Enhancing Social Collaboration in Human-Robot Teams

Sarah Strohkorb Sebo and Brian Scassellati
Yale University
{sarah.sebo, brian.scassellati}@yale.edu

I. Motivation

Recent work in psychology has demonstrated that team performance is significantly driven by factors related to the the social support team members exhibit towards one another. Woolley et al. [18] showed that the collective intelligence of a team across a variety of tasks is significantly correlated with the team's average social sensitivity, assessed using the "reading the mind in the eyes test." Edmondson [2] found that a team's psycholocial safety "a shared belief held by members of a team that the team is safe for interpersonal risk taking," leads to higher team learning, which in turn leads to higher team performance of work teams in a manufacturing company. Additionally, Stubbs and Wolff [14] demonstrated that the emotional intelligence of a military team leader is correlated with the presence of emotionally competent group norms, and the presence of these group norms is significantly related to team performance.

As robots are increasingly adopted and incorporated into everyday social and work environments, it is critical that they are able to contribute positively to group's social support in order to maximize both team member satisfaction and team performance. As an illustrative example, imagine a human-robot team that is assembling a portion of an airplane, where the human team members are putting together and fastening the necessary parts and the robot runs tests to ensure correct assembly. The robot finds several errors in the airplane section assembly and verbally notifies its human teammates using either utterance A, "Abraham failed to correctly attach part x to part y in 3 out of 5 places." or utterance B, "Part x was not correctly attached to part y in 3 out of 5 places. What happened? And how can we improve to avoid this in the future?" Utterance A places direct blame on Abraham for failing to attach part x to part y. Conversely, utterance B notifies the team of the error, does not assign direct blame to Abraham, and solicits feedback on how the team can improve. If the robot continues to communicate in ways like in utterance B, the team would be more likely to solve problems together and regularly give constructive feedback to one another to improve performance.

Through our research, we explore the influence robots can have on social collaboration within a human-robot team and develop ways in which a robot can intelligently contribute to the group's social support. Our specific goals are articulated in the following problem statement.

Problem Statement. As we, robotics researchers, look to equip robots with skills that enable them to work well within human-robot teams, it is important that we focus on developing tools and algorithms for robots to enhance the group's social support. In our work, we (1) demonstrate our exploration of methods for a robot to strengthen social support within a human-robot group and (2) propose a multimodal sensing algorithm to detect relevant social dynamics within a human-robot group so that a robot can strengthen social support more strategically.

II. Background

Research in human-robot interaction (HRI) has begun to explore the influence of a robot's behavior on individuals in a human-robot group context. A robot's gaze in human-robot groups has been shown to alter people's turn-taking behaviors [11], conversational roles [9], and content recall [8]. Physical orientation of a robot can effect the number of followers it has when guiding people in a shopping mall [10] as well as the physical proximity of people to a robot [16]. These physical social signals from a robot – gaze and physical orientation – although may be somewhat subtle, clearly show a significant influence on human behavior in a human-robot group.

In addition to physical behavior, several studies have shown that a robot's social behavior affects human behavior and opinion. Wang et al. [17] found that whether a robot expressed its opinion explicitly or implicitly influenced whether team members changed their opinions to match those of the robot. Tan et al. [15] demonstrated that when a human confederate verbally and physically abused a robot, human participants differed in their responses to the abuse (moving the robot to its original position) and had different perceptions of the severity of the abuse dependent on the robot's reaction to being abused. Correia et al. [1] showed that in a setting where two human participants played a game with two robots, where each human partnered with one robot, participants expressed greater liking, trust, and group identification with a robot who expressed group emotion than a robot expressing individual emotion. Despite the growth of work examining a robot influence on individual's reactionary behavior and opinions in a human-robot group, there has not been a large focus on how robots can also affect the way that humans within the group interact with one another. The work we later present contributes to this growing area of research.

III. Exploring Methods of Shaping Social Support in Human-Robot Teams

As a first investigation into how robots can enhance social collaboration between people, we investigated the influence of task-focused vs relationship-focused discussion questions in a collaborative rocket-building task between two children with a robot companion [12]. We discovered that children who were asked task-focused questions by the robot had higher performance scores in the collaborative game than the children who were asked relationship-focused questions by the robot, however, had a lower perception of their own performance than the children who were asked relationship-focused questions by the robot. This result was somewhat surprising, in that we expected the children with the relationship-focused questions to have better teamwork that would lead to a higher performance, however, because our participants were children and do not focus as much on a task at hand as adults, it is possible that the increased performance with the children who received task-focused questions helped the pair to stay focused on the objective and achieving the goal.

Fig. 1. Through conducting a human-subjects study we demonstrated that a robot's vulnerable utterances positively influence human-to-human communication within a human-robot team [13].

In order to probe more deeply into how a robot influences the dynamics within the team, we designed a study to investigate how vulnerable statements made by a robot influence human teammates' reactions to errors made in a collaborative game played by three adults and one robot [13], see Figure 1. In this study, human participants either interacted with a robot who made a neutral or vulnerable comment at the end of each of 30 rounds. In order to assess the dynamics of the group, we examined the responses of human participants to rounds of the game where one of the participants had made an error - a moment where the team's trust of one another would be tested and revealed. Findings showed that human team members in a group with a robot making vulnerable comments as opposed to neutral comments were more likely to explain the mistake, console, and laugh with their fellow human team members in the aftermath of a team member's error in the game. This finding demonstrated the 'ripple effect' of the robots behavior: the human team members displayed more trusting behavior toward their fellow human team members if the robot they interacted with modeled trust and vulnerability.

In current work, we are interested in investigating which behavior strategies of a robot are most effective at including a member of the human-robot team that feels excluded. Within teams, it is common for team members have more commonalities with and affinity towards some team members over others. However, these divisions can easily work against positive group dynamics that drive optimal team performance. We have constructed an experiment to test whether a robot can help include a teammate who feels excluded from the group. In this experiment, the robot addresses more affirming statements like, "we should bring the screwdriver, good idea Seema," as well as more active listening backchannels like 'yeah' and 'mmhmm' toward those who feel isolated in the group. We are interested to see if these affirming statements by the robot help include the group member that initially feels excluded in the group, measured by post-experiment survey questionnaires, the time spent talking of the participants, and each individual's influence on the choices made by the team.

IV. Detecting Relevant Social Dynamics within Human-Robot Groups

The current state-of-the-art measurements of social support in a team involve administering surveys to team members, as is the case with measuring psychological safety [2], social sensitivity [18], and emotional intelligence [14]. Some work in natural language processing has explored the development of models that can automatically detect agreement [4], consistency of understanding [7], and when key decisions are made [6] during meetings. However, to our knowledge, there exist no tools of measuring and detecting the level of social support exhibited by team members toward one another in real time.

We are interested in building a model that takes as inputs real-time audio and video data from human members of a group and predicts the group's current level of social support. Based on prior work, we believe that detecting verbal conversational features like agreement [4], backchannels [5], and disfluencies and self-repairs [3] will be helpful in predicting social support. We plan on training and testing our model on meeting corpus data, such as the AMI meting corpus[1], and in-person data we collect ourselves.

After building this model to detect a group's current level of social support, we plan on running a human subjects experiment where we use the model's results to inform robot actions within a human-robot team to help improve social support when it is detected to be low. Behaviors the robot could employ include: supporting the ideas of team members, seeking feedback from team members on how to improve, asking for help and offering help, discussing errors in a way that does not blame team members, and soliciting input from any team members who might be in some way excluded from the group.

Acknowledgment

This work was supported by the National Science Foundation award number 104893.

[1] http://groups.inf.ed.ac.uk/ami/corpus/

REFERENCES

[1] Filipa Correia, Samuel Mascarenhas, Rui Prada, Francisco S Melo, and Ana Paiva. Group-based emotions in teams of humans and robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 261–269. ACM, 2018.

[2] Amy Edmondson. Psychological safety and learning behavior in work teams. *Administrative science quarterly*, 44(2):350–383, 1999.

[3] Felix Gervits, Kathleen Eberhard, and Matthias Scheutz. Disfluent but effective? a quantitative study of disfluencies and conversational moves in team discourse. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3359–3369, 2016.

[4] Dustin Hillard, Mari Ostendorf, and Elizabeth Shriberg. Detection of agreement vs. disagreement in meetings: Training with unlabeled data. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology: companion volume of the Proceedings of HLT-NAACL 2003–short papers-Volume 2*, pages 34–36. Association for Computational Linguistics, 2003.

[5] Malte Jung, Jan Chong, and Larry Leifer. Group hedonic balance and pair programming performance: affective interaction dynamics as indicators of performance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 829–838. ACM, 2012.

[6] Been Kim and Cynthia Rudin. Learning about meetings. *Data mining and knowledge discovery*, 28(5-6):1134–1157, 2014.

[7] Joseph Kim and Julie A Shah. Improving team's consistency of understanding in meetings. *IEEE Transactions on Human-Machine Systems*, 46(5):625–637, 2016.

[8] Bilge Mutlu, Jodi Forlizzi, and Jessica Hodgins. A storytelling robot: Modeling and evaluation of human-like gaze behavior. In *2006 6th IEEE-RAS International Conference on Humanoid Robots*, pages 518–523. Citeseer, 2006.

[9] Bilge Mutlu, Toshiyuki Shiwa, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. Footing in human-robot conversations: how robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 61–68. ACM, 2009.

[10] Masahiro Shiomi, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. A larger audience, please!: encouraging people to listen to a guide robot. In *Proceedings of the 5th ACM/IEEE international conference on Human-robot interaction*, pages 31–38. IEEE Press, 2010.

[11] Gabriel Skantze, Martin Johansson, and Jonas Beskow. Exploring turn-taking cues in multi-party human-robot discussions about objects. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, pages 67–74. ACM, 2015.

[12] Sarah Strohkorb, Ethan Fukuto, Natalie Warren, Charles Taylor, Bobby Berry, and Brian Scassellati. Improving human-human collaboration between children with a social robot. In *Robot and Human Interactive Communication (RO-MAN), 2016 25th IEEE International Symposium on*, pages 551–556. IEEE, 2016.

[13] Sarah Strohkorb Sebo, Margaret Traeger, Malte Jung, and Brian Scassellati. The ripple effects of vulnerability: The effects of a robot's vulnerable behavior on trust in human-robot teams. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 178–186. ACM, 2018.

[14] Elizabeth Stubbs Koman and Steven B Wolff. Emotional intelligence competencies in the team and team leader: A multi-level examination of the impact of emotional intelligence on team performance. *Journal of Management Development*, 27(1):55–75, 2008.

[15] Xiang Zhi Tan, Marynel Vázquez, Elizabeth J Carter, Cecilia G Morales, and Aaron Steinfeld. Inducing bystander interventions during robot abuse with social mechanisms. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 169–177. ACM, 2018.

[16] Marynel Vázquez, Elizabeth J Carter, Braden McDorman, Jodi Forlizzi, Aaron Steinfeld, and Scott E Hudson. Towards robot autonomy in group conversations: Understanding the effects of body orientation and gaze. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pages 42–52. ACM, 2017.

[17] Lin Wang, Pei-Luen Patrick Rau, Vanessa Evers, Benjamin Krisper Robinson, and Pamela Hinds. When in rome: the role of culture & context in adherence to robot recommendations. In *Proceedings of the 5th ACM/IEEE international conference on Human-robot interaction*, pages 359–366. IEEE Press, 2010.

[18] Anita Williams Woolley, Christopher F Chabris, Alex Pentland, Nada Hashmi, and Thomas W Malone. Evidence for a collective intelligence factor in the performance of human groups. *science*, 330(6004):686–688, 2010.