

Multi-Party Socially Assistive Robotics: Defining the Moderator Role

Elaine Short
Univ. of Southern California
elaine.g.short@usc.edu

Maja Mataric
Univ. of Southern California
mataric@usc.edu

1. INTRODUCTION

Many types of multi-party interactions include a *moderator*, an agent that is responsible for directing the interaction, but is not necessarily directly participating in the task. In this work, we formalize 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 who is holding an object or a tool, and *social* resources, such as the conversational floor or participants' attention.

The primary application domain for this model is in the domain of socially assistive robotics (SAR), an area of research in which social robots are deployed and studied in applications where they leverage social interaction in order to enable users achieve some goal [2]. We specifically focus on group learning with children, where the moderator ensures that several children working in a group interact successfully while completing a learning task. This abstract presents a preliminary computational model of this definition of moderation, as well as describing in-progress work towards learning the parameters of the model.

2. BACKGROUND

Although this work is broadly applicable to embodied social agents, we focus our evaluation on a robot as the moderator, leveraging results in human-robot interaction (HRI) suggesting that physically embodied agents are more effective for manipulating human behavior, as well as domain-specific research indicating that physically embodied interactions have benefits for task performance. Several studies suggest that the physical embodiment of a robot has important effects on learning and interaction [5, 6, 1].

In the area of socially assistive robotics for children, [3] deployed an autonomous robot in a two-week long field trial in an elementary school. The authors found that in the unstructured interactions with the robot, children frequently interacted with the robot in groups. In a study of multi-

child, multi-robot interaction, [4] found that disengagement cues for children in group interactions are different than those for children in individual interactions with robots, and that a model trained on one type of interaction does poorly at predicting disengagement in the other type of interaction, suggesting that models intended to be used in multi-party interactions need to be trained on similar interactions, motivating our focus on group rather than individual interactions.

3. METHODOLOGY

We define moderation as a decision-making problem, where the moderator is trying to simultaneously achieve two goals: the goal of the task from the perspective of the participants, and some moderation goal governing *how* the task is to be accomplished. There are two key ideas in this formalization: first, that a moderator's actions can be evaluated relative to these two goals; and second, that the moderator's action set consists only of reassigning physical and social resources, which in turn affect the state of the interaction, as shown in Figure 1.




Figure 1: Approach to Moderating Interactions

This approach can be incorporated into any decision-making algorithm that maps states, actions, and rewards to state-action pairs. However, to ground our notation, we will use a probabilistic approach to defining the elements of the model, an approach that has proven successful in robotics for addressing uncertainties of real-world deployment [7].

As explained above, the moderator is simultaneously trying to achieve two goals, between which there may be some tradeoff. However, because the moderator is constrained to only act on interaction resources, the moderator is actually making decisions in the space of *resource distributions*.

Thus we can define the reward function R as a function of the distribution of resources d :

$$R(d) = G_{mod}(d) + \alpha * G_{task}(d) \quad (1)$$

Where α is the parameter governing the relative importance of the two goals, $G_{mod}(d)$ is the degree to which the moderator goal is achieved for a given distribution of resources, and $G_{task}(d)$ is the degree to which the task goal is achieved for that distribution of resources. Similarly, the transition probabilities in this approach are the probability of success of a moderator behavior m in reassigning a resource, $Pr(d'|d, m)$, or, for the null action, the probabilities that each resource is redistributed without moderator intervention.

4. LEARNING MODEL PARAMETERS

4.1 Annotation and Feature Extraction

We focus on a subset of the corpus that allows us to understand 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, speech and toy-related gestures such as holding the toy, bidding to hand the toy off, or bidding to receive the toy were annotated. Additionally, an informal annotation was performed of types of speech behaviors, such as questions, interruptions, and backchannels.

5. PRELIMINARY RESULTS

The goal for the analysis of these data is 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. A preliminary qualitative analysis of the annotations on speech and toy manipulation behaviors in this dataset reveals the following patterns:

Interaction Goals: The two different topics studied as part of this work provide two different interaction goals. In the story task, the participants are explicitly told that the story should either be told jointly or they should take turns telling a story; this means that ensuring all members of the interaction participate is explicitly defined as a goal of the interaction. We qualitatively observe in the story-telling interactions that the conversational floor is held by individuals for much longer than in the naming interactions, and that the participants are more evenly involved in the interaction. This suggests that participants are self-moderating in order to achieve the interaction goal, and supports the use of this dataset as an example of self-moderation.

Conversational Resources: We study two resources in the interaction: the toy and the conversational floor. The conversational floor often alternates between two people for an extended period of time, suggesting that this resource might be best modeled as something that can be held jointly between two participants, as well as being held by a single participant. With regards to the toy, we observe that the

individual controlling the toy is more likely to hold the conversational floor, providing support for manipulating the toy as a resource in order to change the participation of the individuals in the discussion.

Moderation Behaviors: Finally, we observe a number of self-moderation behaviors, including general questions that release the floor to any other member of the group, and directed questions, that release the floor to a specific other individual. We also observe the toy being passed between participants, which as discussed above, appears to affect the balance of the conversation.

6. CONCLUSIONS

In this abstract we present work towards modeling moderation in order to enable a socially assistive robot to regulate learning-focused group interactions with children. Preliminary results from a dataset with adults provide support for this approach as well as insights into how the model can be grounded in human interaction data.

7. 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.

8. REFERENCES

- [1] W. A. Bainbridge, J. Hart, E. S. Kim, and B. Scassellati. The effect of presence on human-robot interaction. In *17th IEEE Int. Symp. Robot Hum. Interact. Commun. 2008. RO-MAN 2008*, pages 701–706, New York, NY, USA, 2008. IEEE.
- [2] D. Feil-Seifer and M. J. Matarić. Defining socially assistive robotics. In *Proc. 2005 IEEE 9th Int. Conf. Rehabil. Robot.*, volume 2005, pages 465–468, Chicago, IL, June 2005.
- [3] T. Kanda, T. Hirano, D. Eaton, and H. Ishiguro. Interactive Robots as Social Partners and Peer Tutors for Children: A Field Trial. *Human-Computer Interact.*, 19(1):61–84, June 2004.
- [4] I. Leite, M. McCoy, D. Ullman, N. Salomons, and B. Scassellati. Comparing Models of Disengagement in Individual and Group Interactions. In *Proc. Tenth Annu. ACM/IEEE Int. Conf. Human-Robot Interact. - HRI '15*, pages 99–105, New York, New York, USA, Mar. 2015. ACM Press.
- [5] D. Leyzberg, S. Spaulding, M. Toneva, and B. Scassellati. The Physical Presence of a Robot Tutor Increases Cognitive Learning Gains. In *34th Annu. Conf. Cogn. Sci. Soc.*, number 1, pages 1882–1887, Sapporo, Japan, 2012.
- [6] J. R. Movellan, M. Eckhardt, M. Virnes, and a. Rodriguez. Sociable robot improves toddler vocabulary skills. In *Human-Robot Interact. (HRI), 2009 4th ACM/IEEE Int. Conf.*, pages 307–308, NY, NY, USA, Mar. 2009. ACM Press.
- [7] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.