Role-Based Coordinating Communication for Effective Human-Robot Task Collaborations

Aaron St. Clair and Maja Matarić
Interaction Lab
University of Southern California
Los Angeles, California, USA
astclair|mataric@usc.edu

Abstract— **This short summary paper briefly describes a method for using the embodied social communication capabilities of a robot to achieve and enhance coordination in human-robot task collaboration scenarios. The approach focuses on planning coordinating social behaviors using the formalism of roles to allow a robot to produce and interpret communicative feedback expressing a desired allocation of duties and to issue positive or negative reinforcement to a person as the task progresses and in response to the inferred future activity of the collaborating partner.**

***Keywords*—Cooperating Autonomous Robots, Human-Robot Cooperation in Teamwork, Cognitive Modeling/Science in HRI.**

I. Introduction

If service robots in human-centered environments are to support collaborative task performance with people, they must effectively convey intentional behavior and adapt their actions in response to feedback and to the actions and preferences of a human teammate. In co-located environments, the robot should be able to make use of natural social communication modalities such as embodied gesture, sounds, and speech, to achieve coordination with its collaborators, provide help when necessary, improve the performance of the team and subjective evaluations of the robot as a teammate. Prior work has demonstrated that robots that account for the actions of their collaborators when deciding what to do are preferred and perceived as more intelligent [1] and that anticipatory action can play an important role in increasing team fluency [2]. Other work has treated the collaborative process as dialog supporting verbal turn-taking and sub-task assignment [3, 4]. We aim to use social communication (speech and embodied gesture) to efficiently and intelligently allocate roles during a collaboration.

II. Using Roles to Coordinate Group Activity

This work seeks to address the role of social communication and feedback in regulating dynamic task performance without imposing a strict dialog structure. In these contexts, there are a wide range of possible scenarios where team performance could benefit from effective use of communication to convey: information in a partially-observable environment, strategic or tactical misalignment of the team, failure and error conditions, load balancing, and positive feedback or encouragement. Initially, we focus on ensuring that team activities are tactically and strategically aligned through the use of context-dependent role inference and allocation. Although people may play various identifiable roles during a group activity, the notion of *role* itself, as well as the process by which a role is assigned or assumed, is a human construct that typically lacks formal definition [8]. Our goal for the human-robot team is to minimize conflict and improve team fluency by using communication to convey, reinforce, or modify the future activity of members of the group. For the purposes of this work, we assume a one-on-one, co-located collaboration between a person and a robot. We also assume that both know the steps required to achieve the collaborative goal and are motivated to work together to achieve the goal as a team. In addition to performing the task, the robot must be able to evaluate the status of the collaboration and decide when and what kind of communication can best be used to improve task performance.

III. Approach

In order to generalize the communication system across different task domains, we must devise a method for factoring the details of the particular activity away from the *role allocation* process. To accomplish this, the robot must have task-independent mechanisms allowing it to: 1) infer the role of the person collaborating with it from multimodal sensor input; 2) evaluate the alignment of its own role with the inferred role of its collaborator; and 3) produce and interpret embodied communication concerned with allocation or modification of roles. Given that agents are goal-directed and have prior knowledge of the activity, we assume we are given a Markov decision process (MDP) formulation of the problem comprising a tuple $M = (S, A, T, R)$ where S is a set of mutually exclusive states encoding the configuration of agents and objects in the world, A is the set of actions the agent can take, and T is the transition function. These types of models have been widely used in robotics for planning, multi-agent coordination, and in human-robot dialog systems [5], and are also frequently used in robot learning from demonstration [6], where the model is learned from some outside source such as a human. Following from the previous assumption that the robot knows how to do the task, we also assume we are given a set of policies, $\Pi(s) = a$, mapping states to actions, that can be used to complete the task but are not necessarily optimal.

A. Defining Roles

To generate sensible role-based coordinating communication from the robot to the person, we first require some information about the person's mental model of the task in terms of how their current actions fit into the team's strategy for achieving the overarching task goals. People are remarkably good at ascribing intent to observed activities, even with minimal contextual cues [7]. Due to the complexity of intentional activity and the myriad factors beyond the constraints of the task that could potentially affect a person's decision-making during the collaboration, it is infeasible to model intention directly. Instead, we simplify this process by computing the likelihood of the person executing each of the provided policies given an observed sequence of their activity. This can be done via an exponential moving average, where more recent observations are given higher weight, allowing our estimate to change in response to changes in the user's strategy. We will use this policy information as our notion of the user's role at a given moment in the task, and will compare it with the robot's own selected policy.

B. Task Performance Evaluation

Once we have an estimate of the user's policy, we must evaluate the projected quality of the collaboration if both agents continue to use the same policies (keep the same roles) and decide whether a change would be worthwhile. Given the policy of each agent, we can compute the expected reward for the person by evaluating the value of each policy weighted by its likelihood from the previous step. As this is a human-robot context we are not necessarily concerned with finding an optimal joint allocation but rather with avoiding and recovering from serious mistakes. To control the overall amount of communication, we must determine threshold reward and certainty values below which communication is unnecessary.

C. Communication Production

After the robot has decided to communicate a desired role allocation in order to improve task performance, it must efficiently convey this desired allocation to its collaborator. Since people working in teams benefit from shared visual information [9], it seems desirable for the robot to use a variety of natural embodied communication modalities such as gaze, deixis, and emblematic or iconic gestures, in addition to speech. These modalities can be categorized as follows:

1) Task-Dependent

Modalities such as illustrating gestures specific to an action, like a hammering motion for instance, are tightly-coupled to the task. To make effective use of task-dependent modalities, the robot must know or learn how specific gestural cues refer to specific roles in the task. This can be accomplished by observing human gesture use during the task, or via a training process in which the robot learns resulting interpretations of its gestures in terms of roles or actions.

2) Task-Independent

Other communication modalities such as attention-focusing cues like gaze and deixis are mainly dependent on the spatial relationships between the collaborators and objects involved in the task. Using these modalities to communicate a desired role allocation requires the robot to know a spatial location in the workspace that evokes the desired role (for example, pointing to a work object like a hammer to indicate a nailing task) and to be able to indicate it accurately. While faster and easier to produce, these modalities may be more ambiguous, especially in cluttered environments, therefore placing the interpretative burden on the collaborator.

IV. Task Categorization and Evaluation

Our goal for this approach is not just to model the communication dynamics of a single, specific task but rather to develop a model that can generalize to other task settings as well. Toward this end we plan on evaluating the methodology in multiple different task environments and with different people drawn from diverse populations. In this section, we will 1) detail a set of factors that we anticipate will impact coordinating communication usage and 2) clarify the set of applicable tasks to which the approach can be applied.

Based on existing work studying human-human and human-robot collaborations we have identified a number of factors that could potentially affect the dynamics of coordinating communication. These can be separated into two overarching categories: characteristics of the user and properties of the task itself. User characteristics that may affect communication tendencies include the following:

- Personality, a determining factor in social behavior, with effects on communication preferences [10].
- Hierarchy or status of the user with respect to collaborator, will impact the willingness of the user to issue and accept role assignments. Capability, the user's ability to perform the task adequately will affect their choice of activities and may induce them to ask for help or defer activities.

The structure of the task itself also has an effect on the efficacy of social communication for coordination. As stated previously, we are limiting our investigation to pairwise, physically-decoupled tasks, that is tasks where the coordination occurs through social communication rather than through the feedback associated with jointly manipulating an object. We are also making the assumption that the robot can capably perform at least some of the task, this assumption limits the scope of the task to activities that a robot could conceivably do autonomously such as simple object manipulation and navigation. In addition to these distinctions we consider Steiner's taxonomy of group tasks [11] to further clarify the types of tasks we anticipate will be amenable to this modeling approach. Steiner defines three categorical dimensions across which tasks are grouped that describe: divisibility of the task, type of goal, and interdependence of each individual's inputs in final output.

1) Divisibility

By relying on role assignment as the means of coordination we implicitly assume that the task is divisible i.e., can be effectively separated into subtasks that each agent can perform separately in contribution to the greater task goals. Unitary tasks, as defined by Steiner, are an interesting special case in which the task is not divisible resulting in one teammate becoming a bystander. In this case, either the robot or the person could not directly perform the task and would become idle. While the approach could support an idle robot by setting the robot's task action set to be empty, the resulting communication generated would only be comprised of recommended actions for the person. This scenario in which the robot monitors an activity and offers feedback is an interesting future direction with applicability in rehabilitation and other domains, but we will not consider it specifically here.

2) Goal types

Two goal types are considered in Steiner's framework optimizing goals, with subjective ratings of the achieved result and maximizing goals, which typically have an objective evaluation metric built-in. From the robot's point of view this distinction is not very useful since ultimately the robot control system requires some quantitative metric with which it can select optimal behaviors in expectation. We will mostly be considering the sorts of assembly tasks that would fit in the maximizing goal category.

3) Interdependence

This area of Steiner's categorization is concerned with how the input from various members is combined to form a single output and how group performance can be related to the best and worst individual performances. We anticipate the addition of the robot to improve performance as compared to either the robot or person alone since the types of tasks the robot can perform are typically divisible into separate subtasks with output combined in an additive manner. Thus we anticipate adding more agents to increase the performance of the group.

V. Experimental Setup & Design

In order to demonstrate that this modeling approach is capable of generalization to many tasks we plan on evaluating its performance in multiple different task environments. To provide a controlled, repeatable environment where we can conduct experiments with people and a robot we have developed an augmented reality task environment. The experimental space consists of a 6x6 meter smart room environment that has been outfitted with sensors and overhead projection capabilities (Figure 1). The calibrated projection system uses a set of ceiling-mounted projectors to display a set of virtual objects or agents on the floor of the experimental space. Each view is warped by a linear transformation to ensure that overlapping areas are correctly registered, minimizing shadows. People and robots in the room are tracked in real-time using a set of Microsoft Kinects. A task simulation system updates the positions and states of the virtual objects and agents in response to sensor information and optional localization data from the robot. This setup allows for quickly prototyping tasks, changing parameters, and running multiple trials of a given task under identical conditions.

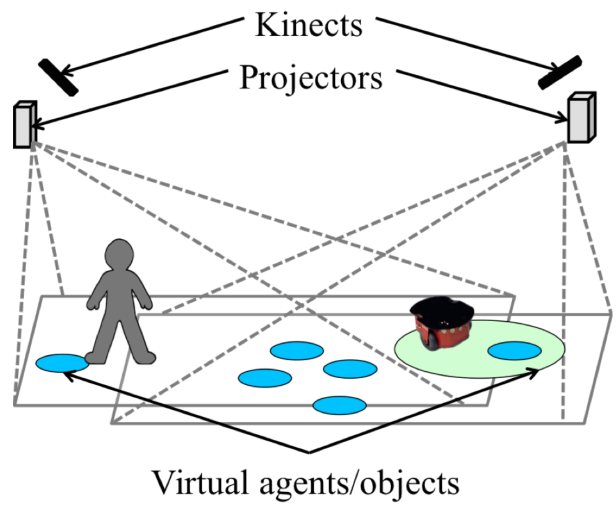

Figure 1. A diagram of the experimental setup. Calibrated overhead projectors overlap to minimize shadowing. The task simulation uses person and robot tracking data from multiple Microsoft Kinects mounted near the ceiling.

Prior work has established the use of these types of systematically abstracted environments, called synthetic tasks, for use in studying human factors [12]. In this case we are using the virtual environment to abstract away physical manipulation of objects, which will instead be handled using positional information, remote buttons, and gesture input collected by a handheld WiiMote. This preserves the properties of physical manipulation for physically-decoupled tasks as items may only be manipulated by one agent at a time and is not outside the realm of robot capability.

The first task that we have implemented is a pseudo-herding task, in which the human-robot team must collect a set of wandering sheep and return them to a pen area. Unlike real herding, which has been previously studied in detail and requires the agents to maintain spatial relationships, the simulated sheep attach themselves to the person or robot when caught. This allows the task to be divisible, and necessitates some coordination regarding who gets which sheep, especially as the number of sheep dwindles. As an additional constraint we have implemented a timed lock that must be switched on by one of the agents at a regular interval to prevent the sheep from escaping the pen, and a light switch, which turns off periodically preventing agents from seeing the sheep. The interdependence of the agents in this task as per Steiner's framework could be categorized as additive since there are a fixed number of sheep that must be herded and each agent's individual actions contribute equally and directly to team performance. The lock and light introduces two periodic task constraints upon which the agents must coordinate their activity to prevent negative results.

VI. Ongoing and Future Work

Using the herding task with twelve sheep as an initial evaluation environment and a Pioneer 2AT as the robot collaborator, we have begun a series of experiments to develop and evaluate the role-based communication system. First, we must establish a performance benchmark for teams of different composition including the robot alone, a person alone, and two-person teams. Collected data from n=6 people has established mean completion times for each group (Table 1). The robot is clearly the slowest performer, due to the lock and light locations and sheep speeds, which are tuned to make the task interesting for people and results in the robot having difficulty keeping the lock triggered. We have observed that all the two-person teams tested thus far have experienced at least one (and sometimes more) instance of letting the timed lock lapse and thus losing collected sheep and undoing all task progress. More investigation is necessary to determine if this is due to a lack of team rapport or lack of familiarity with the task and whether the same will hold in the robot condition.

TABLE I
MEAN COMPLETION TIMES ON SHEEP HERDING TASK

Team Type	Completion Time (minutes)
Robot Alone	10:00+
Person Alone	2:07
Person-Person	1:58

Currently, we are using the person-person data to select the types of roles supported by the robot, to develop recognition systems for estimating the role the person is performing, and to specify the communication actions (verbal phrases and deictic gesture) that the robot will use to refer to the selected roles. This will be accomplished by collecting third party annotations of the collected task performance sequences. The finished communication system will be evaluated in a larger experimental validation, which will feature repeated trials of the task to mitigate the effects of practice. In future work, we are interested in investigating methods for adapting and personalizing communication tendencies in response to user characteristics such as personality. We are also investigating the effects of time-induced stress on people's use of communication, and developing methods to recognize when the user is feeling stressed and allowing the robot to adapt as needed to the situation and communicate effectively.

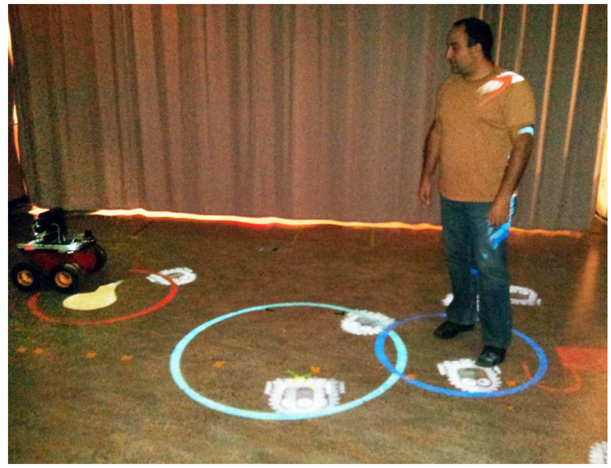

Figure 2. A person and a Pioneer 2AT collaborating on the herding task. The light bulb, pen, and lock are used to play the game. Circles are drawn around the robot and the person (red and blue, respectively).

Acknowledgment

The research reported here was conducted at the Interaction Lab, part of the Center for Robotics and Embedded Systems (CRES) at the University of Southern California. This work was supported in part by National Science Foundation (NSF) grants CNS-0709296, IIS-0803565, and NSF IIS-1117279 and ONR MURI grant N00014-09-1-1031.

References

[1] J. Shah, J. Wiken, B. Williams, and C. Breazeal, "Improved human-robot team performance using Chaski, a human-inspired plan execution system," in *International Conference on Human-Robot Interaction (HRI)*, pp. 29–36, ACM, 2011.

[2] G. Hoffman and C. Breazeal, "Cost-based anticipatory action selection for human–robot fluency," *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 952–961, 2007.

[3] C. Breazeal, C. Kidd, A. Thomaz, G. Hoffman, and M. Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *Intelligent Robots and Systems, 2005. (IROS'05)*, pp. 708–713, 2005.

[4] J. Trafton, N. Cassimatis, M. Bugajska, D. Brock, F. Mintz, and A. Schultz, "Enabling effective human–robot interaction using perspective-taking in robots," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 35, no. 4, pp. 460–470, 2005.

[5] T. Kollar, S. Tellex, D. Roy, and N. Roy, "Toward understanding natural language directions," in *International Conference on Human-Robot Interaction (HRI),* pp. 259–266, IEEE, 2010.

[6] B. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009.

[7] F. Heider and M. A. Simmel. An experimental study of apparent behavior. *American Journal of Psychology*, 57:243–249, 1944.

[8] R. Smith, R. Hixon, and B. Horan, "Supporting flexible roles in a shared space," in Proceedings of the 1998 *ACM conference on Computer Supported Cooperative Work (CSCW)*, pp. 197–206, ACM, 1998.

[9] R. Kraut, S. Fussell, and J. Siegel, "Visual information as a conversational resource in collaborative physical tasks," Human-computer interaction, vol. 18, no. 1, pp. 13–49, 2003.

[10] C. Nass and K. M. Lee, "Does computer-synthesized speech manifest personality? Experimental tests of recognition, similarity-attraction, and consistency-attraction.," *Journal of Experimental Psychology: Applied*, vol. 7, no. 3, p. 171, 2001.

[11] I. Steiner, Group Process and Productivity. New York: Academic Press, 1972.

[12] E. Martin, D. R. Lyon, and B. T. Schreiber, "Designing synthetic tasks for human factors research: An application to uninhabited air vehicles," In *Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting*, pp. 123-127. Santa Monica, CA, 1998.