Artificial Intelligence for Human-Robot Interaction
Papers from the AAAI 2015 Fall Symposium

# Towards Robot Moderators: Understanding Goal-Directed Multi-Party Interactions

**Elaine Short and Maja J. Matarić**
University of Southern California
Department of Computer Science
Los Angeles, California, USA

## Introduction

Socially Assistive Robotics (SAR) is a growing field dedicated to developing models and algorithms that enable robots to help people achieve goals through social interaction (Feil-Seifer and Matarić 2005). Prior work in this field has focused on one-on-one interactions, but there is interest in extending this work to multi-party interactions. We contribute to the study of multi-party SAR by defining the role of *moderator*, an agent that is responsible for directing an interaction, but is not necessarily directly participating in the task. We present a formalization of the task of *moderation* as the process by which a goal-directed multi-party interaction is regulated via manipulation of interaction resources, including both *physical* resources, such as an object or a tool, and *social* resources, such as the conversational floor or participants' attention. Finally, we present preliminary results of an analysis of self-moderated multi-party human-human interaction that support several of the underlying assumptions of this formalization.

## Background

Prior work in SAR has primarily focused on one-on-one interactions, although some work has been done with groups of participants interacting with individual robots (Yamazaki et al. 2012) or pairs of robots (Leite et al. 2015). Matsuyama et al. developed a robot that could facilitate engagement in multi-party interactions, and showed that a robot moderator was accepted by participants (2015). Furthermore, prior results suggest that people, especially children, in groups are motivated to interact with robots (Kanda et al. 2004). Bohus and Horvitz developed algorithms that enable an embodied virtual agent to direct group interactions in a gameplay scenario, demonstrating that autonomous agents can direct multi-party interactions (2009). Through the computational formalization of the moderator role, our work focuses on the general problem of facilitating goal-directed human-human interactions through the presence of a robot.

## Modeling Approach

We model moderation as a decision-making problem, in which the moderator must choose actions that maximize group performance according to two goals: a *task goal* defining **what** is to be accomplished (e.g., receiving the highest possible score on an assignment) and a *moderation goal* defining **how** the task should be accomplished (e.g., all members of the group participate in solving the problem).

The key insight of our approach to modeling moderation is to focus on the distribution (or allocation) of *interaction resources*, which include both social resources such as the conversational floor or participants' attention, and task resources such as access to a task space or holding task objects. This enables the moderator to plan in the space of resource distributions rather than in the space of all possible interaction states, by separately modeling the effect of resource distribution on the interaction state, and the relationship between the interaction state and goal performance.

This approach to formalizing moderation makes several claims about how human-human interactions are moderated:

1. Interaction goals affect resource allocation;
2. Resource allocation affects interaction state; and
3. Moderation behaviors in human-human interactions act on interaction resources.

## Dataset

To explore these claims, we analyzed data from a subset of a previously-collected dataset: the UTEP-ICT multi-cultural multi-party interaction corpus (Herrera et al. 2010). The data examined in this work include three groups of four individuals engaging in two different tasks each, each lasting approximately 9 minutes: first, invent a name for a stuffed toy; then jointly or sequentially tell a story about the toy.

The participants were video recorded from 6 angles around the room, and audio recorded from lapel microphones. Additional survey data were collected, but the analysis of those data is beyond the scope of this work. A subset of the corpus was annotated to allow an examination of the management of both physical and social resources: the American English groups' toy-naming and story-telling interactions. The data were annotated by undergraduate coders for a number of relevant features to the interaction. In one of the groups, technical issues in audio recording rendered the data unusable for this purpose.

For the remaining groups, for each participant, several features were annotated: speech; gaze (towards another par-

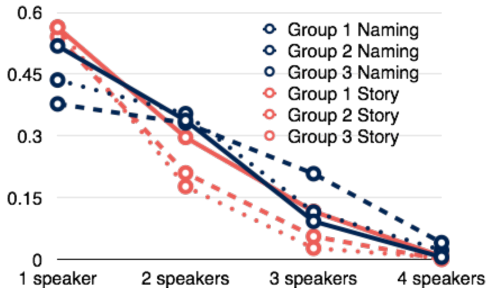


Figure 1: Proportion of time with [n] speakers in 100ms intervals.

ticipant or the toy); and toy-related gestures, including holding the toy, bidding to hand the toy off, or bidding to receive the toy.

Additionally, an informal annotation was performed of types of speech behaviors, such as questions, interruptions, and backchannels. The analysis of these features will be used to develop preliminary models for the task and moderation goals, to define a set of moderation behaviors, and to model the relationship between conversational resources and world state.

## Results

In a preliminary, qualitative, analysis of this dataset, we find some support for the model assumptions:

*Interaction goals affect resource allocation:* We find differences in participant behavior between the naming and story tasks. In the naming task, where the participants are only given the task goal of agreeing on a name, we observed shorter turns and more interruptions. In the story task, where participants were also instructed that they should tell the story jointly, we find longer turns and more even participation. As seen in Figure 1, the story task had more time with only one speaker, and less time with two or more simultaneous speakers.

*Resource allocation affects interaction state:* Additionally, we find that the distribution of resources affects the interaction state: as seen in Figure 2, most participants spoke a greater proportion of the time when they held the toy. This effect was stronger in the story task, suggesting that participants may have been deliberately using the toy to moderate the interaction.

*Moderation behaviors act on interaction resources:* We observed a number of moderation behaviors in the interactions acting on resources such as the conversational floor and the toy. Open (not directed to a particular target) questions were used to release the floor, directed questions were used to pass the floor to an individual, and passing the toy, as discussed above, may have been used to regulate turns.

## Conclusion

In this work, we propose a computational formalization of moderation, from which we generate several hypotheses about moderation in human-human interaction: that the allocation of interaction resources affects the state of the interaction; that moderation behaviors act on interaction resources; and that the goals of the interaction affect resource allocation. In a preliminary analysis of the UTEP-ICT dataset, we find support for further investigation of these hypotheses, encouraging further work using this model. This future work will include formal statistical analysis using a larger dataset, the analysis of gaze patterns in the interaction, and a more detailed study of the effect of holding the toy on the conversational floor. The results of the analysis will be used to parameterize an initial model of moderation for this type of task, to be implemented on a robot and tested with groups of adults in a task analogous to those used in the UTEP-ICT dataset.

## Acknowledgments

This work was supported by the National Science Foundation (Expeditions in Computing IIS-1139148, CNS-0709296, REU Supplements, and GRFP). The authors thank David Traum for providing advice and access to the UTEP-ICT dataset, and Rhianna Lee, Jackson Burgess, Joyce Su, and Stephanie Yi for their help in data coding.

## References

Bohus, D., and Horvitz, E. 2009. Models for Multiparty Engagement in Open-World Dialog. *Comput. Linguist.* 225–234.

Feil-Seifer, D., and Matarić, M. J. 2005. Defining socially assistive robotics. In *Proc. 2005 IEEE 9th Int. Conf. Rehabil. Robot.*, volume 2005, 465–468.

Herrera, D.; Novick, D.; Jan, D.; and Traum, D. 2010. The UTEP-ICT cross-cultural multiparty multimodal dialog corpus. In *Multimodal Corpora Work. Adv. Capturing, Coding Anal. Multimodality (MMC 2010)*.

Kanda, T.; Hirano, T.; Eaton, D.; and Ishiguro, H. 2004. Interactive Robots as Social Partners and Peer Tutors for Children: A Field Trial. *Human-Computer Interact.* 19(1):61–84.

Leite, I.; McCoy, M.; Ullman, D.; Salomons, N.; and Scassellati, B. 2015. Comparing Models of Disengagement in Individual and Group Interactions. In *Proc. Tenth Annu. ACM/IEEE Int. Conf. Human-Robot Interact. - HRI '15*, 99–105. New York, New York, USA: ACM Press.

Matsuyama, Y.; Akiba, I.; Fujie, S.; and Kobayashi, T. 2015. Four-participant group conversation: A facilitation robot controlling engagement density as the fourth participant. *Computer Speech & Language* 33(1):1–24.

Yamazaki, A.; Yamazaki, K.; Ohyama, T.; Kobayashi, Y.; and Kuno, Y. 2012. A Techno-Sociological Solution for Designing a Museum Guide Robot : Regarding Choosing an Appropriate Visitor. In *ACM HRI' 12, Boston, Massachusetts, USA*, 309–316. New York, New York, USA: ACM Press.

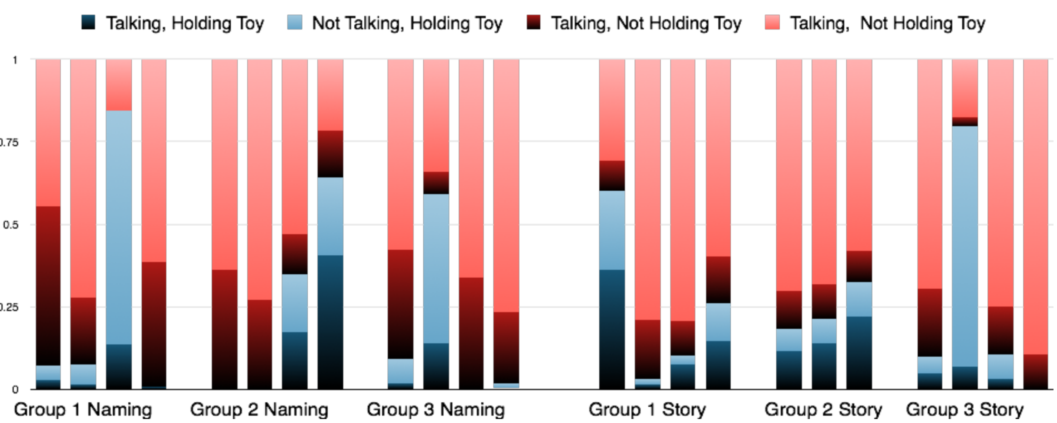


Figure 2: Time spent speaking while holding the toy and while not holding the toy.