

Robot Moderation of a Collaborative Game: Towards Socially Assistive Robotics in Group Interactions

Elaine Short¹ and Maja J Matarić¹

Abstract—This paper presents an algorithm for enabling a robot to act as the moderator in a group interaction centered around a tablet-based assembly game. The algorithm uses one of two different objective functions: one intended to be “performance equalizing”, wherein the robot attempts to equalize scoring among users, and another intended to be “performance reinforcing”, wherein the robot attempts to help the group score as many points as possible. In an evaluation study with ten groups of three participants, we found that the “performance equalizing” algorithm improved task performance and reduced group cohesion, while the “performance reinforcing” algorithm improved group cohesion and reduced task performance.

I. INTRODUCTION

Socially Assistive Robotics (SAR) addresses the challenges of enabling robots to use social interaction to help users achieve goals, especially in activities related to health, wellness, and education [1]. We extend SAR to multi-party interactions by considering the problem of *moderation*, where a robot assists a group to achieve a shared goal. In this paper, the moderation algorithm is presented and evaluated in a collaborative game played by undergraduate students; however, the moderation approach is general and can be used in diverse domains, such as facilitating inter-generational interactions between older adults and their families [2], and children’s group learning, both areas of our ongoing work.

We provide an overview of the moderation approach in a collaborative assembly game, as well as an evaluation study of 30 participants in 10 groups examining the effect of such SAR moderation, comparing two different objective functions: performance-equalizing and performance-reinforcing. We show that the performance-equalizing function increases task performance and decreases group cohesion compared to the performance-reinforcing version. We present an additional analysis suggesting that this affect is at least partially caused by the frequency with which the robot addressed the participants. These results demonstrate that a robot moderator can affect social and task features of group interactions.

II. RELATED WORK

In SAR domains, robots may act as tutors, coaches, or peers to achieve desired outcomes with human user(s), typically through repeated interactions over time. SAR targets those interactions in complex settings such as homes, clinics, and schools, where multiple users may frequently interact in

groups with the single robot. For example, researchers found that children in school preferred group interactions with a robot [3]. Furthermore, group interactions are a standard tool in learning, coaching, and therapy, such as using peers to help children with autism integrate into the playground environment [4], and leveraging family interactions to improve health outcomes of older adults [5], [6].

Prior work in group human-robot interaction has been conducted primarily from two perspectives: first, from the perspective of explicit collaborative human-robot teams (e.g., in manufacturing and the military), and second, from the perspective of implicit loosely interacting human-robot groups. In the first context, the primary purpose of the robot(s) is to provide unique capabilities to the team, such as a drone providing aerial footage of a scene [7]. In such work, researchers consider features such as workload, task performance, situational awareness, and preferences about robot autonomy [8]. In the second context, the robot serves specific social features of the interaction, such as shaping user roles [9], choosing an addressee [10], or serving multiple users [11]. In this second area, researchers consider features such as user speech, gaze, social relationships, and preferences about robot agency. Because SAR is inherently both social and goal-oriented, our work aims to contribute to both of these perspectives, by developing algorithms that enable a robot to act as a *moderator* in multi-party human-robot interactions. In our work, the robot has unique capabilities it shares with users, as in human-robot teaming, but uses them to manage social aspects of the interaction, as in social human-robot interaction.

III. MODERATION ALGORITHM

We developed a moderation algorithm for collaborative tasks in which users have individual subgoals whose completion requires collaboration with others, and where the goal the moderator is supporting at a given time is chosen according to individual user’s performance, based on one of two different objective functions described below.

To moderate the interaction, the robot provides assistance with g_p , a goal belonging to user p , such that

$$g_p = \operatorname{argmin}(U(g)) \quad (1)$$

where g is some goal not achieved in the current game state, and $U(g)$ is the objective function. There are two versions of the model: 1) the performance-reinforcing version, in which the robot chooses a goal that is closest to completion; and 2) the performance-equalizing version, in which the

*This work is supported by the National Science Foundation (NRI IIS-1208500, Expeditions in Computing IIS-1139148, CNS-0709296, CBET-1548502, REU Supplements, and GRFP).

¹All authors are at, University of Southern California, Los Angeles, California, USA elaine.g.short@usc.edu mataric@usc.edu




Fig. 1. Screenshot of one user's game screen

Algorithm 1 Request Generation Algorithm

g the goal, $s_g \in g$ an open subgoal of g , T_u the tokens held by u , $p(t)$ properties already applied to token t

if $p(t) = s_g$ for some $t \in T_u$, t not in the goal area **then**
 Choose Request 5
else
 if $|T_u| = 0$ for all users that can help s_g **then**
 Choose a random s_g
 if It is the first time the moderator helped s_g **then**
 Choose Request 1
 else
 Choose Request 3
 end if
 else
 Let t be the token maximizing $p(t) \cap s_g$
 if User can apply a property in $s_g - p(t)$ **then**
 Choose Request 6
 else
 if It is the first time the moderator helped s_g **then**
 Choose Request 2
 else
 Choose Request 4
 end if
 end if
 end if
end if

robot chooses a goal from the user who has completed the fewest number of goals for the team so far (i.e., the poorest-performing user). Thus, if $s(p)$ is the number of goals achieved by user p , n is the total number of actions needed to complete the goal, and $r(g_p)$ is the remaining distance to the goal g_p , measured by the number of actions needed to achieve the goal g_p :

$$U_{reinforcing}(g_p) = r(g_p) \quad (2)$$

and

$$U_{equalizing}(g_p) = s(p) + (1 - (r(g_p)/n)) \quad (3)$$

Once a goal is chosen, the moderator selects the action that most quickly achieves a sub-goal of the goal, per the behavior generation algorithm shown in Algorithm 2. The 20-second

TABLE I
ROBOT-PARTICIPANT REQUESTS

- | | |
|---|---|
| 1 | Player $\langle name \rangle$ needs a $\langle property \rangle$ token. |
| 2 | Player $\langle name \rangle$ can make $\langle property \rangle$ tokens. |
| 3 | Player $\langle name \rangle$ should give player $\langle name \rangle$ a $\langle property \rangle$ token. |
| 4 | Player $\langle name \rangle$ should give player $\langle name \rangle$ a token to make [into a] $\langle property \rangle$. |
| 5 | Player $\langle name \rangle$ should put the $\langle properties \rangle$ token in the goal area. |
| 6 | Player $\langle name \rangle$ should make a $\langle property \rangle$ token. |

Algorithm 2 Robot Behavior Generation Algorithm

while Elapsed time $< 6\text{min}$ **do**
 Choose goal minimizing $U(g)$
 Choose an appropriate request using Algorithm 1
 Wait 20s
end while

interval used in the algorithm was validated in a pilot study with three groups, and was found by all participants to be an acceptable moderator input rate.

To study this model of robot behavior in an interactive and enjoyable collaborative activity, we developed a cooperative game loosely inspired by SPACETEAM¹, a popular mobile game in which a steady stream of actions needing to be completed is provided to the team, with both information and abilities distributed among the players, so that players need to communicate in order to perform the actions. This game provides a rich context for collaborative group interactions, with opportunities for players to help each other and a high level of entertainment value. Our game uses the same idea of mismatched distribution of needs and resources, but in a simplified context. The goal of the game is for the team to score as many points as possible by creating pairs of colored shapes from gray squares. Each user sits in front of an individual angled tablet that displays the game interface visible only to that user, as seen in Figure 1. None of the users can complete all of their goals without help from some other user, so collaboration is required.

In the game, $s(p)$ is the number of points scored by the user, n is the total number of properties in the goal (in this game, two properties for each of two parts, for a total of four), and $r(g_p)$ is the number of properties remaining to be applied to achieve the goal. To help the users, the robot chooses the appropriate request from Table I, using Algorithm 1, providing information about capabilities or needs the first time it helps with a goal, then providing specific instructions on subsequent attempts.

IV. METHODOLOGY

To evaluate the effects of the two different objective functions for moderation, we conducted a double-blind repeated-measures study with groups of three participants. Participants engaged in “training” sessions with a robot moderator and “testing” sessions without moderation, to examine whether

¹<http://spaceteam.ca/>

the moderator has effects on group performance after moderation. Prior work found that social and task goals of a group may be opposed to each other [12]. In this work, we address the question of whether the performance-equalizing (PE) moderator, while making short-term sacrifices to task performance, might better prepare the group for long-term collaboration than the performance-reinforcing (PR) moderator. Our study addresses the following hypotheses related to this trade-off:

- H1: Task performance (as measured by points scored) will be better with the PR moderator.
- H2: Participants will report greater group cohesion with the PE moderator.
- H3: Participants will speak more with the PE moderator.
- H4: Participants will look around more (at the other participants) in sessions with the PE moderator.

Taking social goals into account maintains intra-group relationships and improves long-term performance. Thus we include the following additional hypotheses:

- H5: Participants will speak more in the testing session following the PE moderator interaction.
- H6: Participants will look around more in the testing session following the PE moderator interaction.
- H7: Participants will have better performance in the testing session following the PE moderator interaction.

Additionally, to understand the participants' views of the specific robot, as well as robots in general, and if those views change after interacting with a moderator robot, we measured self-reported attitudes towards robots in general, and technology acceptance of the specific robot system to address the following hypotheses:

- H8: After moderation, participants will report a more positive attitude towards robots than reported before the interaction.
- H9: Participants will report greater acceptance of the robot as a technology after more sessions with the robot.

A. Experimental Procedure

After completing informed consent, participants were given an initial set of questionnaires to complete (described in Section IV-B), then invited to sit down at the table, as seen in Figure 2(a). Participants were told that they would be playing five six-minute sessions of the game with the robot, and that the robot would provide initial instructions. The experimenter then left the room, and the robot gave a set of simple verbal instructions for the game. After six minutes from the start of the game, the robot said that time was up and the experimenter returned to the room. This procedure was repeated four more times at the higher difficulty, alternating moderated and unmoderated sessions, and with the system randomizing the order of the two utility functions, so that the experimenter was blind to what type of moderation was taking place. After the third and final sessions, both unmoderated, the participants completed additional questionnaires, described in Section IV-B below.

The study used the SPRITE robot [13], shown in Figure 2(b). The robot is approximately 30 cm tall (without fur),




Fig. 2. Robot and setup used in study

has a 10 cm range of motion, and uses a mobile phone to display simple facial expressions. In this study, the robot was covered by soft fur, as seen in Figure 2(b).

B. Measures

Before the first interaction, participants were asked to complete a demographics questionnaire, the interaction and social role subscales of the Negative Attitudes towards Robots (NARS) questionnaire ([14] translated by [15]), and a 10-question Big Five personality inventory [16]. The latter questionnaire will be used to develop models for personalized moderation, and is not analyzed as part of this study.

After each interaction with the robot, participants completed the Group Cohesiveness Scale [17], the Unified Theory of Acceptance and Use of Technology (UTAUT) questionnaire [18] as modified by [19], and repeated the NARS questionnaire. Additionally, after the final interaction, participants completed a set of questions based on those used in the UTEP-ICT dataset [20]; this questionnaire includes questions about the behavior of the other participants in the interaction, additional demographic details, and questions about the relationships between participants. Long-form questions were asked about perceptions of the task and space was provided for any additional comments.

In addition to the self-report measures, a number of behavioral measures were collected. First, participants' actions in the game were recorded, along with the number of points scored by each participant. Three USB cameras recorded each participant's face, and the OpenFace library was used to annotate video from the cameras for head pose [21], [22]. Headset microphones were used to obtain audio recordings of each participant's voice, and automatic voice activity detection was used to annotate the data, using the method from Van Segbroeck, Tsiartas, and Narayanan [23].

V. RESULTS

Thirteen groups of three participants were recruited for this study from university freshman computer science classes and department mailing lists. Two groups were excluded from the analysis due to malfunctions in the robot's programming. One group was run with only two participants due to a no-show, and is also excluded from the analysis. Of the ten remaining groups (30 participants), there were 19 male and 11 female participants. 17 identified as Asian (including




Fig. 3. Score and group cohesion in the interaction. (** $p < 0.05$)

South Asian), 6 as white, 2 Hispanic/Latin, and 4 mixed-race or other. Participants’ average age was 18.97 years ($SD = 1.067$). Participants had limited prior experience with robots, rating their prior exposure to any kind of robot as 1.93 on a 4-point Likert scale ($SD = 1.01$). 27 of the participants had lived in the United States longer than any other country, for an average of 17.87 years ($SD = 3.14$), the remaining 3 participants had lived in East Asia the longest. All participants spoke fluent English. Six of the groups saw the performance-reinforcing moderator first, and four groups saw the performance-equalizing first. Due to technical issues, one group did not have audio or video recording, and one group did not have audio recording; those data were not included in the video or audio analyses.

A. Performance and Group Cohesion

To test hypothesis $H1$, we compared participants’ performance during the moderated conditions with a paired t-test, and found a significant difference in participants’ scores in the two conditions [$t(29) = -2.0997, p < 0.05$]. However, the scores were higher in the performance-equalizing condition ($M = 6.60, SD = 1.83$) than in the performance-reinforcing condition ($M = 6.03, SD = 1.81$). Relative to hypothesis $H7$, the score for performance-equalizing in the post-moderation session ($M = 7.133, SD = 1.81$) was also higher than the score for performance-reinforcing in the post-moderation session ($M = 6.83, SD = 1.98$), as seen in Figure 3(a). Testing hypothesis $H2$ with a paired t-test, we found a significant difference in participants’ group cohesion scores between conditions [$t(29) = 2.3794, p < 0.05$]. Group cohesion, however, was higher in the performance-reinforcing condition ($M = 3.39, SD = 0.44$) than in the performance-equalizing condition ($M = 3.26, SD = 0.46$). These results can be seen in Figure 3(b).

B. Speech and Gaze

To examine participant speech during the interaction, we first filtered the speech by times when a given participant was the only speaker, eliminating cross-microphone speech. To test $H3$, we compared speech during the two post-moderation sessions with a paired t-test and found no significant differences. There were also no significant differences in speech during the moderation sessions ($H3$). To determine how participants’ head orientation changed during the sessions,




Fig. 4. Average standard deviation of head pitch and yaw (** $p < 0.05$, * $p = 0.06$)

we calculated the variance on head pitch and roll for each session for each participant, then performed the statistical tests. This allowed us to normalize for the mean head position, since each participant sat at a slightly different angle to the camera. There were no significant differences between the post-moderation conditions ($H6$), however, during the moderation conditions ($H4$), there was significantly higher variance in participants’ head pose in the yaw direction [$t(26) = 2.5553, p < 0.05$], and a trend towards significance in the pitch direction [$t(26) = 1.9555, p = 0.06$], with more variance in the performance-reinforcing condition (head pitch: $M = 0.174, SD = 0.086$, head yaw: $M = 0.182, SD = 0.074$) than in the performance-equalizing condition (head pitch: $M = 0.153, SD = 0.072$, head yaw: $M = 0.165, SD = 0.060$), as seen in Figure 4.

C. Attitude Towards The System

Using an ANOVA to test $H8$, we found a significant effect of order on the NARS Negative “Attitudes toward Situations and Interactions with Robots” subscale [$F(2, 58) = 14.70359, p < .001$], but no differences in the “Negative Attitudes toward Social Influence of Robots” and “Negative Attitudes toward Emotions in Interaction with Robots” subscales. A post hoc test with Bonferroni correction showed a significant difference between the baseline ($M = 1.41, SD = 0.44$) and post-study ($M = 1.72, SD = 0.55$) scores on the “Attitudes toward Situations and Interactions with Robots” subscale ($p_{adj} = 0.044$), and a difference trending on significance between the baseline and mid-study ($M = 1.70, SD = 0.46$) scores ($p_{adj} = 0.058$). Analyzing the UTAUT questionnaire results to test $H9$, we observed that five of the subscales (Facilitating Conditions, Perceived Enjoyment, Perceived Sociability, Social Presence, and Trust) had modest increases from the mid-experiment to post-experiment questionnaires, and none of the subscales had any decreases from mid- to post-experiment. However, after correcting for the number of tests, there were no significant differences.

D. Effect of Moderator Speech

To better understand the differences between conditions, we performed an additional analysis examining the effects of the number of times the robot addressed a participant, calculated by the number of times the robot spoke that participant’s name. Based on prior work suggesting that




Fig. 5. Changes in participant behavior as a result of changes in the number of times participants were addressed by the robot

participants in a robot-moderated multi-party interaction who were not addressed by the robot had lower group cohesion scores [24], we included two additional hypotheses:

- H10: Robot speech towards participants will be correlated with greater group cohesion.
- H11: Robot speech towards participants will be correlated with greater helpfulness towards other participants.

We examined correlations between this value and both self-reported group cohesion and participant helpfulness, as measured by the number of subgoals that a participant achieved for someone else. We found positive correlations with robot mentions of a participant and group cohesion scores ($r(52) = 0.28, p < 0.05$) and helpfulness ($r(52) = 0.29, p < 0.05$).

Furthermore, to control for the fact that participants' behavior affected how the robot chose whom to address, we examined the *change* in the cohesion and helpfulness scores based on the *change* in the number of times the robot addressed the participant, between the two different moderation conditions. In that analysis, we found that the change in group cohesion scores was positively correlated with the change in the number of times the participant was addressed by the robot ($r(25) = 0.44, p < 0.05$) and that the change in participant helpfulness was positively correlated with the change in the number of times the participant was addressed by the robot ($r(25) = 0.41, p < 0.05$).

E. Attitudes Towards Moderated Activity

The game itself was entertaining to the participants, with 10 out of the 30 participants specifically stating that the game was "fun" when asked to write about their favorite part of the task or in the additional comments section. However, several participants felt that they did not pay attention to the robot, or that its advice was not as helpful as it could have been. This varied between conditions, with groups taking more of the advice given by the robot in the performance-reinforcing condition ($M = 13.1, SD = 2.13$) than in the performance-equalizing condition ($M = 11.2, SD = 2.74$), with the




Fig. 6. Number of the robot's suggestions taken by the group, out of 18 ($p = 0.067$)

difference trending on significance [$t(9) = 2.0827, p = 0.067$]. This result can be seen in Figure 6.

VI. DISCUSSION

We developed an algorithm that enables a socially assistive robot to moderate a group activity based on task features and evaluated the algorithm in an assembly game task, with two different utility functions (performance-reinforcing and performance-equalizing), predicting that improved task performance would come at the cost of decreased social cohesion, and vice versa. The results show that the predicted trade-off did occur, but evidence of improved social features, including increased group cohesion scores and more looking around (as measured by variance in head pose), were found in the performance-reinforcing rather than performance-equalizing condition, in contrast to the predictions made in *H1*, *H2*, and *H4*. Further analysis found evidence that this result may have been caused by differences in how the robot addressed the participants: in the performance-equalizing condition the robot was more likely to repeatedly address the participant with the lowest performance. We found that the more the robot spoke to participants, the higher group cohesion they reported (supporting *H10*), and the more they helped the other participants in the group (supporting *H11*). However, the robot's assistance did come at a cost: the performance of the group was higher in the unmoderated testing sessions (*H6* and *H7* unsupported), likely due to fact that the game required frequent verbal coordination between players, and any robot speech decreased the amount of time the players could discuss their strategy, although no significant differences in speech were found between sessions (*H3* and *H5* unsupported). Further research is needed to determine whether the long-term social benefits of a robot moderator offset short-term decreases in task performance. Additionally, participants attitudes towards robots on the "Attitudes toward Situations and Interactions with Robots" subscale of the NARS became more negative (*H8* unsupported), and their scores on the UTAUT inventory did not change (*H9* unsupported), raising the question of whether the robot increases group cohesion by increasing participants' identification with the other human users in contrast to the robot moderator. Finally, the game task we used was very successful at getting participants involved and excited about working together as a group. The robot's approach to moderating the interaction resulted in participants completing over half of the robot's

suggestions, although there are opportunities for improving the timing and salience of the robot's suggestions.

VII. CONCLUSION

This work presented an algorithm for enabling a robot to act as a moderator, applied it to a collaborative game task, and evaluated it in that task context. Although, as expected, we found trade-offs between social cohesion and task performance, the two utility functions for the moderation algorithm resulted in the opposite of the predicted effects. The “performance-equalizing” utility function increased task performance and decreased group cohesion, and “performance-reinforcing” utility function decreased task performance and increased group cohesion, an effect that may have been mediated by the frequency with which the robot addressed the different participants in the interaction. This work provides important evidence that a robot moderator is able to positively affect the social dynamics of a group, even without the use of speech recognition, and that *to whom* the moderator is talking and *how often* it is talking may be key features to include in future moderation algorithms. Such moderation algorithms have important applications in multi-party socially assistive robotics, where a robot helps a group achieve some goal or goals relating to the health, wellness, or education of the group members.

ACKNOWLEDGMENT

The authors thank David Traum for providing advice and access to the UTEP-ICT dataset, Dale Short and George Schaertl for their assistance with robot design and assembly, and Rhianna Lee, Kara Douville, and Shayna Goldberger for their help with robot behavior and preparation for the user study.

REFERENCES

- [1] D. Feil-Seifer and M. J. Matarić, “Defining socially assistive robotics,” in *Proceedings of the IEEE 9th Int. Conf. on Rehabilitation Robotics*, vol. 2005, Chicago, IL, June 2005, pp. 465–468.
- [2] E. Short, K. Swift-Spong, H. Shim, K. M. Wisniewski, D. Kim Zak, S. Wu, E. Zelinski, and M. J. Matarić, “Understanding Social Interactions with Socially Assistive Robotics in Intergenerational Family Groups,” in *IEEE Int. Symposium on Robot and Human Interactive Communication (RO-MAN 2017)*. Lisbon, Portugal: IEEE, 2017.
- [3] T. Kanda, R. Sato, N. Saiwaki, and H. Ishiguro, “A Two-Month Field Trial in an Elementary School for Long-Term HumanRobot Interaction,” *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 962–971, Oct. 2007.
- [4] K. Pierce and L. Schreibman, “Increasing complex social behaviors in children with autism: effects of peer-implemented pivotal response training,” *Journal of Applied Behavior Analysis*, vol. 28, no. 3, pp. 285–295, 1995.
- [5] T. E. Seeman, “Health promoting effects of friends and family on health outcomes in older adults,” *American journal of health promotion : AJHP*, vol. 14, no. 6, pp. 362–70.
- [6] R. A. Jablonski, D. Reed, and M. L. Maas, “Care Intervention for Older Adults with Alzheimer’s Disease and Related Dementias: Effect of Family Involvement on Cognitive and Functional Outcomes in Nursing Homes,” *Journal of Gerontological Nursing*, vol. 31, no. 6, pp. 38–48, June 2005.
- [7] R. W. Beard, T. W. McLain, M. A. Goodrich, and E. P. Anderson, “Coordinated target assignment and intercept for unmanned air vehicles,” *IEEE transactions on robotics and automation*, vol. 18, no. 6, pp. 911–922, 2002.
- [8] M. Gombolay, A. Bair, C. Huang, and J. Shah, “Computational Design of Mixed-Initiative Human-Robot Teaming that Considers Human Factors: Situational Awareness, Workload, and Workflow Preferences,” *Int. Journal of Robotics Research*, 2017.
- [9] D. Bohus and E. Horvitz, “Facilitating multiparty dialog with gaze, gesture, and speech,” in *Int. Conf. on Multimodal Interfaces and the Workshop on Machine Learning for Multimodal Interaction on - ICMI-MLMI ’10*. New York, New York, USA: ACM Press, Nov. 2010, p. 1.
- [10] A. Yamazaki, K. Yamazaki, T. Ohyama, Y. Kobayashi, and Y. Kuno, “A Techno-Sociological Solution for Designing a Museum Guide Robot : Regarding Choosing an Appropriate Visitor,” in *Proceedings of the 7th annual ACM/IEEE Int. Conf. on Human-Robot Interaction (HRI 2012)*. Boston, MA: ACM Press, Mar. 2012, pp. 309–316.
- [11] M. Foster, A. Gaschler, and M. Giuliani, “Two people walk into a bar: Dynamic multi-party social interaction with a robot agent,” in *Multimodal interaction*. New York, New York, USA: ACM Press, Oct. 2012, pp. 3–10.
- [12] S. Viller, “The Group Facilitator: A CSCW Perspective,” in *Proceedings of the Second European Conf. on Computer-Supported Cooperative Work ECSCW 91*, L. Bannon, M. Robinson, and K. Schmidt, Eds., Amsterdam, The Netherlands, 1991, pp. 81–95.
- [13] E. Short, D. Short, Y. Fu, and M. J. Matarić, “SPRITE: Stewart platform robot for interactive tabletop engagement,” Department of Computer Science, University of Southern California,” Tech Report, Mar 2017. [Online]. Available: <http://robotics.usc.edu/publications/967/>
- [14] T. Nomura, T. Suzuki, T. Kanda, and K. Kato, “Measurement of negative attitudes toward robots,” *Interaction Studies*, vol. 7, no. 3, pp. 437–454, 2006.
- [15] C. Bartneck, T. Nomura, T. Kanda, T. Suzuki, and K. Kennsuke, “A cross-cultural study on attitudes towards robots,” in *Proc. HCI Int. Conf.*, Las Vegas, NV, Jan. 2005.
- [16] B. Rammstedt and O. P. John, “Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German,” *Journal of Research in Personality*, vol. 41, no. 1, pp. 203–212, Feb. 2007.
- [17] T. Wongpakaran, N. Wongpakaran, R. Intachote-Sakamoto, and T. Boripuntakul, “The Group Cohesiveness Scale (GCS) for psychiatric inpatients,” *Perspectives in Psychiatric Care*, vol. 49, no. 1, pp. 58–64, Jan. 2013.
- [18] F. D. Davis, “Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology,” *Management Information Systems Quarterly*, vol. 13, no. 3, p. 319, Sept. 1989.
- [19] M. Heerink, B. Kröse, V. Evers, and B. Wielinga, “Assessing Acceptance of Assistive Social Agent Technology by Older Adults: the Almere Model,” *Int. Journal of Social Robotics*, vol. 2, no. 4, pp. 361–375, Sept. 2010.
- [20] D. Herrera, D. Novick, D. Jan, and D. Traum, “The UTEP-ICT cross-cultural multiparty multimodal dialog corpus,” in *Multimodal Corpora Workshop: Advances in Capturing, Coding and Analyzing Multimodality (MMC 2010)*, Valletta, Malta, 2010.
- [21] T. Baltrusaitis, P. Robinson, and L.-P. Morency, “OpenFace: An open source facial behavior analysis toolkit,” in *2016 IEEE Winter Conf. Appl. Comput. Vis.* IEEE, mar 2016, pp. 1–10.
- [22] —, “Constrained Local Neural Fields for Robust Facial Landmark Detection in the Wild,” in *2013 IEEE Int. Conf. Comput. Vis. Work.* IEEE, dec 2013, pp. 354–361.
- [23] M. Van Segbroeck, A. Tsiartas, and S. Narayanan, “A robust frontend for VAD: exploiting contextual, discriminative and spectral cues of human voice,” in *INTERSPEECH*, 2013, pp. 704–708.
- [24] E. Short, K. Sittig-Boyd, and M. J. Matarić, “Modeling Moderation for Multi-Party Socially Assistive Robotics,” in *IEEE Int. Symp. Robot Hum. Interact. Commun. (RO-MAN 2016)*. New York, NY: IEEE, 2016.