipulation of physical objects. Equally important is the fact that the actual recognition of some non-verbal cues on virtual agents is harder for humans [31]. Therefore, the interaction people have with robots, as opposed to virtual characters, is closer to the one they have with humans in subtle aspects, such as attribution of personal space [3], self-disclosure [18], empathy felt [5, 22, 40], or even brain activity [16].

Nevertheless, there are also works reporting no significant differences between the results obtained with robots and their virtual versions [5, 24, 39, 14], which support and reinforce the need for further research on this topic. Indeed, there are several particularities about the embodiment of a robot that still puzzle



(a) Embodied condition

(b) Disembodied condition

Fig. 1: Manipulation of the embodiment

researchers. Are they a matter of being physically-embodied in the real world as opposed to being virtually-embodied? Are they a matter of the physical presence and proximity to the user? Or are they related to the degree of social presence, agency or realism of the agent?

The actual definition of embodiment is also not clear yet and, in particular, to what extent an artificial agent can be considered embodied. For instance, we are witnessing a widespread use of intelligent *disembodied* voice assistants in our daily lives (e.g., Cortana, Google’s assistant, Siri), and some would argue about their intelligence [34], as well as their lack of embodiment [9]. Can we consider that voice itself is an anthropomorphic feature and that it is enough for humans to create an embodied mental model of an agent? Even when voice is not present, the actions of an autonomous agent may be enough to attribute agency and some degree of anthropomorphism. We posit that the study of embodiment lacks (to the best of our knowledge) proper comparisons with disembodied or quasi-disembodied agents. Such comparisons are particularly relevant when these extremes, i.e. embodied and disembodied agents, can both objectively execute certain tasks. In those situations, the degree of embodiment may be a mere facilitator of the social interaction, and may, in turn, affect the execution of that task. This paper contributes to this scope of literature by comparing how people perceive these two extreme levels of embodiment: physically-embodied (and co-present) agents and disembodied agents (Figure 1).

Furthermore, interactive agents will enter our lives in several contexts, such as the workspace, and their roles will require collaborative capabilities and, in most of those situations, they will even act as teammates. Within this vision, the appearance

of these cooperative machines was also postulated as one of the many open challenges [38]. Our paper addresses this question by exploring the impact of embodied affordances of socially interactive teammates in a cooperative task. If previous findings support the general idea that people comply more with social and moral norms around robots (compared to virtual agents), we also expected that people would cooperate, trust and identify more with a team of robots (compared to a team of virtual agents). Consequently, we also analysed how people perceive an artificial teammate that does not comply with collective norms and how the perception of such teammate is affected by its embodied affordances.

We conduct a user study using a mixed-design, where we manipulate the embodiment (between-subjects) and the degree of cooperation employed by each agent (within-subjects). As a result, each participant engages in a team setting with two autonomous agents that display opposite strategies to play a Collective Risk Dilemma¹: one is prosocial (high degree of cooperation) and the other is selfish (low degree of cooperation). Generally, we hypothesised that the embodied condition increases the degree of cooperation of the participants and would improve their subjective evaluation of the team. Moreover, we also expected that, in the embodied condition, the perception of the prosocial agent would be more positive and the perception of the selfish agent would be more negative, compared to the disembodied condition.

Our results suggest that the evaluation of the team was, to some extent, positively influenced by the embodiment. Surprisingly, the perception of the agents showed opposite results to what we have hypothesised. Finally, the degree of cooperation by the participants was similar between the two embodiment conditions. Based on the obtained results we drew considerations for the design of social behaviour of embodied agents, as well as for group interactions or team settings between humans and agents.

II. EMBODIMENT

There is no consensual definition of embodiment, especially when artificial agents come into play. For instance, Ziemke identified six notions of embodiment [45]. Although most of these notions try to discern whether a living body is required, there is a particular one that Ziemke identified as an orthogonal perspective – the *social embodiment*. It was initially proposed by Barsalou et al. and is focused on the interplay of embodiment and the social interaction [4]. The *social embodiment* requires an instantiation that mirrors “the state of the body” and has a central role during the interaction. However, such instantiation does not have to necessarily occur in the real world, as the notions of physical-, organismoid- or organismic embodiment do, and therefore considers as well virtually-embodied agents.

This paper adopts the aforementioned *social embodiment* notion and identifies as “disembodied” any agent that is not instantiated in a body, i.e. without any virtual or physical

shape. Consequently, it identifies any robot or virtual agent as embodied (physically- or virtually-embodied, respectively), as long as they possess some external and visible shape or representation for the user. Nevertheless, we would like to emphasise the close relation between *social embodiment* and the *structural coupling* defined by Quick et al. [35], in which an agent is considered embodied as long as mutual perturbations between it and its environment are possible. Later, on top of this notion, Dautenhahn et al. have proposed that the degree of embodiment can be measured by quantifying those mutual perturbations between the agent and the environment [9]. Interestingly, this notion does not require the agent to possess a body and the authors have actually raised the question “(...) *what, if anything, is special about material embodiment?*”. In that perspective, our definitions of embodied and disembodied agents could also be mapped into a higher and lower degree of embodiment, respectively. Due to clarity and practical reasons, we will keep the notion of *social embodiment*.

III. RELATED WORK

Two recent surveys related to the impact of embodiment on Human-Computer Interaction reported that the majority of the experimental studies generally claim co-present robots are rated more positively and provide more satisfactory interactions than virtually embodied agents [26, 10]. Nonetheless, some contradictory findings pointed out relevant variables and important considerations for the understanding of how the embodiment affects the dynamics and perceptions of people.

On the one hand, some works found evidence that a robot, when compared to a virtual or virtually displayed agent, was perceived as more enjoyable, watchful, helpful [44, 11, 18], socially present [11, 20], preferable [18, 28, 39], or generally evaluated it with higher scores [18, 21, 11]. Similarly, interacting with a robot can also improve the performance on the task [21, 25, 12], or provide higher levels of enjoyment [17, 33, 21] and engagement [17, 18, 19]. On the other hand, others found non-significant differences between a robot and a virtual agent on its trustworthiness [3, 39], perceived intelligence [14, 20], social presence [14, 27, 42], or even its general evaluation of the social attributes [39, 27]. Furthermore, there is also evidence that interacting with either a robot or a virtual agent can lead to similar levels of enjoyment and engagement [5, 17, 18], as well as similar performance on the task [28, 14].

One of the first investigations that brought divergent results was a set of two user studies conducted by Jung & Lee [24]. In the first experiment, participants were assigned to interact with either the AIBO robot or a virtual version of it. During the interaction, this pet-shaped agent performed singing and dancing and then entered a reactive mode that responded to touch inputs —in the robot condition— or to mouse clicks —in the virtual condition. The results of this first experiment showed participants in the robot condition attributed more social presence to it and rated the interaction more positively. Interestingly, the second experiment replicated the previous one except that the reactive capabilities in both conditions were disabled. The results of this second experiment showed

¹A variant of Public Goods Games, commonly used by game theorists.

opposite results with the virtual agent being the most socially present and the interaction with it being rated more positively. Therefore, the authors underlined the importance of tactile communication in physically embodied agents, i.e. robots.

Similarly, Kidd & Breazeal pioneered an experimental comparison between a co-present and a tele-present robot [17]. They found non-significant differences in the engagement of participants in these two conditions, while previously finding that a robot was more engaging than an animated character in a similar experiment. The authors further attributed these results to “*the fact that the robot is a real, physical thing, as opposed to the fictional animated character in the screen*”.

Later, Li performed a meta-analysis by surveying 33 experimental works on this topic and highlighted that comparisons between a co-located robot and a virtually embodied agent were mixing two variables —physical presence and embodiment [26]. He suggested a clear distinction between examining physical presence (i.e., by comparing co-present and tele-present robots) and embodiment (i.e., by comparing digitally-displayed robots and virtual counterparts). Interestingly, the two previously mentioned arguments by Jung & Lee and Kidd & Breazeal can be mapped into Li’s critique. The first one emphasised the aspect of physicality that can only be fully exploited in a co-present robot, while the latter suggested the embodiment itself may induce different degrees of realism that do not occur when comparing between the interaction with a robot and with a video-feed of it.

Social presence is another important aspect when analysing the impact of the embodiment. Although we have previously mentioned investigations that reported different levels of social presence for different levels of embodiment, other works could not find differences on social presence nor on their main behavioural or perception measures when comparing a robot and virtual version of it [14, 27, 42]. Hoffmann & Krämer have further postulated that the similar levels of social presence in both conditions might explain the absence of additional differences [14], which in turn supports the notion of embodiment by Dautenhahn et al. [9].

Another consideration highlighted in the work of Hoffmann & Krämer is the interplay of embodiment with the scenario. Their experiment manipulated not only the embodiment, but also the scenario by having a persuasive-conversational scenario and a task-oriented scenario where participants had to manipulate objects in the real-world [14]. The authors reported an interesting interaction effect between the two independent variables on the perceived competence of the agent. The co-present robot was perceived as the most competent in the task-oriented scenario, while the virtual agent was the most competent in the conversational scenario. Remarkably, if we revisit the reported findings regarding the obtained performances of people, the same pattern seems to appear. In tasks that require manipulations in the real world, most results suggest co-present robots outperform their virtual counterparts [44, 21, 14, 12]. In conversational or role-playing tasks, most findings point to no differences between a robot and a virtual agent [18, 39, 14, 27]. One might consider an exception to be

the case of learning tasks where the interaction is mediated by touch-screens or displays, in which more investigation is needed to evaluate a trend. On the one hand, Kennedy et al. found no significant differences between a robotic tutor and a virtually embodied one [15]. On the other hand, Leyzberg et al. found that a robotic tutor outperformed the virtual and the speech-only versions [25], and the authors attributed this result to a higher perceived authority due to its physical presence.

Finally, we would like to emphasise a recent work by Kontogiorgos et al., which is closer to our investigation due to their *smart speaker* condition with an embodied conversational agent with speech-only behaviours [20]. Their experimental study compared three levels of embodiment: a humanoid robot with verbal and non-verbal behaviour; the same humanoid robot only with verbal behaviour; and the *smart speaker*, a device without anthropomorphic design features (i.e., Alexa) that only had verbal behaviours. Results suggested the expression of non-verbal behaviours was able to increase conversational turns and attentional allocation at the cost of task performance. Additionally, the robot without non-verbal behaviours was rated with higher levels of social presence than the *smart speaker*, suggesting the anthropomorphic design features of the embodiment may mediate the social presence. Although the perceived intelligence was similar across the three conditions, it seems that in a task where people are required to manipulate physical objects, situation awareness may be conveyed by both non-verbal behaviour and anthropomorphic design features.

Overall, from the revised literature, we notice a lack of user studies within group interactions or multi-party collaborations. Furthermore, most embodiment manipulations did not consider minimalist embodiments or even totally disembodied agents. The novelty of our paper is to address these two aspects, by exploring the differences between teaming up with robots and with disembodied agents.

IV. EXPERIMENTAL SETUP

This user study aims at exploring the effects of embodiment on human perceptions of robotic teammates, as well as their degree of cooperation with the team of robots. In our experiment, participants were asked to team up with two autonomous agents and play a collaborative game. We manipulated, in a mixed-design, the embodiment of those agents (disembodied vs. robot), the outcome of the game (team wins vs. team loses), and the strategy of each agent (selfish vs. prosocial).

The agents are fully autonomous and interact in a social manner during the game. Those interactions were developed with the SERA toolkit [36] and are reactive to events from the game engine. The verbal and non-verbal utterances were previously scripted but they are chosen by an emotional decision making framework [29]. The code of the game and the scripted social behaviours of the agents are openly available².

A. Task

The main task was to play a collaborative game, called *For The Record*, which is extensively described in [37]. The

²<https://github.com/SamGomes/for-the-record>

game frames a Collective Risk Dilemma and uses a musical metaphor making it more entertaining and pleasing. Players form a “a band of musicians” and their goal is to record as many successful albums together. During the game, each player has two individual skills: *instrument* and *marketing*. The *instrument* is used to help the team create a good album, while *marketing* is used to collect individual revenue in the game. In the end, players are ranked according to their individual profit, but only if the team won the game (reached the threshold), otherwise all players lose their accumulated profits. The skills of *instrument* and *marketing* are used to throw dice and, consequently, to determine the contribution to the album (collective goal) and receive individual profit (individual goal), respectively. In each phase of the game, each player can throw as many dice as their level on that particular skill. The mixed-motive decision occurs when each player has to decide which skill they want to upgrade, which happens once per round. By upgrading the *instrument*, a player increases the likelihood of team’s success. By upgrading the *marketing*, on the other hand, a player increases the likelihood of his individual success.

Another important aspect of this game is the uncertainty factor due to the outcomes being determined by digital dice rolls. Players may increase their probability to succeed collectively or individually (by cooperating or defecting, respectively), but the outcome of their actions is ultimately controlled by an uncertain event. For the purpose of our experimental study, the outcomes of the dice were secretly manipulated to ensure a particular team outcome, given the experimental condition assigned to the participant.

B. Independent Variables

1) *Embodiment (between-subjects)*: The embodiment was manipulated in two levels. Participants would either play with two physically-embodied and co-present robots (Figure 1a) or with two disembodied agents (Figure 1b). The agents were fully autonomous in both conditions. In terms of behaviour, the robots in the embodied condition performed verbal and non-verbal behaviour (gaze, posture, animated expression of emotions). The disembodied agents only had verbal behaviour in the form of silent speech bubbles and used the exact same speech acts as the robots in the embodied condition. The speech bubbles of each agent would always be displayed in distinct positions to distinguish which one is speaking.

2) *Agents’ Playing Strategy (within-subjects)*: The two agents played the game with distinct strategies: the *prosocial* agent chooses to upgrade the *instrument* skill in each new round (cooperation); the *selfish* agent chooses to upgrade the *marketing* skill in each new round (defection). As each participant was exposed to the two strategies, this was a within-subjects variable. The agents were distinguished by their names in the game interface and, accordingly, in the questionnaires (*Emys* - selfish strategy; *Glin* - prosocial strategy).

3) *Outcome (between-subjects)*: The manipulation of the outcome had two levels in a between-subjects design: each participant would either win or lose the game. This was achieved by secretly manipulating the value obtained in the (digital)

dice throws during the game. The rationale for this control is based on previous findings suggesting that the outcome of a collaborative task significantly affects the perception of robotic teammates [8]. Note that by scripting the outcomes of the both dice, the *instrument* and the *marketing*, we can control both if the team wins or loses the game, as well as the ranks of each player, respectively. The *selfish* agent was always at the first place and the *prosocial* agent at last. The participant’s punctuation could be closer to one of the two agents, according to strategies employed by the participant, but would always end up at second place in the game.

C. Hypotheses

Our hypotheses are based on the premise that the interaction people have with robots is closer to the one they have with humans when compared with virtual or virtually-displayed agents. Overall, it seems that either the social presence or the physical presence of a robot leads people to attribute more agency, authority and realism to it, as well as they seem to comply more with social and moral norms around robots [3, 5, 18]. Consequently, we expect people to tolerate less the artificial teammate that is more selfish (and to praise more the prosocial teammate) when it is embodied compared to when it is disembodied. The following hypotheses are aligned with these results, in the sense that they assume that the presence of embodiment improves the perception of the team, the social attributes of the agents, and the degree of cooperation of the participants.

H1 Participants will perceive the team more positively in the embodied condition compared to the disembodied condition.

H2a Regardless of the embodiment, the perception of the prosocial agent will be more positive than the selfish agent.

H2b The perception of the prosocial agent will be more positive when it is embodied than when it is disembodied.

H2c The perception of the selfish agent will be more negative when it is embodied than when it is disembodied.

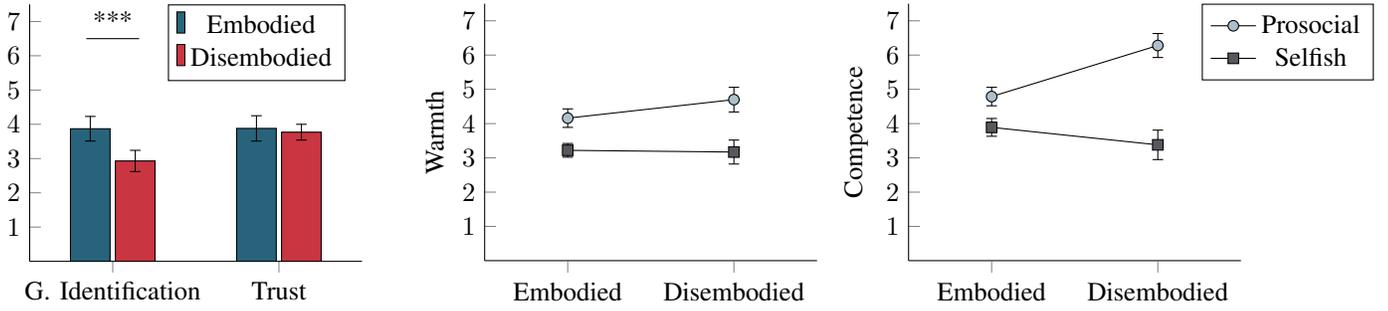
H3 Participants will be more prosocial in the embodied condition compared to the disembodied condition.

D. Dependent Measures

We assessed participants’ subjective evaluation of each agent in terms of their social attributes: warmth, competence and discomfort from RoSAS [7]. They were also asked to select only one agent for hypothetical future games, between the prosocial and the selfish one. Participants evaluated the quality and satisfaction with the team at a group level by reporting their subjective trust [1] and identification towards the team [23]. Finally, we analysed the objective cooperation rate of the participants by counting how many times they chose to cooperate, out of 4 decisions.

E. Procedure

The procedure of the experiment consisted of: (A) a briefing about the experiment, game rules and a demonstration of how to play the game; (B) a session of 5 rounds of the collective dilemma with two autonomous agents on the team;



(a) Main effects of the embodiment on the perception of group identification ($p < 0.001$) and group trust ($p = 0.517$). (b) Interaction effect between embodiment and the agents' strategy on the perception of warmth ($p = 0.035$). (c) Interaction effect between embodiment and the agents' strategy on the perception of competence ($p = 0.001$).

Fig. 2

(C) a final self-assessment questionnaire; and (D) a debriefing. However, the embodied condition was collected in the lab, while the disembodied was run on an online platform. The main difference on the procedures was in the initial briefing (A): for the embodied condition, a researcher explained the game rules in a verbal manner; and for the disembodied condition, it was an online tutorial that participants had to read. Participants in the embodied condition took approximately 30 minutes to complete the experiment while in the disembodied condition they took approximately 40 minutes. Also, the participation reward, given at the debriefing phase (D), was a fixed amount of 5\$ for participants in the disembodied condition and a cinema ticket of a similar price for participants in the embodied condition, regardless of their individual outcomes in the game.

F. Sample

In the embodied condition, 70 participants were recruited in person at the facilities of an energy company. For the disembodied condition, a total of 89 participants were recruited on MTurk. However, after validating the attention checks, out of those 89, only 41 valid participants remained. As such, out of the initial sample of 159, the statistical analysis was executed on a data set of 111 participants. Regarding the independent variable of the outcome, we collected 57 that won the game (35 in the embodied condition and 22 in the disembodied condition) and 54 that lost the game (35 in the embodied condition and 19 in the disembodied condition). Overall, in terms of gender, there were 47 females, 63 males and 1 participant chose "other". The ages ranged from 22 to 63 ($M = 36.060$, $SD = 10.816$).

As our two embodiment conditions were collected separately i.e., one in the lab and the other on MTurk, we checked for differences in the distributions of demographic variables. Both the age and gender of participants was similar across the four experimental groups, revealing identical demographics in our conditions. Furthermore, several experimental studies analysed the validity of comparing MTurk with laboratory subjects [43, 30, 6], even for economic games [2], showing comparable results in these two samples. Overall, given this

body of evidence we are convinced that it is quite unlikely that different results would be obtained if we opted for not using MTurk for the disembodied condition.

V. RESULTS

A. Group measures

1) **Group Identification:** The reliability of the group identification scale was $\alpha = 0.813$, which indicates a good internal consistency. Through a two-way ANOVA, we observed that embodiment had a significant medium-sized main effect on the participants' group identification perception ($F(1, 106) = 15.589$, $p < 0.001$, $r = 0.358$). Participants perceived higher group identification when collaborating with the embodied agents ($M_{Embodied} = 3.866$, $SD = 1.320$; $M_{Disembodied} = 2.932$, $SD = 1.482$) revealing that group identification was positively influenced by the presence of embodiment (Figure 2a). Game result also had a significant medium-sized main effect on the participants' group identification perception ($F(1, 106) = 24.913$, $p < 0.001$, $r = 0.436$). Participants perceived higher group identification when winning the game ($M_{Victory} = 4.056$, $SD = 1.367$; $M_{Loss} = 2.956$, $SD = 1.321$). No significant interaction effect was perceived between embodiment and game result ($F(1, 106) = 3.769$, $p = 0.055$, $r = 0.184$).

2) **Group Trust:** The reliability of the group trust scale was $\alpha = 0.945$, which indicates an excellent internal consistency. A two-way ANOVA showed no significant main effect of the embodiment ($F(1, 106) = 0.423$, $p = 0.517$, $r = 0.063$, Figure 2a), nor of the game result ($F(1, 106) = 3.201$, $p = 0.076$, $r = 0.170$). Moreover, no significant interaction effect was found between embodiment and game result ($F(1, 106) = 0.733$, $p = 0.394$, $r = 0.084$).

B. Perception of the agents

1) **Warmth:** The internal consistency for the warmth dimension was close to acceptable $\alpha = 0.693$. No main effects were observed for embodiment ($F(1, 106) = 2.574$, $p = 0.112$, $r = 0.155$) nor game result ($F(1, 106) = 1.422$, $p = 0.236$, $r = 0.114$). However, a large-size main effect of the agents' strategy could be observed ($F(1, 106) = 75.661$, $p <$

0.001, $r = 0.645$). The warmth attributed to the prosocial was significantly ranked higher than the warmth attributed to the selfish over all conditions ($M_{\text{prosocial}} = 4.364$, $SD = 1.152$; $M_{\text{selfish}} = 3.201$, $SD = 0.941$).

A two-way interaction effect (see Figure 2b) between strategy and embodiment was also found ($F(1, 106) = 4.566$, $p = 0.035$, $r = 0.202$). This indicates that the warmth values of the two strategies, prosocial and selfish, were affected differently by the embodiment. We performed a contrast analysis between different levels of embodiment for each agent, the prosocial and the selfish, using Mann-Whitney U tests. For the selfish agent, no significant difference was found on its warmth when it was embodied ($U = 1407$, $Z = -0.171$, $p = 0.864$, $r = 0.016$; $M_{\text{selfish-E}} = 3.218$, $SD = 0.833$) and disembodied ($M_{\text{selfish-D}} = 3.171$, $SD = 1.110$). However, the attribution of warmth to the prosocial agent was significantly different when it was embodied ($U = 1089$, $Z = -2.015$, $p = 0.044$, $r = 0.192$; $M_{\text{prosocial-E}} = 4.162$, $SD = 1.115$) and disembodied ($M_{\text{prosocial-D}} = 4.703$, $SD = 1.147$). It reveals the warmth of the prosocial agent increased when that agent was disembodied, increasing as well the discrepancy with the warmth attributed to the selfish agent.

No significant two-way interaction effect was observed between strategy and the game result ($F(1, 106) = 1.243$, $p = 0.267$, $r = 0.110$) and no three-way interaction effect was observed between strategy, embodiment and game result ($F(1, 106) = 1.964$, $p = 0.164$, $r = 0.134$).

2) **Competence:** The competence dimension had a good internal consistency ($\alpha = 0.809$). A small-sized main effect was observed for embodiment ($F(1, 106) = 9.376$, $p = 0.003$, $r = 0.285$) indicating that the disembodied agents' competence was significantly ranked higher than the embodied agents' competence ($M_{\text{Disembodied}} = 4.827$, $SD = 1.231$; $M_{\text{Embodied}} = 4.342$, $SD = 1.093$). A small-size main effect was also observed for game result ($F(1, 106) = 4.405$, $p = 0.038$, $r = 0.200$), indicating that participants who won the game significantly ranked the competence of the agents higher than participants who lost the game ($M_{\text{Victory}} = 4.672$, $SD = 1.320$; $M_{\text{Loss}} = 4.367$, $SD = 1.144$). Additionally, a large-sized main effect of strategy could be observed ($F(1, 106) = 145.730$, $p < 0.001$, $r = 0.761$) indicating that the competence attributed to the prosocial was significantly higher than the competence attributed to the selfish over all conditions ($M_{\text{prosocial}} = 5.344$, $SD = 1.318$; $M_{\text{selfish}} = 3.701$, $SD = 1.208$).

A large-size two-way interaction effect was also observed (see Figure 2c) between strategy and embodiment ($F(1, 106) = 41.909$, $p = 0.001$, $r = 0.532$), revealing that the competence attributed to embodied and disembodied agents was affected differently by their strategies. More precisely, the competence attributed to the selfish agent was significantly lower when it was disembodied ($U = 1049$, $Z = -2.373$, $p = 0.018$, $r = 0.225$; $M_{\text{selfish-D}} = 3.378$, $SD = 1.357$) than when it was embodied ($M_{\text{selfish-E}} = 3.894$, $SD = 1.075$). However, the attribution of competence to a prosocial agent signifi-

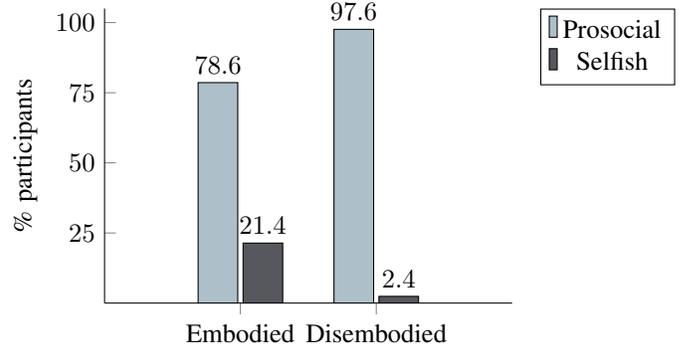


Fig. 3: Association between the embodiment and the number of participants that preferred each agent ($p = 0.005$).

cantly increased when it was disembodied ($U = 449.5$, $Z = -5.974$, $p < 0.001$, $r = 0.570$; $M_{\text{prosocial-D}} = 6.276$, $SD = 1.104$) compared to when it was embodied ($M_{\text{prosocial-E}} = 4.790$, $SD = 1.111$).

A medium-sized two-way interaction effect between strategy and game result was also found ($F(1, 106) = 11.633$, $p = 0.001$, $r = 0.315$). This indicates that the competence values of the two strategies, prosocial and selfish, were affected differently by the game result. When participants won the game, the competence attributed to the prosocial was significantly higher than to the selfish ($Z = -4.065$, $p < 0.001$, $r = 0.539$). However, that difference was even more noticeable when they lost the game ($Z = -5.999$, $p < 0.001$, $r = 0.816$).

Finally, no significant three-way interaction effect was observed between strategy, embodiment and game result ($F(1, 106) = 1.298$, $p = 0.257$, $r = 0.110$).

3) **Discomfort:** The items of the discomfort dimension had a good internal consistency ($\alpha = 0.817$). No main effects were observed for embodiment ($F(1, 106) = 0.004$, $p = 0.949$, $r < 0.032$) nor game result ($F(1, 106) = 2.697$, $p = 0.104$, $r = 0.158$). Nevertheless, a small-size main effect of strategy was found ($F(1, 106) = 59.160$, $p < 0.001$, $r = 0.598$), indicating that the discomfort attributed to the prosocial was lower than the discomfort attributed to the selfish ($M_{\text{selfish}} = 2.946$, $SD = 1.364$; $M_{\text{prosocial}} = 1.846$, $SD = 1.042$).

Additionally, we neither observed a significant two-way interaction effects between strategy and embodiment ($F(1, 106) = 0.555$, $p = 0.458$, $r = 0.071$), nor between strategy and game result ($F(1, 106) = 2.008$, $p = 0.159$, $r = 0.138$), nor a significant three-way interaction effect between strategy, embodiment and game result ($F(1, 106) = 0.007$, $p = 0.933$, $r < 0.032$).

4) **Partner Selection:** Using a Fisher's Exact test, we found a significant association between the level of embodiment and the number of times participants selected each agent for future games (Figure 3, $\chi^2(1) = 7.558$, $p = 0.005$). In the embodied condition, 78.6% of participants (55 out of 70) selected the prosocial robot for future games, while in the disembodied condition, that preference was even more salient with 97.6% of participants (40 out of 41) choosing the prosocial agent.

C. Cooperation Rate

Finally, through a Mann-Whitney U Test, we analyzed the actions performed by the participants when playing the game alongside embodied and disembodied agents. No significant differences were observed ($Z = -0.565, p = 0.572, r = 0.055$), meaning that the amount of times players decided to cooperate did not significantly varied when playing alongside embodied ($M = 2.400, SD = 0.875$) or disembodied agents ($M = 2.439, SD = 0.950$). Overall, the average cooperation rate was 2 times out of 4.

VI. DISCUSSION

In **H1**, we have hypothesised that *participants would perceive the team more positively in the embodied condition*. The results of our study only partially validate this hypothesis, since a significant difference in the predicted direction was only found for group identification, but not for group trust. It seems the degree of identification towards the team can be positively affected by the embodied presence of the robots and the richer social interaction they might provide. The group trust, on the other hand, seems to have been less affected by our manipulation of embodiment.

Due to the close relation between these two measures, which are usually correlated, it was surprising that their results were not aligned. However, the subjective nature of these group perceptions can reflect both the execution of the task or a certain degree of satisfaction with the social interaction of the team. Therefore, our findings suggest group identification might be more associated with the social interaction, while group trust might reflect more the actions performed for the task. We speculate that in our collaborative game, due to the uncertainty factor, participants have evaluated the task execution by considering the strategies of all members. For instance, the fact that the team had an unconditional defector in both conditions could have damaged participants' trust.

Furthermore, considering the results we obtained in these group measures, we would like to highlight their importance as an additional consideration for future research on human-robot group interactions. Due to the known relation between trust and the performance of a robot [13], it would be interesting to also analyse the relation between group trust and the perception of competence of the team as a unified entity. In our user study, this group measure can be hard to infer only from the individual perceptions of competence of each agent. Not only because we have not assessed participants' perceptions of self competence, but also because this group perception might not necessarily follow a linear structure. In the future, we would like to further explore this idea by considering more group measures.

We predicted in **H2a** that *regardless of the embodiment, the perception of the prosocial agent will be more positive than the perception of the selfish agent*. We found support for this hypothesis as the prosocial agent was indeed rated with higher levels of both warmth and competence, and lower levels of discomfort when compared to the selfish agent.

In **H2b**, we have hypothesised that *the perception of the prosocial agent will be more positive in the embodied condition*. Our results do not support this hypothesis and we actually verified the opposite situation. The prosocial agent was rated with higher levels of both warmth and competence when it was disembodied than when it was embodied.

A similar situation occurred with **H2c**, in which we hypothesised that *the perception of the selfish agent will be more negative in the embodied condition*. The attribution of competence revealed an opposite result to what we have predicted, with the embodied selfish agent being rated as more competent than the disembodied selfish agent. Regarding the other social attributes, warmth and discomfort no significant differences were found.

Our hypotheses regarding the perceptions of the agents, both **H2b** and **H2c**, were based on the assumption that the social behaviour of the embodied robots would increase the expectations for them to have an adequate behaviour regarding social and moral norms. This assumption is supported by the actual behaviour people display around physically-embodied agents compared to virtually-embodied agents [3, 18]. As a result, we expected that the higher degree of embodiment would actually highlight the strategies of the agents, by improving the perception of a prosocial strategy and by worsening the perception of a selfish strategy. On the contrary, our results suggest that the robotic embodiment reduced, rather than increased, the gap between the perceptions of the two strategies. It seems that the lack of an embodiment made participants more focused on the agents' strategies and, conversely, in the embodied condition participants did not pay as much attention to the strategies due to other factors being considered. In fact, the social behaviours of the embodied agents were designed in a neutral tone so that their only difference was the strategy to play the game. However, it might be the case that participants felt that there was a disconnect between their strategy and the way they utilized its embodiment to convey that strategy. Another possible interpretation is that the embodiment made the actual social interaction more salient over the task-related aspects, as if a new dimension in their evaluation came into play. Again, as the robotic agents, both the prosocial and the selfish, employed similar non-verbal behaviours, it might have in turn decreased the difference between the perceptions of their social attributes.

In **H3**, we have predicted that *participants will be more prosocial in the embodied condition*. We expected that either the higher degree of embodiment would inhibit participants' selfishness, or the lower degree of embodiment would promote it. However, our results did not support this hypothesis.

On the one hand, we believe that there might be a relationship between this result and the one obtained for group trust. In other words, the fact that participants equally trusted their teams in both embodiment conditions, might justify the absence of significant differences in the cooperation rate. Again, the presence of an unconditional defector on the team, the selfish agent, might have played a significant role in the strategy of the participants. On the other hand, we would like

to point out that our experimental scenario contained only four decision points, which might be insufficient to notice significant differences. Future investigations using similar collective dilemmas should contain more rounds to clarify this issue.

Finally, regarding the independent variable of the outcome, its main effect on group identification and competence, as well as the interaction effect with the agents' strategy are both inline with previous findings.

VII. CONCLUSION

This paper explores the role of embodiment in collective dilemmas with artificial agents. In this type of task, opting for cooperation usually incurs some cost or loss for oneself and, therefore, defection might seem more enticing. Within a wider research vision of creating agents that enhance human cooperation [32], this paper is focused on the effect that the embodiment of those agents can have. In particular, we compared robots that use several embodied modalities (i.e., gaze, posture, animations) with disembodied agents that socially interact only with textual comments. Such comparison is not only novel, but also relevant considering the widespread use of intelligent disembodied voice assistants (e.g., Cortana, Google's assistant, Siri). *What is the impact of adding embodied affordances to interactive social agents in cooperative tasks?*

To analyse this question, we conducted an empirical study where participants were asked to form a team with two autonomous agents. One of the agents embraced a prosocial strategy (i.e., by always choosing to cooperate), while the other adopted a selfish strategy (i.e., by always opting to defect). We also manipulated the level of embodiment in two extreme levels: a robot and a disembodied agent. We analysed how participants (1) evaluated the team, (2) perceived each agent, and (3) their degree of cooperation towards the team.

Firstly, regarding group measures, participants identified themselves more with the team of embodied agents than with the team of disembodied agents. However, no significant difference was found on the trust towards the team between the two embodiment conditions. Secondly, in terms of the agents' perceptions, the main effect of the strategy in the three social attributes revealed that participants clearly distinguished between the most and the least collaborative strategies, by positively rating the prosocial agent and by negatively rating the selfish one. Nevertheless, when comparing the social attributes across different levels of embodiment, the results suggest that the strategies of the embodied agents were partially masked by their social behaviours. Thirdly and finally, our results did not show a significant difference from our embodiment manipulation on participants' degree of cooperation).

The following considerations for human-robot teams can be derived from the obtained results. Firstly, a misalignment of valences between the actions of an agent and its social behaviours can change the perceptions of those agents. In our user study, the social behaviour of both robotic agents was similarly neutral in valence, which might have caused a negative impact on the perception of a prosocial agent and improved the perception of a selfish agent. This constitutes an

important consideration for the design of social behaviour in embodied agents. As future work, it would be interesting to explore if an agent executing prosocial actions but displaying negative social behaviours, rather than neutral ones, would be rated more negatively than an agent executing selfish actions but displaying positive social behaviour.

A second consideration we can draw from our results is that the embodied affordances of the agents seem to lead people to consider additional aspects during the interaction. Those aspects seem to be related to the social interaction that is beyond task-related actions, and can be established through those embodied affordances, i.e. gaze or facial expressions. According to our results, the embodiment may drive both positive and negative facets of human-robot interaction. On the one hand, it showed a positive impact on the degree of identity with the team. Group identification constitutes an important process of group settings and is reported, in the social sciences, as being further associated with other relevant group measures, such as team performance [41]. On the other hand, it was also able to partially mitigate the strategies that each agent took in the collective dilemma. If we focus on the perceptions of the selfish agent, it raises a negative perspective on the impact of embodiment. The presence of embodied features seems to have decreased the emphasis on its non-cooperative actions and to have improved how humans perceived that agent. This concern is particularly relevant if we consider that, in the future, socially intelligent agents can portray the selfish intentions of their owners, e.g. an individual or a company. By having these considerations in mind, practitioners can design adequate affordances and behaviours of social agents according to the goals and nature of the task.

One limitation in the results we obtained is that they might not generalise to virtual agents as our embodied condition refers to a robot, i.e. physically-embodied agent. To address this issue, a follow-up study should accommodate another condition with virtually-embodied agents. A final limitation that is common in the analysis of group interactions is the nested design between the embodiment and the strategy. The perception of an agent with a certain strategy can differ according to the strategy of the other agent on the team. Considering in our design we only had opposite strategies, to address this issue in future studies, other conditions where both agents employ the same strategy would be necessary.

Overall, embodied interactions seem to introduce more aspects into the evaluation of agents, possibly related to a higher social presence and/or a richer social interaction. Although they might improve how humans relate with agents in team settings, if these social aspects can positively mask questionable or even immoral decisions, a dark side of embodiment may emerge.

ACKNOWLEDGMENTS

This work was supported by national funds of Fundação para a Ciência e a Tecnologia (UIDB/50021/2020, SFRH/BD/118031/2016 and SFRH/BD/143460/2019), and the ANIMATAS Project (H2020-MSCA-ITN-2017/765955).

REFERENCES

- [1] Kathleen Allen and Richard Bergin. Exploring trust, group satisfaction, and performance in geographically dispersed and co-located university technology commercialization teams. In *In Proceedings of the NCHIA 8th Annual Meeting: Education that Works*, pages 18–20, 2004. URL <https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.201.407>.
- [2] Ofra Amir, David G Rand, et al. Economic games on the internet: The effect of \$1 stakes. *PloS one*, 7(2), 2012. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3283743/>.
- [3] Wilma A Bainbridge, Justin W Hart, Elizabeth S Kim, and Brian Scassellati. The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics*, 3(1):41–52, 2011. URL <https://link.springer.com/article/10.1007/s12369-010-0082-7>.
- [4] Lawrence W Barsalou, Paula M Niedenthal, Aron K Barbey, and Jennifer A Ruppert. Social embodiment. *Psychology of learning and motivation*, 43:43–92, 2003. URL <https://www.sciencedirect.com/science/article/pii/S0079742103010119?via%3Dihub>.
- [5] Christoph Bartneck. Interacting with an embodied emotional character. In *Proceedings of the 2003 international conference on Designing pleasurable products and interfaces*, pages 55–60. ACM, 2003. URL <https://dl.acm.org/doi/10.1145/782896.782911>.
- [6] Christoph Bartneck, Andreas Duenser, Elena Moltchanova, and Karolina Zawieska. Comparing the similarity of responses received from studies in amazon’s mechanical turk to studies conducted online and with direct recruitment. *PloS one*, 10(4), 2015.
- [7] Colleen M Carpinella, Alisa B Wyman, Michael A Perez, and Steven J Stroessner. The robotic social attributes scale (RoSAS): development and validation. In *ACM/IEEE Int. Conf. on Human-Robot Interaction*, 2017. URL <https://dl.acm.org/doi/10.1145/2909824.3020208>.
- [8] Filipa Correia, Samuel F Mascarenhas, Samuel Gomes, Patrícia Arriaga, Iolanda Leite, Rui Prada, Francisco S Melo, and Ana Paiva. Exploring prosociality in human-robot teams. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 143–151. IEEE, 2019. URL <https://ieeexplore.ieee.org/abstract/document/8673299>.
- [9] Kerstin Dautenhahn, Bernard Ogden, and Tom Quick. From embodied to socially embedded agents—implications for interaction-aware robots. *Cognitive Systems Research*, 3(3):397–428, 2002. URL <https://www.sciencedirect.com/science/article/pii/S1389041702000505>.
- [10] Eric Deng, Bilge Mutlu, Maja J Matarić, et al. Embodiment in socially interactive robots. *Foundations and Trends® in Robotics*, 7(4):251–356, 2019. URL <https://www.nowpublishers.com/article/Details/ROB-056>.
- [11] Juan Fasola and Maja J Matarić. A Socially Assistive Robot Exercise Coach for the Elderly. *J. Hum.-Robot Interact.*, 2(2):3–32, June 2013. doi: 10.5898/JHRI.2.2. Fasola. URL <https://doi.org/10.5898/JHRI.2.2.Fasola>.
- [12] Marina Fridin and Mark Belokopytov. Embodied robot versus virtual agent: Involvement of preschool children in motor task performance. *International Journal of Human-Computer Interaction*, 30(6):459–469, 2014. URL <https://www.tandfonline.com/doi/abs/10.1080/10447318.2014.888500>.
- [13] Peter A Hancock, Deborah R Billings, Kristin E Schaefer, Jessie YC Chen, Ewart J De Visser, and Raja Parasuraman. A meta-analysis of factors affecting trust in human-robot interaction. *Human factors*, 53(5):517–527, 2011. URL <https://journals.sagepub.com/doi/10.1177/0018720811417254>.
- [14] Laura Hoffmann and Nicole C Krämer. Investigating the effects of physical and virtual embodiment in task-oriented and conversational contexts. *International Journal of Human-Computer Studies*, 71(7-8): 763–774, 2013. URL <https://www.sciencedirect.com/science/article/abs/pii/S1071581913000591>.
- [15] James Kennedy, Paul Baxter, and Tony Belpaeme. Comparing robot embodiments in a guided discovery learning interaction with children. *International Journal of Social Robotics*, 7(2):293–308, 2015. URL <https://link.springer.com/article/10.1007/s12369-014-0277-4>.
- [16] Soheil Keshmiri, Hidenobu Sumioka, Ryuji Yamazaki, and Hiroshi Ishiguro. Differential Effect of the Physical Embodiment on the Prefrontal Cortex Activity as Quantified by Its Entropy. *Entropy*, 21(9):875, 2019. URL <https://www.mdpi.com/1099-4300/21/9/875>.
- [17] Cory D Kidd and Cynthia Breazeal. Effect of a robot on user perceptions. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, volume 4, pages 3559–3564. IEEE, 2004. URL <https://ieeexplore.ieee.org/abstract/document/1389967>.
- [18] Sara Kiesler, Aaron Powers, Susan R Fussell, and Cristen Torrey. Anthropomorphic interactions with a robot and robot-like agent. *Social Cognition*, 26(2): 169–181, 2008. URL <https://guilfordjournals.com/doi/10.1521/soco.2008.26.2.169>.
- [19] Kangsoo Kim, Arjun Nagendran, Jeremy N Bailenson, Andrew Raij, Gerd Bruder, Myungho Lee, Ryan Schubert, Xin Yan, and Gregory F Welch. A Large-Scale Study of Surrogate Physicality and Gesturing on Human-Surrogate Interactions in a Public Space. *Frontiers in Robotics and AI*, 4:32, 2017. URL <https://www.frontiersin.org/articles/10.3389/frobt.2017.00032/full>.
- [20] Dimosthenis Kontogiorgos, Andre Pereira, Olle Andersson, Marco Koivisto, Elena Gonzalez Rabal, Ville Vartiainen, and Joakim Gustafson. The effects of anthropomorphism and non-verbal social behaviour in virtual assistants. In *Proceedings of the 19th ACM International*

- Conference on Intelligent Virtual Agents*, pages 133–140. ACM, 2019. URL <https://dl.acm.org/doi/abs/10.1145/3308532.3329466>.
- [21] Hatice Kose-Bagci, Ester Ferrari, Kerstin Dautenhahn, Dag Sverre Syrdal, and Chrystopher L Nehaniv. Effects of embodiment and gestures on social interaction in drumming games with a humanoid robot. *Advanced Robotics*, 23(14):1951–1996, 2009. URL <https://www.tandfonline.com/doi/abs/10.1163/016918609X12518783330360>.
- [22] Sonya S Kwak, Yunkyung Kim, Eunho Kim, Christine Shin, and Kwangsu Cho. What makes people empathize with an emotional robot?: The impact of agency and physical embodiment on human empathy for a robot. In *2013 IEEE RO-MAN*, pages 180–185. IEEE, 2013. URL <https://ieeexplore.ieee.org/abstract/document/6628441>.
- [23] Colin Wayne Leach, Martijn Van Zomeren, Sven Zebel, Michael LW Vliek, Sjoerd F Pennekamp, Bertjan Doosje, Jaap W Ouwerkerk, and Russell Spears. Group-level self-definition and self-investment: a hierarchical (multicomponent) model of in-group identification. *Journal of personality and social psychology*, 95(1):144, 2008. URL <https://psycnet.apa.org/doiLanding?doi=10.1037%2F0022-3514.95.1.144>.
- [24] Kwan Min Lee, Younbo Jung, Jaywoo Kim, and Sang Ryong Kim. Are physically embodied social agents better than disembodied social agents?: The effects of physical embodiment, tactile interaction, and people’s loneliness in human–robot interaction. *International journal of human-computer studies*, 64(10):962–973, 2006. URL <https://www.sciencedirect.com/science/article/abs/pii/S1071581906000784>.
- [25] Daniel Leyzberg, Samuel Spaulding, Mariya Toneva, and Brian Scassellati. The physical presence of a robot tutor increases cognitive learning gains. In *Proceedings of the annual meeting of the cognitive science society*, volume 34, 2012. URL <https://escholarship.org/uc/item/7ck0p200>.
- [26] Jamy Li. The benefit of being physically present: A survey of experimental works comparing copresent robots, telepresent robots and virtual agents. *International Journal of Human-Computer Studies*, 77:23–37, 2015. URL <https://www.sciencedirect.com/science/article/abs/pii/S107158191500004X>.
- [27] Mike Lighthart and Khiet P Truong. Selecting the right robot: Influence of user attitude, robot sociability and embodiment on user preferences. In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 682–687. IEEE, 2015. URL <https://ieeexplore.ieee.org/abstract/document/7333598>.
- [28] Rosemarijn Looije, Anna van der Zalm, Mark A Neerinx, and Robbert-Jan Beun. Help, I need some body the effect of embodiment on playful learning. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, pages 718–724. IEEE, 2012. URL <https://ieeexplore.ieee.org/abstract/document/6343836>.
- [29] Samuel Mascarenhas, Manuel Guimarães, Rui Prada, João Dias, Pedro A Santos, Kam Star, Ben Hirsh, Ellis Spice, and Rob Kommeren. A virtual agent toolkit for serious games developers. In *2018 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–7. IEEE, 2018. URL <https://ieeexplore.ieee.org/abstract/document/8490399>.
- [30] Winter Mason and Siddharth Suri. Conducting behavioral research on Amazon’s Mechanical Turk. *Behavior research methods*, 44(1):1–23, 2012. URL <https://link.springer.com/article/10.3758/s13428-011-0124-6>.
- [31] Ali Mollahosseini, Hojjat Abdollahi, Timothy D Sweeny, Ron Cole, and Mohammad H Mahoor. Role of embodiment and presence in human perception of robots’ facial cues. *International Journal of Human-Computer Studies*, 116:25–39, 2018. URL <https://www.sciencedirect.com/science/article/abs/pii/S1071581918301745>.
- [32] Ana Paiva, Fernando P Santos, and Francisco C Santos. Engineering pro-sociality with autonomous agents. In *Thirty-Second AAI Conference on Artificial Intelligence*, 2018. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/viewPaper/16799>.
- [33] André Pereira, Carlos Martinho, Iolanda Leite, and Ana Paiva. iCat, the chess player: the influence of embodiment in the enjoyment of a game. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 3*, pages 1253–1256. International Foundation for Autonomous Agents and Multiagent Systems, 2008. URL <https://dl.acm.org/doi/10.5555/1402821.1402844>.
- [34] Rolf Pfeifer and Christian Scheier. *Understanding intelligence*. 2001. URL <https://dl.acm.org/doi/book/10.5555/559357>.
- [35] Tom Quick, Kerstin Dautenhahn, Chrystopher L Nehaniv, and Graham Roberts. On bots and bacteria: Ontology independent embodiment. In *European Conference on Artificial Life*, pages 339–343. Springer, 1999. URL https://link.springer.com/chapter/10.1007/3-540-48304-7_45.
- [36] Tiago Ribeiro, André Pereira, Eugenio Di Tullio, and Ana Paiva. The sera ecosystem: Socially expressive robotics architecture for autonomous human-robot interaction. In *2016 AAI Spring Symposium Series*, 2016. URL <https://www.aaai.org/ocs/index.php/SSS/SSS16/paper/viewPaper/12748>.
- [37] Fernando P Santos, Samuel F Mascarenhas, Francisco C Santos, Filipa Correia, Samuel Gomes, and Ana Paiva. Outcome-based Partner Selection in Collective Risk Dilemmas. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1556–1564. International Foundation for Autonomous Agents and Multiagent Systems, 2019. URL <https://dl.acm.org/doi/10.5555/3306127.3331873>.
- [38] Isabella Seeber, Eva Bittner, Robert O Briggs, Triparna de Vreede, Gert-Jan De Vreede, Aaron Elkins, Ronald

- Maier, Alexander B Merz, Sarah Oeste-Reiß, Nils Randrup, et al. Machines as teammates: A research agenda on ai in team collaboration. *Information & management*, 57(2):103174, 2020.
- [39] Elena Márquez Segura, Michael Kriegel, Ruth Aylett, Amol Deshmukh, and Henriette Cramer. How do you like me in this: User embodiment preferences for companion agents. In *International Conference on Intelligent Virtual Agents*, pages 112–125. Springer, 2012. URL https://link.springer.com/chapter/10.1007/978-3-642-33197-8_12.
- [40] Stela H Seo, Denise Geiskkovitch, Masayuki Nakane, Corey King, and James E Young. Poor thing! Would you feel sorry for a simulated robot? A comparison of empathy toward a physical and a simulated robot. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 125–132. IEEE, 2015. URL <https://ieeexplore.ieee.org/abstract/document/8520653>.
- [41] Anit Somech, Helena Syna Desivilya, and Helena Lidogoster. Team conflict management and team effectiveness: The effects of task interdependence and team identification. *Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior*, 30(3):359–378, 2009. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/job.537>.
- [42] Sam Thellman, Annika Silvervarg, Agneta Gulz, and Tom Ziemke. Physical vs. virtual agent embodiment and effects on social interaction. In *International Conference on Intelligent Virtual Agents*, pages 412–415. Springer, 2016. URL https://link.springer.com/chapter/10.1007/978-3-319-47665-0_44.
- [43] Kyle A Thomas and Scott Clifford. Validity and Mechanical Turk: An assessment of exclusion methods and interactive experiments. *Computers in Human Behavior*, 77:184–197, 2017. URL <https://www.sciencedirect.com/science/article/pii/S074756321730506X>.
- [44] Joshua Wainer, David J Feil-Seifer, Dylan A Shell, and Maja J Matarić. Embodiment and human-robot interaction: A task-based perspective. In *RO-MAN 2007-The 16th IEEE International Symposium on Robot and Human Interactive Communication*, pages 872–877. IEEE, 2007. URL <https://ieeexplore.ieee.org/abstract/document/4415207>.
- [45] Tom Ziemke. What’s that thing called embodiment? In *Proceedings of the 25th Annual Conference of the Cognitive Science Society*, 2003. URL <https://escholarship.org/uc/item/60w6v9jz>.