

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 (?). 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 computational 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 (?) or pairs of robots (?). Furthermore, prior results suggest that people, especially children, in groups are motivated to interact with robots (?). ? 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. Through the computational formalization of the moderator role, our work focuses directly on facilitating human-human interactions through the presence of a robot, in contrast to prior work in which the human-robot interaction is the focus.

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 (?). The data examined in this work include three groups of four individuals engaging in two different tasks each, each lasting approximately 9 minutes:

1. Invent a name for a stuffed toy;
2. 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, the following features were annotated:

- Speech;




Figure 1: Time spent speaking while holding the toy and while not holding the toy.

- Gaze (towards another participant or the toy); and
- Toy-related gestures: 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 analysis of this dataset, we find 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 2, 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 1, 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: Qualitatively, we observed a number of moderation behaviors in the interactions, including open (not directed to a particular target) questions, used to release the floor, directed questions used to pass the floor to an individual, and the passing of the toy, which, as discussed above, may have been used to regulate the conversational floor.




Figure 2: Proportion of time with [n] speakers in 100ms intervals.

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 these hypotheses, encouraging further work using this model. This future work will include the analysis of gaze patterns in the interaction, as well as 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.