Noname manuscript No.
(will be inserted by the editor)

When A Robot Is Your Teammate

Cohesive Alliances With Social Robots In Multi-party Settings

Filipa Correia · Francisco S. Melo · Ana Paiva

Received: date / Accepted: date

Abstract Creating effective teamwork between humans and robots involves not only addressing their performance as a team but also sustaining the quality and sense of unity among teammates, also known as cohesion. This paper explores the research problem of: how can we endow robotic teammates with social capabilities to improve the cohesive alliance with humans? By defining the concept of a human-robot cohesive alliance in the light of the multidimensional construct of cohesion from the social sciences, we propose to address this problem through the idea of multifaceted human-robot cohesion. We present our preliminary effort from previous works to examine each of the five dimensions of cohesion: social, collective, emotional, structural and task. We finish the paper with a discussion on how human-robot cohesion contributes to the key questions and ongoing challenges of creating robotic teammates. Overall, cohesion in human-robot teams might be a key factor to propel team performance and it should be considered in the design, development and evaluation of robotic teammates.

Keywords Human-Robot Interaction · Teamwork · Cohesion · Collaboration · Multi-party

1 Introduction

As robotic systems become part of our lives in various domains or environments, e.g. domestic (Christensen, 2003), industrial (Guiochet et al., 2017), public spaces (Jensen et al., 2005; Kanda et al., 2009), education (Belpaeme et al., 2018), health (Breazeal, 2011), they are naturally expected to engage in collaborative settings or even integrate into teams along with humans (Groom and Nass, 2007; Seeber et al., 2020). In human-robot teams, robots share common goals with humans and they need to act together to achieve those goals. In those situations, *how can*

Filipa Correia[1,2] · Francisco S. Melo[1] · Ana Paiva[1]
[1]INESC-ID, Instituto Superior Técnico, Universidade de Lisboa
[2]ITI, LARSyS, Instituto Superior Técnico, Universidade de Lisboa
Lisbon, Portugal
E-mail: filipacorreia@tecnico.ulisboa.pt

robotic teammates be effective collaborators? An immediate question that might occur is: what exactly is the *effectiveness* of a human-robot team, and how can it be measured? There might be several different ways of assessing effectiveness, depending on the purpose or the goals of the team. One possibility is maximising the performance achieved by the team, by looking at the final outcome. There are, however, other possibilities such as examining other aspects or factors that sustain the quality of the team and may, in turn, mediate performance (De Visser et al., 2020).

We took inspiration from the social sciences, and we looked at what characterises an interpersonal team. For instance, the social psychologist Forsyth refers to teams as "unified, cohesive groups" (Forsyth, 1990). He pointed out that members of a team are usually committed to a *common goal* in which individual success is only a consequence of collective success. There is a strong *interdependence* of the members' efforts that requires a *coordinated* interaction. In terms of *structure*, teams usually present highly adaptive skills that allow a revision of their norms and procedures in order to improve their functioning. Finally, the *cohesion* of a team is also a prominent characteristic, which is pointed by many authors as the most important aspect of groups and teams (Lott and Lott, 1965; Cartwright, 1968; Piper et al., 1983; Cota et al., 1995; Carron and Brawley, 2000; Eys et al., 2009).

Cohesion is the integrity, solidarity and unity that maintains a group together or in a *cohesive alliance*. It is associated with satisfaction, performance, and productivity. For instance, Mullen and Cooper found an interesting bidirectional relationship between cohesion and performance, which suggests that not only do cohesive teams tend to perform better but also the success of the team leads to increased cohesion (Mullen and Copper, 1994). The multi-dimensionality of the cohesion construct suggests that there are five possible courses that lead to a *cohesive alliance* (Forsyth, 1990) (see Figure 1): the task, the social bonds, the collective entity, and the structure and emotions. *Task cohesion* emerges from the shared commitment towards a common goal, while *social cohesion* develops from the attractions among the members. The identification of the members with the group and the degree of belongingness constitutes *collective cohesion*. The *structural cohesion* derives from the norms, roles and relationships that link the members of the group. Finally, *emotional cohesion* is the intensity of the members to express group feelings.

We are particularly interested in this type of satisfactory and effective coalition in mixed teams of humans and robots. Considering both how humans establish cohesive alliances in their interpersonal groups or teams, and the multi-dimensionality of this construct, we define a cohesive alliance between humans and robots as *a coalition in which the relation between group members, both humans and robots, emerges from at least one of the dimensions of cohesion and results in a shared sense of unity by all group members*. In the light of this definition, we argue that the design and evaluation of robotic teammates can take into account these five dimensions in order to establish human-robot cohesion or a cohesive alliance between humans and robots. This paper sets out a novel exploration of multifaceted human-robot cohesion.

Another innovative aspect of our work is the exploration of human-robot teamwork in multi-party small groups. The structure of human-robot teams might vary from dyadic teams to larger coalitions where several humans and several robots act

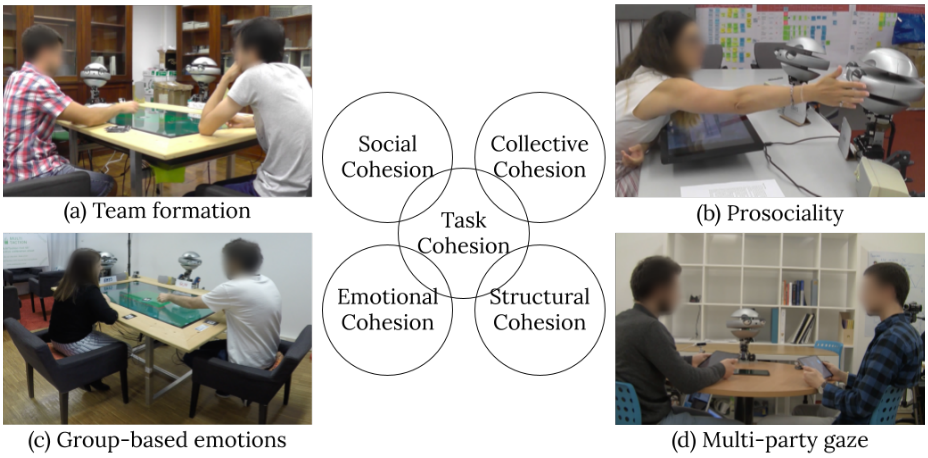

Fig. 1 Addressing the multi-dimensionality of the cohesion construct.

together towards a common goal. Most Human-Robot Interaction (HRI) literature on the topic of human-robot teamwork is focused on dyadic interactions (Hoffman and Breazeal, 2007; Dragan et al., 2015; Huang and Mutlu, 2016; Chang et al., 2018; Shayganfar et al., 2019; Hoffman, 2019). Few examples explored teamwork and collaboration within multi-party groups of humans and robots, i.e. groups with more than one person and/or more than one robot (Jung et al., 2013; Fraune et al., 2017b; Salomons et al., 2018; Traeger et al., 2020). This paper specifically contributes to the development of robotic teammates in multi-party interactions of small groups, in our use cases, the groups have either three or four elements. Hence the general research problem we address is: *how can we endow a robotic teammate with social capabilities to improve the cohesive alliance in a multi-party setting with humans?*

In the following sections, we start in Section 2 by reviewing the state-of-art on the topics of teamwork and collaboration, multi-party interactions, and cohesion. Then, we propose that cohesion can be used as a framework to create robotic teammates, presented in Section 3, and we examine four use cases of our previous work to support this framework in Sections 4-7. Each use case specifically explores one of the above-mentioned dimensions of cohesion, as Figure 1 shows. Specifically, Section 4 focuses on membership preferences and team formation, which are related to aspects of social cohesion. In Section 5, we describe a second use case exploring how can robotic teammates portray different levels of collective cohesion. The third use case, which is described in Section 6, focuses on the robot's expression of group-based emotions, which is related to emotional cohesion. The last use case, which is presented in Section 7, explores different perceptual capabilities in robotic teammates that correspond to different structural team configurations. Additionally, task cohesion is also inherently present across all the projects, as we have chosen scenarios where humans are asked to form a team with robots.

Throughout the paper, we argue that the three core stages to create robotic teammates hold its (1) design, (2) development, and (3) evaluation. As a result, the following sections are organised according to these three central levels, including the discussion in Section 8. Overall, this paper contributes with a framework to create cohesive alliances between humans and social robots in team settings. Moreover, it discusses four related use cases that illustrate and support the feasibility of the proposed framework. Finally, the paper ends with a discussion on the future challenges that aim at encouraging researchers to pursue some of the possible avenues in order to create cohesive alliances between humans and robotic teammates.

2 Related Work

This paper lies at the intersection of three topics in HRI, which are teamwork, multi-party interactions and cohesion.

2.1 Teamwork and Collaboration

One of the main concerns in human-robot teamwork is how to support joint work and create uent and coordinated actions (Ho man and Breazeal, 2004). To that end, important capabilities can be an adaptation to the human teammate (Ho man and Breazeal, 2007; Shah et al., 2011) or anticipatory behaviours (Huang and Mutlu, 2016; Iqbal and Riek, 2017). Other relevant social behaviours that enhance human-robot teamwork include the presence of backchanneling (Jung et al., 2013), gaze (Kshirsagar et al., 2020), or other non-verbal social cues (Breazeal et al., 2005). Nevertheless, robots can also assume additional roles in these types of interactions, for instance, robots can be a proxy of human teammates through telepresence (Stoll et al., 2018), or mediate interpersonal con icts among other human teammates (Shen et al., 2018).

2.2 Multi-party Interactions

In multi-party interactions between humans and robots, there are some contradicting ndings suggesting that well-known group phenomena in interpersonal groups might or might not apply to human-robot groups, such as the case of the discontinuity e ect (Chang et al., 2012; Fraune et al., 2019b), the conformity e ect (Brandstetter et al., 2014; Salomons et al., 2018, 2021), or the ingroup versus outgroup bias (Fraune et al., 2017b; Steain et al., 2019). Beyond explicit comparisons with human-human groups, when robots adopt new roles such as being a facilitator or a mediator, they can have a positive impact on the interaction among humans according to their social capabilities. Many examples showed robotic mediators can foster the inclusion among human group members through verbal support (Sebo et al., 2020), the portrayal of vulnerability (Strohkorb Sebo et al., 2018; Traeger et al., 2020), gaze behaviours (Mutlu et al., 2012; Gillet et al., 2021), attentive postures (Tennent et al., 2019), resource allocation (Jung et al., 2018). By exploring other roles for robots in multi-part settings, researchers examined other types of social behaviours, for instance, displaying empathy according to the emotional climate of the group (Alves-Oliveira et al., 2019), or adopting di erent navigation strategies in a crowded environment (Mavrogiannis et al., 2019).

2.3 Cohesion

We identi ed four lines of research within the literature that explores cohesion in human-robot interactions. First, robotic facilitators try to enhance the cohesion of an interpersonal group. Examples include a narrative companion robot with the capability to make social narrative links (Uchida et al., 2020), and a moderation algorithm explored that could be either performance equalizing or performance reinforcing (Short and Mataric, 2017). Second, the exploration of how cohesive groups of robots are perceived (Fraune et al., 2017a). Third, how does the cohesion of interpersonal groups a ect their interaction with a robot (Fraune et al., 2019a). Fourth, how to develop social behaviour for robots that accounts for cohesion. One example is the work by Sebo et al. exploring the impact of supporting interpersonal cohesiveness through verbal comments, compared to supporting task performance (Strohkorb et al., 2016). Sathyamoorthy et al. proposed a novel metric for estimating cohesion and applied it in a navigation method that maintains the perceived group cohesion (Sathyamoorthy et al., [n.d.]). Similarly, Landol and Dragan introduced an algorithm for an autonomous car that considers the driving style of the surrounding cars (Landol and Dragan, 2018).

However, Abrams and Rosenthal-von der Pütten have recently proposed the Identi cation-Cohesion-Entitativity (I-C-E) framework, which highlights important considerations to study cohesion in interactions with robots. Their paper reviews the theoretical foundation to address the concepts of ingroup identi cation, cohesion, and entitativity in HRI (Abrams and Rosenthal-von der Pütten, 2020). The authors clarify the similarities, di erences and relations between these constructs and point out that they have been used interchangeably in the literature.

Our paper bridges the gap among these three topics of teamwork, multi-party interactions and cohesion by exploring how can robotic teammates form cohesive alliances in group interactions. The speci c tasks in which we examine human-robot cohesion include not only teamwork and shared work between humans and robots, but also multi-party interactions with either more than one robot or more than one human. Our approach to addressing this research problem is also novel in its multifaceted vision of cohesion for human-robot interactions.

3 Cohesion as a framework to create robotic teammates

We argue that creating robotic teammates is an iterative process over the feedback loop between: design, development and evaluation. Such theory embraces three core questions: (1) how can we design robotic teammates?; (2) how can we develop robotic teammates?; and (3) how can we evaluate robotic teammates? Although these three questions can be addressed separately, they also have a seamless relationship between them. In particular, any contribution in either design or development is only consummated after performing an adequate evaluation supporting their e ectiveness. The design stage does not require development if researchers use techniques such as Wizard of Oz, in which the robot is being controlled by a human operator. However, addressing the autonomy of robots in the development stage either holds a design contribution or is inspired by previous design contributions. Finally, results from the evaluation stage constitute a reinforcement to

Fig. 2 Cohesion as a framework for the creation of robotic teammates.

iterate and improve the design and the development Figure 2 portrays the described symbiotic relationship between the stages of creating a robotic teammate.

Designing robots is usually associated with the shape, embodiment or other physical features of the robots. However, designing social behaviours for robots is equally important in scenarios that require the robot to interact with humans, as in human-robot teamwork. An ideal design should convey the right social a ordances considering the behaviours that the robot can perform. In the particular case of designing robotic teammates, it is important to explore which social behaviours robots need in order to enhance teamwork, how they express those behaviours, and which embodied features allow those behaviours. As a result, we propose that the two main elements to take into account while designing robotic teammates are both morphology and behaviour.

In terms of development, for robots to act autonomously in human-robot teams, they need to be endowed with three core skills. Firstly, they require perceptive skills to observe dynamic changes in the environment, which include changes not only related to the task but also related to other human teammates. Secondly, robots need cognitive skills to reason about the information they observe in a team context, and about the expected impact of their own actions. Generally, such skills involve an understanding of group/team processes. Thirdly, robots require expressive skills to facilitate teamwork with humans. Therefore, we argue that robotic teammates should hold autonomous capabilities at three di erent levels: perception, cognition and expression.

Finally, the evaluation of robotic teammates involves the tools and methods to perform teamwork between humans and robots and to assess their quality. An important element to consider at this stage is the scenario that should undertake teamwork at the same time it should highlight and elicit the speci c behaviours being evaluated. As a result, di erent evaluation setups require di erent capabilities by the robotic teammates. Another important element is the measures which can be subjective or objective instruments to assess teamwork or the inherent facets of teamwork. The measures should similarly be tailored to assess the speci c behaviours or team processes being evaluated.

Stage	Collective Cohesion	Emotional Cohesion	Social Cohesion	Structural Cohesion	Task Cohesion
Designing Robotic Teammates	- Behaviours that portray belongingness to the team - Physical characteristics or accessories that signal belongingness	- Behaviours that express group feelings	- Behaviours that foster attraction between teammates	- Behaviours that acknowledge roles, norms and relationships in the team	- Behaviours that comply and endorse the common goals
Developing Robotic Teammates	- Perceive social cues of association and dissociation on teammates - Interpret and reason about belongingness - Express association (and dissociation)	- Perceive emotional cues on teammates - Interpret and reason about group feelings - Express emotional cues	- Perceive liking and disliking cues on teammates Interpret and reason about membership preferences and rejections - Express liking (and disliking) cues	- Perceive social cues with structural information - Recognise and reason about roles, norms and relationships - Express cues with structural information	- Perceive social cues related to the task accomplishment - Reason about individual and group goals in the task - Express social cues to endorse common goals or task accomplishment
Evaluating Robotic Teammates	- Measures of belongingness or identi cation	- Measures of group emotional state	- Measures of attraction or preferences	- Measures of roles (e.g. leadership)	- Measures of trust or uency

Table 1 Guiding the design, development and evaluation of robotic teammates in the light of the dimensions of cohesion.

The novelty of our framework is the adoption of cohesion as a central element to guide the creation of robotic teammates, as shown in Figure 2. We propose that the five dimensions of cohesion can lead the stages of design, development and evaluation. Table 1 clarifies this proposal by providing examples for each dimension of cohesion. Moreover, each of the following sections dives into a row of the table in order to detail and analyse how cohesion contributes to the design, development and evaluation of robotic teammates.

3.1 Cohesion in the design of robotic teammates

We propose that designing the social behaviours of robotic teammates in the light of cohesion can strengthen human-robot teamwork. Based on the dimensions of cohesion, one should opt to endow robotic teammates with behaviours that portray belongingness (i.e. collective cohesion), express group feelings (i.e. emotional cohesion), foster attractions (i.e. social cohesion), acknowledge the roles, norms and relationships (i.e. structural cohesion), and/or endorse the common goals of the teammates (i.e. task cohesion).

When using an off-the-shelf robot, the concrete design of such behaviours must take into account the morphology of the robot and its available multimodal channels. For instance, if the robot possesses any sort of natural language skill (e.g. speech, text), it may portray belongingness by using the pronoun "we" or, if it possesses robotic arms and hands, it may perform inclusive gestures widening and opening its hands towards the teammates. If the morphology of the robot is also a target of the design stage, one should first choose which behaviours may be relevant and adequate for the specific task, and then design its embodiment accordingly. In the case of collective cohesion, we additionally identified the usage of physical accessories that signal belongingness or identification, such as a team colour to elicit the minimal group paradigm (Kuchenbrandt et al., 2011).

3.2 Cohesion in the development of robotic teammates

The dimensions of cohesion can also guide the development of autonomous social behaviour for robotic teammates, which is crucial to deploy robotic teammates in real-world settings. In terms of development, our framework considers three categories of capabilities: perceptive, cognitive and expressive.

Expressive capabilities account for autonomous triggering mechanisms of social cues that can be interpreted by human teammates as attempts to enhance the cohesion of the team. The specific social cues used by the robotic teammate might, in turn, be associated with each dimension of cohesion. Beyond natural language that can easily and explicitly elicit each dimension of cohesion, one can use, for instance, the robot's posture to convey association (i.e. collective cohesion) or liking cues (i.e. social cohesion). Gaze is another example of a powerful non-verbal behaviour that can communicate common goals (i.e. task cohesion) through joint attention cues, or even roles (i.e. structural cohesion) such as leader/subordinate (Capozzi et al., 2019) or addressee/bystander (Mutlu et al., 2012).

The perceptive skills include computational mechanisms to autonomously recognise relevant social cues from human teammates. Once again, the social cues that the robotic teammate can perceive might be more related to a specific dimension of cohesion according to their nature. Following the previous examples, the robot might recognise the body posture of its human teammates as an association/dissociation cue (i.e. collective cohesion) or as a liking/disliking cue (i.e. social cohesion), and also the gaze of the human teammates to infer the implicit roles through cognitive capabilities (i.e. structural cohesion).

Being aware of the mere occurrence of a social cue is not very informative per se for the robotic teammate. For that reason, perception generally requires some degree of cognition as it allows the robotic teammate to interpret the meaning of that social cue in the context of the task and of the teamwork. In some cases, it is hard to separate perception and cognition, for instance, if the robot has a machine learning model that classifies either liking or disliking behaviours based on several multimodal input channels, the model itself combines perceptive and cognitive capabilities. Hence, when considering cohesion in the cognition of a robotic teammate, we refer to the capabilities of interpreting and reasoning about perceptions, whether they involve the degree of belongingness (i.e. collective cohesion), group feelings (i.e. emotional cohesion), membership preferences (i.e. social cohesion), common goals (i.e. task cohesion), or norms and roles (i.e. structural cohesion). Furthermore, such cognitive capabilities imply that the robotic teammate possesses a notion of the human teammates and of the team as a whole.

A final remark that must be mentioned is the fact that, in the design of social behaviours for robotic teammates, our framework assumes the robot's ultimate goal is to enhance and foster cohesive alliances. However, in the development of perceptive skills, the robotic teammates should be aware of social cues that reveal the deterioration of cohesion by the human teammates. A similar negative facet may be included in the robot's expressive capabilities, due to a necessity of signalling a lack of unity in the team, or due to other individual interests by the owner of the robot.

3.3 Cohesion in the evaluation of robotic teammates

Defining what effective human-robot teamwork is might not be simple due to the specific goals of the task or the properties of the interactive scenario. Generally, effective teamwork refers to a satisfactory execution of actions by the team. We propose to evaluate the effectiveness of human-robot teamwork under the umbrella of cohesion. The dimensions of cohesion suggest a broad range of measures to assess teamwork, from performance-related aspects such as trust or fluency (i.e. task cohesion), to relational-related aspects such as attractions (i.e. social cohesion) or identification (i.e. collective cohesion).

The measurement instruments are usually divided into subjective and objective scales (Gaillard et al., 2014). While subjective questionnaires are prone to bias and conscious answers, they also provide richer beliefs and might more easily present relationships with other measures. Objective instruments, on the other hand, produce reliable scores and can be effortlessly tailored to human-robot interactions, as opposed to subjective scales that are usually adapted from psychology. Ideally, both types of measures should be used together as they might further inform the cognitive capabilities of robotic teammates, as in the case of real-time measures.

4 A use case on social cohesion

Social cohesion in interpersonal groups emerges from the relations and attractions between team members, at an individual level, or among the whole team, at a group level. Nonetheless, it is not clear how those attractions develop when there is a social robot on the team. Our rst use case discusses a previous work that addressed social cohesion in human-robot teams by exploring preferences for certain behavioural traits in robotic partners (Correia et al., 2017b). Speci cally, we used the goal orientation theory to design two distinct robotic characters and we assessed which robot people prefer to form a team with. By looking at membership preferences between humans and robots, we explored social cohesion.

The setup uses a Portuguese trick-taking card game called Sueca that is played by four elements, two humans and two robots, split into two teams. Each human-robot team has the goal of beating the numerical score of the other team. The two created robotic partners were fully autonomous and used an architecture for emotional agents described in the extended version of the previously mentioned paper (Correia et al., 2018b). Although the robots portrayed distinct goal orientations in their social behaviours (as described below), they used the same algorithm to play the card game, which is extensively described in another work (Correia et al., 2017a).

The goal orientation theory suggests that people will present either a learning-oriented goal (i.e. an interest in learning something) or a performance-oriented goal (i.e. an interest in the result and what judgements will emerge from it) (Eison, 1979; Dweck, 1986). Teams consisting of individuals with a learning orientation are reported to show high levels of mutual support behaviours and high quality of interaction, team e cacy and commitment Porter (2005). Hence, we have speculated that a robot with a learning goal orientation would foster social cohesion with its human teammate. Based on this theory, we created a repertoire of multimodal utterances, grouped by the situations that can happen during a game session, so that each robot displays distinct goal orientations. For the performance-oriented robotic character, its social behaviours were built based on a competitive perspective, always in pursuit of the best score. For example, when its team is losing the trick, it can express the following utterance: <gaze(partner)> This cannot continue like this! <animation(sadness)> You have to play better!" . By contrast, the utterances of the learning-oriented robot mirror a more relational perspective, verbalising more support and encouragement towards its partner. For example, when its team is losing the trick, it can express the following utterance: <gaze(partner)> No worries, next time we will do better! <animation(wink)> .

In an experimental study, we investigated which of the robotic characters would be chosen by the participants as a partner for future games. Each session of the study was run with two human participants, who did not know each other beforehand (see Figure 1a). The procedure was divided into 3 consecutive games, so that each participant rst partnered with the other human (game 1), then partnered with one of the two robots (game 2), and nally partnered with the other robot (game 3). Participants were asked which robotic character they preferred at two points in time: (1) before having partnered with either robot and (2) after having played with both robots as partners.

The partner choices seemed to be guided by di erent factors depending on the context of the participants. At the end of the rst game, when the participants had experienced both robots as opponents and had not yet created a partner relationship with either robot, they seemed to choose their partners based solely on character. At that time, the learning-oriented robot was the preferred partner. This nding suggests that teams whose members prioritise relational features are perceived more positively in rst encounters.

However, at the end of the nal game, the participants' choices became less clear, calling attention to other factors that came into play. It seems that personal characteristics and team performance took higher precedence after participants had experienced partner relationships with the robots. The participants seemed to be a ected by their own level of competitiveness, with more competitive people preferring the more competitive robot, i.e. the performance-oriented character. Additionally, although both autonomous robots played the game using the same algorithm and the di erence between the numbers of victories achieved by both teams was not signi cant, there was an association between the team performance and the chosen robot. It was observed that people tend to prefer the winning team.

Lastly, we performed a behavioural analysis on the data collected in the aforementioned Oliveira et al. (2018, 2020), using a coding scheme based on Bales Interaction Process Analysis Bales (1950). We found that the structure of the group's communication is di erent according to the goal orientation of the robot and according to their roles. Overall, the occurrence of di erent behavioural patterns in competitive and collaborative interactions with robots suggests that there might be a relationship between social cohesion and structural cohesion in human-robot interactions.

4.1 Contributions of the rst use case within the proposed framework

The goal of this use case was to explore the social cohesion in human-robot teamwork relations, i.e. how humans develop social attractions towards robotic teammates. Speci cally, we have investigated the impact of the goal orientation theory on such membership preferences. Therefore, we have designed two robotic teammates whose speech acts were either portraying a more performance-oriented or competitive behaviour versus a more learning-oriented or relationship-driven behaviour. We showed empirically that these traits can lead robotic teammates to be perceived di erently in terms of social attributes, and we further demonstrated their impact on how humans choose robotic partners for a team.

In order to foster social cohesion in rst and brief teamwork encounters, robots should prioritise the expression of learning-oriented behaviours over performance-oriented ones. However, fostering social cohesion in repeated or longer interactions is more complex and involves other factors beyond the goal orientation of the robot, e.g. team performance. Nevertheless, our results suggest that whenever robotic teammates display a certain goal orientation, they should take into account the goal orientation of the human and adapt accordingly. Although we found evidence that there are other relevant aspects in uencing social cohesion, especially in repeated interactions, we can overall claim that the attraction towards robotic teammates can be fostered by the behavioural traits of the robots.

By using the proposed framework to create robotic teammates, the contributions of the rst use case are illustrated in Table 2. At a design level, the goal orientation theory can guide the design of social behaviours for robotic teammates

Stage	Social Cohesion	Task Cohesion
Designing Robotic Teammates	Behaviours: - Using goal orientation theory in multimodal utterances	
Developing Robotic Teammates		Cognition: - Emotional architecture to reason about the team goals
Evaluating Robotic Teammates	Measures: - Choice of robot for hypothetical future game	

Table 2 Analysis of the rst use case (Correia et al., 2017a,b, 2018b) through the proposed framework of Section 3.

in multi-party competitive settings, because it supports human-robot attractions. Moreover, at an evaluation level, the choice of a robotic character for a hypothetical future was used as a measure to assess social cohesion. Finally, the architecture and algorithms employed by the autonomous robots in this use case further contribute to the development of cognitive skills for robotic teammates.

5 A use case on collective cohesion

Collective cohesion re ects the degree of identi cation with the group. It is usually mentioned as the replacement of the I by the we and, when this feeling becomes stronger, it may even a ect one's actions with the collective goals taking priority over the individual ones. Therefore, collective cohesion goes hand in hand with prosociality and this second use case aims at creating robots that foster prosocial behaviours Paiva et al. (2018, 2021). We speci cally explored how robotic teammates can express di erent levels of collective cohesion and how it a ects human-robot teamwork (Correia et al., 2019).

The scenario used in this work was called For The Record , a collective risk dilemma mapped into an entertaining game. In this game, a team of two social robots and one person formed a musical band and had a goal of avoiding the band's collapse. The dilemma of this game happens in every round when each player individually chooses which skill to upgrade, between their instrument or marketing skills. By upgrading the instrument skill, a player is fostering the chances of the band's success, i.e. to cooperate. By upgrading the marketing skill, on the other hand, a player increases the chances of reaching the best individual score among other players, i.e. to defect. Although the game has uncertain results (e.g. dice throws), a player can generally maximise the collective goal by cooperating. At the same time, individuals may opt to defect and free ride on the e orts of others. A detailed description of the game and a theoretical analysis are available in another previous work (Santos et al., 2020).

We conducted a user study using the For The Record game, where a team of one human participant and two robotic teammates played on a touch screen (see Figure 1b). The two robotic players di ered in the strategies they apply to play the game. One of them always defects by upgrading the marketing skill, which we call the sel sh robot (low collective cohesion), while the other always cooperates by upgrading its instrument skill in every round, which we call the prosocial robot (high collective cohesion). Beyond their difference in the level of collective cohesion to play the game, their verbal and non-verbal behaviours remained similar. Additionally, we also manipulated the outcome of the game (through the digital dice throws) to either be a victory or a loss for the team. This experimental set-up allowed us to investigate the following three research questions: (1) How do people perceive robotic teammates portraying different levels of collective cohesion?; (2) How can the perception of those robotic teammates be affected by the outcome of the team?; and (3) Does the outcome of the team affect the collective cohesion reported by human teammates?

Participants were first introduced to the game, then they played the game with the robots, and they finished the experiment with a questionnaire. The questionnaire included the perceived social attributes of warmth, discomfort and competence of each robot (Carpinella et al., 2017). In terms of group measures, participants reported the degree of trust with the team (Allen and Bergin, 2004) and the identification with the group (Leach et al., 2008). Finally, they were also asked which of the two robots (prosocial vs selfish) would they prefer for a hypothetical future game.

The results showed that the prosocial robot was rated as warmer and caused less discomfort than the selfish robot, which suggests that the display of a prosocial strategy by the robotic partner enhanced the perception of its social attributes. However, the competence attributed to each robot was affected by the game result. Participants considered the selfish robot as less competent only in the losing condition, which suggests the negative effect of losing the game highlighted the difference between the robots. We also verified a similar result in the preference of a robotic partner, which suggests participants only preferred the prosocial robot over the selfish robot in the losing condition. In other words, people do not seem to mind the selfish behaviours of robotic teammates as long as they do not cause the team to lose the game. Another important result of the study is that the success of the team had a significant positive effect on group identification but not on group trust.

5.1 Contributions of the second use case within the proposed framework

In this case study, we looked at collective cohesion from two perspectives. First, we compared the degree of collective cohesion reported by participants when they won or lost the game. We found that the outcome for the human-robot team had an impact on the sense of unity reported by human teammates, which was higher when the team won compared to when the team lost. While this result does not inform the design nor the development of robotic teammates, we successfully instantiated measures of collective cohesion, namely group trust and group identification.

Second, we have also explored how robotic teammates are perceived when they display a high versus a low level of collective cohesion in their actions. Participants have ascribed more positive social attributes to the robotic teammate that compromises individual benefits in favour of the collective welfare. Overall, our results support the successful manipulation of collective cohesion and, therefore, the actions that a robotic teammate takes in a collective risk dilemma can induce different levels of perceived collective cohesion by humans.

Stage	Collective Cohesion	Task Cohesion
Designing Robotic Teammates	Behaviours: - Exhibit prosocial actions (e.g. favouring collective goals over individual goals	
Developing Robotic Teammates		Cognition: - Emotional architecture to reason about the team goals
Evaluating Robotic Teammates	Scenarios: - Collective Risk Dilemmas Measures: - Actions in Collective Risk Dilemmas - Group Identification	

Table 3 Analysis of the second use case (Correia et al., 2019) through the proposed framework of Section 3.

By using the proposed framework to create robotic teammates, the contributions of the second use case are illustrated in Table 3. At the design stage, it seems desirable to consider prosocial actions in the design of robotic teammates, such as favouring collective goals over individual ones. At the stage of evaluation, we first highlight the effective use of collective risk dilemmas (e.g. For The Record) as scenarios to elicit new facets of human-robot teamwork. Secondly, we assessed collective cohesion with the subjective measure of group identification, which refers to the degree of belongingness towards the group.

6 A use case on emotional cohesion

Emotional cohesion refers to the intensity of the members to express group feelings. But how can a robotic teammate express group feelings? We took inspiration from both the group identity theory and the group-based emotions, which are a particular group of emotions that reflects a high level of identification with the group. This use case explored how can robotic teammates autonomously express group feelings by proposing a computational model of group-based emotions (Correia et al., 2018a).

Our computational model is grounded on a recent psychological model of group-based emotions, proposed by Goldenberg et al. (Goldenberg et al., 2016). The first, and the most important, component of the model is the Self-Categorisation component, which is responsible for managing the current context of the interaction as well as the social groups that are present in that context, if any, and their members. These elements constitute a social layer on top of the physical reality that is being perceived by the robot's sensors. Based on the Self-Categorisation Theory (Hornsey, 2008; Turner et al., 1987), when the robot detects a presence of an out-group then its own in-group identity will become more salient.

The emotional appraisal is the second component of the model. This component is responsible for generating emotions in response to the events that occur within the current social context. An event can correspond to a performance of an action or to a change in a property of the environment. For each event perceived, the emotional component performs a series of value judgements about that event in relation to the robot. Then, a set of emotions are synthesised in accordance with those judgements, which in our case correspond to the set of variables from the Ortony-Clore-Collins (OCC) theory (Ortony et al., 1990). For instance, when someone performs an action that is considered blameworthy, then that person is inclined to feel shame. However, if the blameworthy action is performed by another person, then a reproach emotion is felt instead by the observer.

While many researchers have already been able to integrate an emotional appraisal component in a social robot's architecture, the innovative aspect of our model lies in the notion that our appraisal component is capable of considering a social group as the actor of an event even if, in reality, all actions are being performed by individuals. This is the result of introducing the Self-Categorisation component before the appraisal takes place in order to determine whether the robot sees itself and others acting based on their individual or their group identity. In the latter case, actions performed by individuals that are sharing the same group identity in the current context are appraised as if they are actions performed by the robot itself. Consequently, in a context where the robot is performing a team-based activity and one of its partners performs a blameworthy action, the appraisal component will generate in the robot a group-based emotion of shame, rather than a reproach emotion towards its partner.

Finally, the last component of the model is the Emotional Response component, which is responsible for managing how the robot expresses the emotions that result from the appraisal process. This process must take into account the di erent possibilities that are a orded by the robot's embodiment. Assuming the robot has the ability to change its facial expression and body posture then these are matched to the current emotional state of the robot. In addition to non-verbal signals, the dialogue acts chosen by the robot are also in uenced by its emotions.

To evaluate the proposed model, we used the same card game scenario as in Section 4 and conducted a user study where two participants and two robots formed two human-robot teams to play the card game (see Figure 1c). One of the robotic partners was using the Self-Categorisation component and was, therefore, expressing group-based emotions, while the other robot skipped the step of Self-Categorisation and was producing the traditional individual-based emotions. Participants were randomly assigned to one robotic partner and, after playing a session of the card game, they were asked to report their subjective evaluation of their robotic teammate, using the Godspeed questionnaire (Bartneck et al., 2009), and of their whole team, using trust (Allen and Bergin, 2004) and group identi cation (Leach et al., 2008).

Overall, our results showed that participants that partnered with the robot expressing group-based emotions rated their robotic partner as more likeable compared to participants that partnered with the robot expressing individual-based emotions. In terms of group measures, participants also attributed higher levels of group identi cation and group trust towards their teams, when their robotic partner was expressing group-based emotions.

6.1 Contributions of the third use case within the proposed framework

In this use case, we explored emotional cohesion by proposing a model for the autonomous expression of group feelings, i.e. group-based emotions. We were particularly interested in addressing both the computational techniques to express these emotions autonomously and assessing the impact of this type of emotion in human-robot teamwork. We found support in our user study that the expression of group feelings by robotic teammates intensi es the reported cohesion by humans with their teams. Not only does this empirical support validates the model we proposed to express group-based emotions in robotic characters, but it also demonstrates that these types of emotions are a desirable design feature of robotic teammates. Overall, the results from our user study strongly support that a robot can display emotional cohesion through group-based emotions and that the expression of those emotions can in turn foster the collective cohesion reported by human teammates.

Stage	Emotional Cohesion	Task Cohesion
Designing Robotic Teammates	Behaviours: - Using group-based emotions in multimodal utterances	
Developing Robotic Teammates	Cognition: - Modeling self-categorisation Expression: - Autonomous expression of group-based emotions	Cognition: - Emotional architecture to reason about the team goals
Evaluating Robotic Teammates		Measures: - Group trust

Table 4 Analysis of the third use case (Correia et al., 2018a) through the proposed framework of Section 3.

By using the proposed framework to create robotic teammates, the contributions of the third use case are illustrated in Table 4. At the design stage, it seems desirable to consider group-based emotions in the design of social behaviours for robotic teammates. At the stage of development, we highlight as a cognitive skill the process of self-categorisation that allows the robot to perceive and identify certain social groups (i.e. establishing in-groups and out-groups). Such a process can, in turn, trigger the autonomous expression of this speci c type of emotion. Finally, the evaluation stage includes the assessment of task cohesion with a subjective scale of group trust, intended to capture the satisfaction with the team and trustworthiness that the common goal will be accomplished.

7 A use case on structural cohesion

Structural cohesion represents the roles, norms, or interactions among members of a group, and the topology of those structures is associated with the degree of cohesion of the group. Taking inspiration from this dimension of cohesion that is usually explored with network analysis, we speculated on the degree of connectivity that a robotic teammate should consider in its perceptive skills. For instance, should a robotic teammate perceive communicative acts only towards itself? Or should it also perceive communicative acts between other pairs of team members? Furthermore, are those structural communicative behaviours of a robotic teammate able to enhance human-robot teamwork? The last use case has addressed these research questions by focusing on the nonverbal gaze behaviour of a robotic teammate in a multi-party team setting.

We created a silent coordination task by envisioning a shared workspace between humans and robots e.g. an assembly line in which coordinated behaviour is mainly achieved through non-verbal behaviours e.g. due to a noisy environment. In such a situation, the gaze behaviours of a robotic teammate during a silent coordination task hold important cognitive functions of gaze, such as monitoring and communicative. Therefore, we developed an autonomous gaze system for a robot and proposed two possible approaches for the gaze behaviours of robotic teammates in multi-party settings. The system perceives the gaze behaviours of human teammates and either responds only to gazes towards the robot itself Responding Gazes, or additionally responds to any gaze between the human teammates Attentive Gazes. The Responding Gazeof the robotic teammate attempts to establish mutual gaze, while the Attentive Gaze additionally attempts to establish joint attention behaviours.

The evaluation was done through a third-person perspective over an online recruitment platform, in which participants were asked to watch videos of human-robot team interaction and to rate their perceptions of teamwork regarding each video (see Figure 1d). We created video scripts with gaze instructions for the human teammates for 10 short scenes, in which the goal was to illustrate possible situations that may occur in the silent coordination task. We recorded 3 times each of the 10 scenes (one per condition) using two human actors to impersonate the gaze behaviours of the scripts. The robot's gaze behaviours were fully autonomous and varied between conditions. In each condition the robot was reacting in one of the following three ways:

- Baseline (B) The robot does not perceive and, therefore, does not respond to human gaze behaviours. It only blinks its eyelids;
- Responding Gaze (RG) - The robot perceives human gaze behaviours towards itself and responds by gazing back at that team member. It also blinks as in the Baseline condition;
- Attentive Gaze (AG) - The robot perceives human gaze behaviours towards any target and responds by either gazing back at the gazer, in the case of the robot being the target, or by gazing at the same team member that is currently being gazed at by another. It also blinks as in the Baseline condition.

Participants were randomly assigned to one of the three conditions of robotic gaze behaviour and then randomly assigned to watch 5 of the 10 videos of that condition. For each video they saw, they had to report their perception of teamwork, in terms of coordination, uency, teammate traits and work alliance (Ho man, 2019).

The results of our user study revealed that the robot's gaze behaviours in both RG and AG conditions positively a ected how the robot is perceived as a teammate, as well as the perceived uency of its team. The perception of uency was sensitive to the di erence between the two experimental conditions with attentive gazes being rated as more uid than responding gazes. We also found that perceptions of coordinated teamwork were highly a ected by the number of human teammates performing gaze behaviours. Speci cally, a human-robot team in which the two humans display gaze behaviours is perceived as more coordinated, more uent, and more united than a human-robot team in which only one human teammate displays gaze behaviours. Finally, by looking at scenes with gaze behaviours by only 2 teammates (while the third one stared at an object), we only found di erences in the perceived coordination among 2 humans compared to perceived coordination among 1 human and 1 robot. It seems the perceived coordination of the human-robot team was highly a ected by the gaze behaviours of the human teammates, regardless of the robot's gaze strategy. The perceived uency of the human-robot team revealed, however, to be more sensitive to the robot's gaze behaviours.

7.1 Contributions of the fourth use case within the proposed framework

This work explored structural cohesion by proposing two heuristics for the gaze behaviour of robotic teammates that map di erent structural communicative behaviours. The responding gaze behaviour has a lower degree of connectivity in the sense that the robot only perceives (and reacts) to gaze behaviours towards itself. The attentive gaze behaviour has a higher degree of connectivity as the robot also perceives (and reacts) to gazes between other teammates. These two gaze behaviours support two di erent communication structures among the team, which might, in turn, a ect di erently the coordination between team members. Our empirical study showed that the perception of teamwork is better when a robotic teammate employs these two gaze behaviours, compared to a baseline in which the robot does not perceive nor react to any gaze behaviour. Moreover, the attentive gazes seem to lead to higher levels of perceived uency in the human-robot team. Overall, the preliminary ndings of this use case suggest there is a relation between the structural cohesion employed by a robotic teammate and the perception of teamwork.

Stage	Structural Cohesion	Task Cohesion
Designing Robotic Teammates	Behaviours: - Mutual gaze and joint attention in multi-party settings	
Developing Robotic Teammates	Perception: - Autonomous detection of gazes from human teammates Expression: - Autonomous reaction to gazes by human teammates	
Evaluating Robotic Teammates		Measures: - Fluency

Table 5 Analysis of the fourth use case through the proposed framework of Section 3.

By using the proposed framework to create robotic teammates, the contributions of the fourth and last use case are illustrated in Table 5. This use case proposes that the design of gaze behaviours for robotic teammates should take into account aspects of structural cohesion. For instance, by reacting to gazes between the two human teammates, the robot attempts to establish joint attention and augments the structural connectivity of its team. At the stage of development, this use case contributes to an autonomous detection and reaction system of gaze behaviours. Finally, the evaluation stage includes the assessment of task cohesion with a subjective scale of uency, which captures the dimensions of coordination and e ciency of the team.

8 Discussion and future challenges

This paper proposes a novel approach to enhance robotic teammates through the multifaceted cohesion construct from the social sciences, presented in Section 3. We demonstrate this vision by presenting some examples from our previous work, discussed in Sections 4-7, where each one explores at least one dimension of cohesion. However, creating robotic teammates that are e ective collaborators is still an ongoing e ort and involves several key questions. The following sections summarise our contributions and highlight future challenges.

8.1 How can we design robotic teammates?

This paper suggests the well-known dimensions of cohesion from interpersonal group relations can be used to guide the design of robotic teammates, not only in their social capabilities but also in their morphology.

In our use cases, we have rst shown that to foster attraction with human teammates, robots can display character traits, such as goal orientations. Similarly, robots can be perceived with more positive social attributes when they portray strong collectivism in their actions, hindering their individual goals. Another example evidenced the positive impact of robotic teammates portraying group-based emotions to express group feelings. Finally, preliminary ndings on gaze behaviours in multi-party team settings suggested that the communication structure of gazes between human-robot teammates can a ect the perception of teamwork. These examples demonstrate how social capabilities can be designed when having in mind the dimensions of cohesion.

Although our use cases are focused on social capabilities,addressing the design of robotic teammates' morphology is also an essential issue. Especially because the design of behaviours should go hand in hand with the design of morphology. For instance, future challenges must tackle which multimodal modalities are required to perform certain behaviours while ensuring the embodiment of robotic teammates conveys the right social a ordances for a speci c task.

8.2 How can we develop autonomous robotic teammates?

In order for robots to act autonomously in human-robot teams, they need to be endowed with three core skills. Firstly, they require perceptive skills to observe dynamic changes in the environment, which include changes not only related to the task but also related to other human teammates (e.g. intentions). Secondly, robots need cognitive skills to reason about the information they observe in a team context, and about the expected impact of their own actions. Generally, such skills involve an understanding of group/team processes. Thirdly, robots require expressive skills to facilitate teamwork with humans (e.g. conveying intentions or internal states). Therefore, we argue that robotic teammates should hold autonomous capabilities at three di erent levels: perceptive, cognitive and expressive.

In terms of autonomous perceptive skills, the work discussed in Section 7 explores the perception of gaze behaviours from human teammates. Speci cally, in a triadic team with one robot and two humans, we investigated whether a robotic teammate should perceive gaze behaviours between the two human teammates or if it should only perceive gaze behaviours towards itself. Regarding cognitive skills, we highlight the work discussed in Section 6, in which we proposed a social categorization step before the agent's appraisal process. Such a mechanism allowed the robotic agent to autonomously reason about group identities. It is important to mention that group identity is usually more salient in multi-party competitive settings, as in our card game scenario. Additionally, in terms of expressive skills, we contributed with heuristics to guide the robot's gaze behaviour in multi-party settings (Section 7), and a mechanism to express group feelings according to the previously mentioned salient group identities (Section 6). Finally, all the presented works use an architecture for emotional agents that illustrates cognitive capabilities in terms of task cohesion. A robotic teammate that performs an emotional appraisal of task-related events can e ciently reason about the task using emotions, regardless of how the emotional responses are designed.

Most researchers have been evaluating speci c research questions in controlled environments, and some even employ wizard of oz techniques, or other semi-autonomous approaches, mainly to overcome the di culties of perceiving human teammates. We believe creating more advanced perceptive skills regarding human-robot teamwork will raise the bar for autonomous robotic teammates .

Additionally, as one of the assumptions in Group Dynamics literature is that groups evolve over time (Forsyth, 1990), we also highlight as a promising approach the creation of adaptive capabilities for robotic teammates to cope with dynamic changes . Such mechanisms would allow robotic teammates to adapt their policies according to relevant changes in the environment and, as a result, maintain longer relations. For instance, a robotic teammate capable of predicting the cohesion of its team, or the cohesive alliance expressed by each team member, could explore online techniques to cope with what is happening. Once again, predicting cohesion might as well bene t from the multifaceted construct we presented in this paper, by considering for instance: attractions, prosocial actions, group feelings, and/or the communication structure of the team.

8.3 How can we evaluate human-robot teamwork?

While assuming the importance of cohesion in human-robot teamwork, it is important not only to address how the robotic teammates should express their sense of unity with the team but also to understand and evaluate how human teammates do it. The works we have presented showed examples of evaluation metrics that contribute to the assessment of cohesion, such as measuring attractions and membership preferences (Section 4), prosocial actions in collective risk dilemmas (Section 5), self-reported group identi cation (Section 6), perception of coordination and uency (Section 7).

Although we argue cohesion is an important measure to consider in human-robot collaboration, we also acknowledge the importance of assessing the e ectiveness of teamwork through outcome-related measures such as performance. We have indeed found evidence in our previous work that there is a relationship between performance and the attraction with teammates (Section 4), and that the team's outcome has an impact on both the reported collective cohesion and the perception of the teammates (Section 5). Our argument is instead that human-robot cohesion might mediate the performance of these teams, similarly to interpersonal teams (Cartwright, 1968; Mullen and Copper, 1994; Beal et al., 2003).

Lastly, we also acknowledge that this paper considers cohesion as a broader umbrella of satisfaction and a sense unity of among the team members, which holds two implicit limitations. First, we are blending many speci c constructs together (Abrams and Rosenthal-von der Pütten, 2020), that require further scrutiny. And second, we are directly transferring constructs from social sciences to human-robot interaction. Therefore, one of the biggest challenges in terms of evaluation is the creation of proper measures tailored to teamwork between humans and robotic machines . Robots as teammates certainly introduce new facets to the interaction beyond the well-known interpersonal relationships. Although adapting current subjective questionnaires from psychology can constitute a starting point and can provide informative knowledge, they might be insu cient in the long run.

8.4 Can a human-robot team form a cohesive alliance?

Three important concerns should be taken into account when targeting cohesive alliances between humans and robots in teamwork settings. First, the presented uses cases indeed support thatcohesion can be used as a framework to create robotic teammates in the three stages of design, development and evaluation . However, we acknowledge there might be other relevant constructs to guide such processes, e.g. trust (De Visser et al., 2020).

Second, we recognise that the concept itself of cohesive alliances between humans and robots can change over time, considering the presence of robots in our society is increasing every day. For instance, Abrams and Rosenthal-von der Pütten raised the question of what it means to identify with a human-robot group considering this type of mixed group is not yet established in our daily lives (Abrams and Rosenthal-von der Pütten, 2020). As a result, we suggest future research to embrace and account for how the mere presence of robotic technologies can in uence the way we establish teamwork with robots .

Last, another concern related to addressing cohesion in human-robot teamwork is the cultural aspects. Speci cally because one of the dimensions to analyse cultural di erences between populations is the degree of individualism versus collectivism, which can easily translate into di erent notions of groups and teams (Zhou and Shi, 2011). As a result, we believe human-robot cohesion might be di erently expressed in di erent cultures .

8.5 Concluding Remarks

Overall, we believe that e ective teamwork can be supported and fostered by a shared sense of unity among team members. In other words, cohesion in human-robot teams might be a key factor to propel team performance and it should be considered in the design, development and evaluation of robotic teammates.

Acknowledgments

This work was supported by national funds from the Portuguese Fundação para a Ciência e a Tecnologia through INESC-ID (UIDB/50021/2020) and LARSYS (UIDB/50009/2020), and the ARCADE project (PTDC/CCI-COM/30274/2017). Filipa Correia also acknowledges her PhD grant by Fundação para Ciência e a Tecnologia (SFRH/BD/118031/2016). Lastly, the authors thank Greg Trafton for his thorough review and insightful suggestions, which were crucial to re ning and improving the paper.

References

Anna MH Abrams and Astrid M Rosenthal-von der Pütten. 2020. I C E Framework: Concepts for Group Dynamics Research in Human-Robot Interaction. International Journal of Social Robotics (2020), 1 17.

Kathleen Allen and Richard Bergin. 2004. Exploring trust, group satisfaction, and performance in geographically dispersed and co-located university technology commercialization teams. In In Proceedings of the NCIIA 8th Annual Meeting: Education that Works . 18 20. https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.201.407

Patrícia Alves-Oliveira, Pedro Sequeira, Francisco S Melo, Ginevra Castellano, and Ana Paiva. 2019. Empathic robot for group learning: A eld study. ACM Transactions on Human-Robot Interaction (THRI) 8, 1 (2019), 1 34.

Robert F Bales. 1950. Interaction process analysis; a method for the study of small groups. (1950).

Christoph Bartneck, Dana Kuli¢, Elizabeth Croft, and Susana Zoghbi. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. International journal of social robotics 1, 1 (2009), 71 81.

Daniel J Beal, Robin R Cohen, Michael J Burke, and Christy L McLendon. 2003. Cohesion and performance in groups: a meta-analytic clari cation of construct relations. Journal of applied psychology 88, 6 (2003), 989.

Tony Belpaeme, James Kennedy, Aditi Ramachandran, Brian Scassellati, and Fumihide Tanaka. 2018. Social robots for education: A review. Science robotics 3, 21 (2018), eaat5954.

Jürgen Brandstetter, Péter Rácz, Clay Beckner, Eduardo B Sandoval, Jennifer Hay, and Christoph Bartneck. 2014. A peer pressure experiment: Recreation of the Asch conformity experiment with robots. In 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems IEEE, 1335 1340.

Cynthia Breazeal. 2011. Social robots for health applications. In 2011 Annual international conference of the IEEE engineering in medicine and biology society. IEEE, 5368 5371.

Cynthia Breazeal, Cory D Kidd, Andrea Lockerd Thomaz, Guy Ho man, and Matt Berlin. 2005. E ects of nonverbal communication on e ciency and robustness in human-robot teamwork. In 2005 IEEE/RSJ international conference on intelligent robots and systems IEEE, 708 713.

Francesca Capozzi, Cigdem Beyan, Antonio Pierro, Atesh Koul, Vittorio Murino, Stefano Livi, Andrew P Bayliss, Jelena Ristic, and Cristina Becchio. 2019. Tracking the Leader: Gaze Behavior in Group Interactions. Iscience 16 (2019), 242 249.

Colleen M Carpinella, Alisa B Wyman, Michael A Perez, and Steven J Stroessner. 2017. The robotic social attributes scale (RoSAS): development and validation. In ACM/IEEE Int. Conf. on Human-Robot Interaction . https://dl.acm.org/doi/10.1145/2909824.3020208

Albert V Carron and Lawrence R Brawley. 2000. Cohesion: Conceptual and measurement issues. Small group research 31, 1 (2000), 89 106.

Dorwin Cartwright. 1968. The nature of group cohesiveness. Group dynamics: Research and theory 91 (1968), 109.

Mai Lee Chang, Reymundo A Gutierrez, Priyanka Khante, Elaine Schaertl Short, and Andrea Lockerd Thomaz. 2018. E ects of integrated intent recognition and communication on human-robot collaboration. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 3381 3386.

Wan-Ling Chang, Jeremy P White, Joohyun Park, Anna Holm, and Selma 'abanovi¢. 2012. The e ect of group size on people's attitudes and cooperative behaviors toward robots in interactive gameplay. In RO-MAN, 2012 IEEE . IEEE, 845 850.

Henrik I Christensen. 2003. Intelligent home appliances. In Robotics Research Springer, 319 327.

Filipa Correia, Patrícia Alves-Oliveira, Tiago Ribeiro, Francisco S Melo, and Ana Paiva. 2017a. A social robot as a card game player. In 13th AAAI Conference on Arti cial Intelligence and Interactive Digital Entertainment .

Filipa Correia, Samuel Mascarenhas, Rui Prada, Francisco S Melo, and Ana Paiva. 2018a. Group-based Emotions in Teams of Humans and Robots. In Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction . ACM, 261 269.

Filipa Correia, Samuel F Mascarenhas, Samuel Gomes, Patrícia Arriaga, Iolanda Leite, Rui Prada, Francisco S Melo, and Ana Paiva. 2019. Exploring prosociality in human-robot teams. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI) . IEEE, 143 151. https://ieeexplore.ieee.org/abstract/document/8673299

Filipa Correia, So a Petisca, Patr cia Alves-Oliveira, Tiago Ribeiro, Francisco S Melo, and Ana Paiva. 2017b. Groups of humans and robots: Understanding membership preferences and team formation. In Robotics: Science and Systems, RSS, Vol. 17.

Filipa Correia, So a Petisca, Patrícia Alves-Oliveira, Tiago Ribeiro, Francisco S Melo, and Ana Paiva. 2018b. I Choose... YOU! Membership preferences in human robot teams. Autonomous Robots (2018), 1 15.

Albert A Cota, Charles R Evans, Kenneth L Dion, Lindy Kilik, and R Stewart Longman. 1995. The structure of group cohesion. Personality and social psychology bulletin 21, 6 (1995), 572 580.
Ewart J De Visser, Marieke MM Peeters, Malte F Jung, Spencer Kohn, Tyler H Shaw, Richard Pak, and Mark A Neerincx. 2020. Towards a theory of longitudinal trust calibration in human robot teams. International journal of social robotics 12, 2 (2020), 459 478.
Anca D Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S Srinivasa. 2015. E ects of robot motion on human-robot collaboration. In Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction. ACM, 51 58.
Carol S Dweck. 1986. Motivational processes a ecting learning. American psychologist 41, 10 (1986), 1040.
James Arthur Eison. 1979. The development and validation of a scale to assess di ering student orientations towards grades and learning. Ph.D. Dissertation. University of Tennessee, Knoxville.
Mark Eys, Todd Loughead, Steven R Bray, and Albert V Carron. 2009. Development of a cohesion questionnaire for youth: The Youth Sport Environment Questionnaire. Journal of Sport and Exercise Psychology 31, 3 (2009), 390 408.
Donelson R Forsyth. 1990. Group dynamics. (1990).
Marlena Fraune, Selma 'abanovi¢, and Takayuki Kanda. 2019a. Human Group Presence, Group Characteristics, and Group Norms a ect Human-Robot Interaction in Naturalistic Settings. Frontiers in Robotics and AI 6 (2019), 48.
Marlena R Fraune, Yusaku Nishiwaki, Selma 'abanovi¢, Eliot R Smith, and Michio Okada. 2017a. Threatening Flocks and Mindful Snow akes: How Group Entitativity A ects Perceptions of Robots. In Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction. ACM, 205 213.
Marlena R Fraune, S 'abanovi¢, and Eliot R Smith. 2017b. Teammates rst: Favoring ingroup robots over outgroup humans. In RO-MAN 2017. The 26th IEEE International Symposium on Robot and Human Interactive Communication, Submitted.
Marlena R Fraune, Steven Sherrin, Selma 'abanovi¢, and Eliot R Smith. 2019b. Is Human-Robot Interaction More Competitive Between Groups Than Between Individuals?. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE, 104 113.
Vinciane Gaillard, Axel Cleeremans, and Arnaud Destrebecqz. 2014. Dissociating conscious and unconscious learning with objective and subjective measures. Clinical EEG and neuroscience 45, 1 (2014), 50 56.
Sarah Gillet, Ronald Cumbal, André Pereira, José Lopes, Olov Engwall, and Iolanda Leite. 2021. Robot Gaze Can Mediate Participation Imbalance in Groups with Di erent Skill Levels. In 2021 ACM/IEEE International Conference on Human-Robot Interaction (TO APPEAR). IEEE.
Amit Goldenberg, Eran Halperin, Martijn van Zomeren, and James J Gross. 2016. The process model of group-based emotion: Integrating intergroup emotion and emotion regulation perspectives. Personality and Social Psychology Review 20, 2 (2016), 118 141.
Victoria Groom and Cli ord Nass. 2007. Can robots be teammates?: Benchmarks in human robot teams. Interaction Studies 8, 3 (2007), 483 500.

Jérémie Guiochet, Mathilde Machin, and Hélène Waeselynck. 2017. Safety-critical advanced robots: A survey. Robotics and Autonomous Systems 94 (2017), 43 52.

Guy Ho man. 2019. Evaluating uency in human robot collaboration. IEEE Transactions on Human-Machine Systems 49, 3 (2019), 209 218.

Guy Ho man and Cynthia Breazeal. 2004. Collaboration in Human-Robot Teams. AIAA 1st Intelligent Systems Technical Conference (2004), 1 18. https://doi.org/10.2514/6.2004-6434

Guy Ho man and Cynthia Breazeal. 2007. E ects of anticipatory action on human-robot teamwork e ciency, uency, and perception of team. In Proceedings of the ACM/IEEE international conference on Human-robot interaction . ACM, 1 8.

Matthew J Hornsey. 2008. Social identity theory and self-categorization theory: A historical review. Social and Personality Psychology Compass 2, 1 (2008), 204 222.

Chien-Ming Huang and Bilge Mutlu. 2016. Anticipatory robot control for e cient human-robot collaboration. In The eleventh ACM/IEEE international conference on human robot interaction. IEEE Press, 83 90.

Tariq Iqbal and Laurel D Riek. 2017. Coordination dynamics in multihuman multirobot teams. IEEE Robotics and Automation Letters 2, 3 (2017), 1712 1717.

Björn Jensen, Nicola Tomatis, Laetitia Mayor, Andrzej Drygajlo, and Roland Siegwart. 2005. Robots meet humans-interaction in public spaces. IEEE Transactions on Industrial Electronics 52, 6 (2005), 1530 1546.

Malte F Jung, Dominic DiFranzo, Brett Stoll, Solace Shen, Austin Lawrence, and Houston Claure. 2018. Robot Assisted Tower Construction-A Resource Distribution Task to Study Human-Robot Collaboration and Interaction with Groups of People. arXiv preprint arXiv:1812.09548 (2018).

Malte F Jung, Jin Joo Lee, Nick DePalma, Sigurdur O Adalgeirsson, Pamela J Hinds, and Cynthia Breazeal. 2013. Engaging robots: easing complex human-robot teamwork using backchanneling. In Proceedings of the 2013 conference on Computer supported cooperative work ACM, 1555 1566.

Takayuki Kanda, Masahiro Shiomi, Zenta Miyashita, Hiroshi Ishiguro, and Norihiro Hagita. 2009. An a ective guide robot in a shopping mall. In Proceedings of the 4th ACM/IEEE international conference on Human robot interaction . ACM, 173 180.

Alap Kshirsagar, Melanie Lim, Shemar Christian, and Guy Ho man. 2020. Robot Gaze Behaviors in Human-to-Robot Handovers. IEEE Robotics and Automation Letters 5, 4 (2020), 6552 6558.

Dieta Kuchenbrandt, Friederike Eyssel, Simon Bobinger, and Maria Neufeld. 2011. Minimal group-maximal e ect? evaluation and anthropomorphization of the humanoid robot NAO. In International conference on social robotics. Springer, 104 113.

Nicholas C Landol and Anca D Dragan. 2018. Social Cohesion in Autonomous Driving. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 8118 8125.

Colin Wayne Leach, Martijn Van Zomeren, Sven Zebel, Michael LW Vliek, Sjoerd F Pennekamp, Bertjan Doosje, Jaap W Ouwerkerk, and Russell Spears. 2008. Group-level self-de nition and self-investment: a hierarchical (multicomponent) model of in-group identi cation. Journal of personality and social psychology 95, 1 (2008), 144. https://psycnet.apa.org/doiLanding?doi=10.

1037%2F0022-3514.95.1.144
Albert J Lott and Bernice E Lott. 1965. Group cohesiveness as interpersonal attraction: A review of relationships with antecedent and consequent variables. Psychological bulletin 64, 4 (1965), 259.
Christoforos Mavrogiannis, Alena M Hutchinson, John Macdonald, Patrícia Alves-Oliveira, and Ross A Knepper. 2019. E ects of Distinct Robot Navigation Strategies on Human Behavior in a Crowded Environment. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI) . IEEE, 421 430.
Brian Mullen and Carolyn Copper. 1994. The relation between group cohesiveness and performance: An integration. Psychological bulletin 115, 2 (1994), 210.
Bilge Mutlu, Takayuki Kanda, Jodi Forlizzi, Jessica Hodgins, and Hiroshi Ishiguro. 2012. Conversational gaze mechanisms for humanlike robots.ACM Transactions on Interactive Intelligent Systems (TiiS) 1, 2 (2012), 1 33.
Raquel Oliveira, Patrícia Arriaga, Patrícia Alves-Oliveira, Filipa Correia, So a Petisca, and Ana Paiva. 2018. Friends or Foes?: Socioemotional Support and Gaze Behaviors in Mixed Groups of Humans and Robots. In Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction . ACM, 279 288.
Raquel Oliveira, Patrícia Arriaga, Filipa Correia, and Ana Paiva. 2020. Looking Beyond Collaboration: Socioemotional Positive, Negative and Task-Oriented Behaviors in Human Robot Group Interactions. International Journal of Social Robotics 12, 2 (2020), 505 518.
Andrew Ortony, Gerald L Clore, and Allan Collins. 1990. The cognitive structure of emotions. Cambridge university press.
Ana Paiva, Filipa Correia, Raquel Oliveira, Fernando P. Santos, and Patrícia Arriaga. 2021. Empathy and Prosociality in Social Agents. Handbook on Socially Interactive Agents (2021).
Ana Paiva, Fernando P Santos, and Francisco C Santos. 2018. Engineering pro-sociality with autonomous agents. In Thirty-Second AAAI Conference on Arti cial Intelligence . https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/viewPaper/16799
William E Piper, Myriam Marrache, Renee Lacroix, Astrid M Richardsen, and Barry D Jones. 1983. Cohesion as a basic bond in groups.Human Relations 36, 2 (1983), 93 108.
Christopher OLH Porter. 2005. Goal orientation: e ects on backing up behavior, performance, e cacy, and commitment in teams. Journal of Applied Psychology 90, 4 (2005), 811.
Nicole Salomons, Sarah Strohkorb Sebo, Meiying Qin, and Brian Scassellati. 2021. A Minority of One against a Majority of Robots: Robots Cause Normative and Informational Conformity. ACM Transactions on Human-Robot Interaction (THRI) 10, 2 (2021), 1 22.
Nicole Salomons, Michael van der Linden, Sarah Strohkorb Sebo, and Brian Scassellati. 2018. Humans conform to robots: Disambiguating trust, truth, and conformity. In Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction . ACM, 187 195.
Fernando P Santos, Samuel Mascarenhas, Francisco C Santos, Filipa Correia, Samuel Gomes, and Ana Paiva. 2020. Picky losers and carefree winners prevail in collective risk dilemmas with partner selection. Autonomous Agents and

Multi-Agent Systems 34, 2 (2020), 1 29.
Adarsh Jagan Sathyamoorthy, Utsav Patel, Moumita Paul, Nithish K Sanjeev Kumar, Yash Savle, and Dinesh Manocha. [n.d.]. CoMet: Modeling Group Cohesion for Socially Compliant Robot Navigation in Crowded Scenes. ([n. d.]).
Sarah Sebo, Ling Liang Dong, Nicholas Chang, Michal Lewkowicz, Michael Schutzman, and Brian Scassellati. 2020. The In uence of Robot Verbal Support on Human Team Members: Encouraging Outgroup Contributions and Suppressing Ingroup Supportive Behavior. Frontiers in Psychology 11 (2020), 3584.
Isabella Seeber, Eva Bittner, Robert O Briggs, Triparna de Vreede, Gert-Jan De Vreede, Aaron Elkins, Ronald Maier, Alexander B Merz, Sarah Oeste-Reiÿ, Nils Randrup, et al. 2020. Machines as teammates: A research agenda on AI in team collaboration. Information & management 57, 2 (2020), 103174.
Julie Shah, James Wiken, Brian Williams, and Cynthia Breazeal. 2011. Improved human-robot team performance using chaski, a human-inspired plan execution system. In Proceedings of the 6th international conference on Human-robot interaction. ACM, 29 36.
Mahni Shayganfar, Charles Rich, and Candace Sidner. 2019. Appraisal Algorithms for Relevance and Controllability in Human-Robot Collaboration. In 2019 IEEE International Conference on Humanized Computing and Communication (HCC). IEEE, 31 37.
Solace Shen, Petr Slovak, and Malte F Jung. 2018. Stop. I see a con ict happening.: A robot mediator for young children's interpersonal con ict resolution. In Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction. ACM, 69 77.
Elaine Short and Maja J Mataric. 2017. Robot moderation of a collaborative game: Towards socially assistive robotics in group interactions. In 2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN). IEEE, 385 390.
Andrew Steain, Christopher John Stanton, and Catherine J Stevens. 2019. The black sheep e ect: The case of the deviant ingroup robot. PloS one 14, 10 (2019).
Brett Stoll, Samantha Reig, Lucy He, Ian Kaplan, Malte F Jung, and Susan R Fussell. 2018. Wait, Can You Move the Robot?: Examining Telepresence Robot Use in Collaborative Teams. In Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction. ACM, 14 22.
Sarah Strohkorb, Ethan Fukuto, Natalie Warren, Charles Taylor, Bobby Berry, and Brian Scassellati. 2016. Improving human-human collaboration between children with a social robot. In 2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN). IEEE, 551 556.
Sarah Strohkorb Sebo, Margaret Traeger, Malte Jung, and Brian Scassellati. 2018. The Ripple E ects of Vulnerability. Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction - HRI '18 (2018), 178 186. https://doi.org/10.1145/3171221.3171275
Hamish Tennent, Solace Shen, and Malte Jung. 2019. Micbot: A peripheral robotic object to shape conversational dynamics and team performance. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE, 133 142.
Margaret L Traeger, Sarah Strohkorb Sebo, Malte Jung, Brian Scassellati, and Nicholas A Christakis. 2020. Vulnerable robots positively shape human conver-

sational dynamics in a human–robot team. *Proceedings of the National Academy of Sciences* 117, 12 (2020), 6370–6375.

John C Turner, Michael A Hogg, Penelope J Oakes, Stephen D Reicher, and Margaret S Wetherell. 1987. *Rediscovering the social group: A self-categorization theory.* Basil Blackwell.

Takahisa Uchida, Hiroshi Ishiguro, and Peter Ford Dominey. 2020. Improving Quality of Life with a Narrative Robot Companion: II–Creating Group Cohesion via Shared Narrative Experience. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 906–913.

Wencang Zhou and Xuli Shi. 2011. Special Review Article: Culture in groups and teams: A review of three decades of research. *International Journal of Cross Cultural Management* 11, 1 (2011), 5–34.