

# How Robot Verbal Feedback Can Improve Team Performance in Human-Robot Task Collaborations

Aaron St. Clair  
Interaction Lab  
University of Southern California  
Los Angeles, CA 90089  
astclair@usc.edu

Maja Matarić  
Interaction Lab  
University of Southern California  
Los Angeles, CA 90089  
mataric@usc.edu

## ABSTRACT

We detail an approach to planning effective verbal feedback during pairwise human-robot task collaboration. The approach is motivated by social science literature as well as existing work in robotics and is applicable to a variety of task scenarios. It consists of a dynamic, synthetic task implemented in an augmented reality environment. The result is combined robot task control and speech production, allowing the robot to actively participate and communicate with its teammate. A user study was conducted to experimentally validate the efficacy of the approach on a task in which a single user collaborates with an autonomous robot. The results demonstrate that the approach is capable of improving both objective measures of team performance and the user's subjective evaluation of both the task and the robot as a teammate.

## Categories and Subject Descriptors

H.1.2 [Models and Principles]: User/Machine Systems—*human factors, software psychology*; H.5.2 [Information Interfaces and Presentation]: User Interfaces—*evaluation/methodology, natural language*; I.2.9 [Artificial Intelligence]: Robotics—*operator interfaces*

## Keywords

Human-robot collaborations; natural language

## 1. INTRODUCTION

Robots are becoming increasingly competent at performing tasks such as navigation and manipulation in real-world human environments including homes, offices, and other settings where they are around people. This expanded range of capabilities makes possible many scenarios in which robots not only take action around people but also collaborate with them to accomplish a task together. In a home environment a robot might help a person clean up a messy room, in the office a robot could help to rearrange furniture, and in a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
HRI'15, March 2–5, 2015, Portland, Oregon, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.  
ACM 978-1-4503-2883-8/15/03 ...\$15.00.  
<http://dx.doi.org/10.1145/2696454.2696491>.

manufacturing setting a robot might help assemble a product. These collaborative scenarios present a number of challenges; not only must the robot avoid physical collision with the human, it must also be a productive part of the collaboration accounting for the human's activity when deciding what to do next, and providing a means for the human to understand and control its behavior.

In this paper we consider the problem of coordinating social communication during the course of a pairwise human-robot task collaboration. Effective feedback is important in these scenarios since it has a direct impact on the situational awareness of the robot's collaborator. We present a general methodology allowing a robot to produce social communication (currently speech), to efficiently and intelligently allocate roles during a task collaboration. The primary goal of the method is to support the user's natural collaboration modalities and preferences and ensure applicability across many typical human-robot task collaboration scenarios. The approach generates three types of situated verbalizations aimed at providing useful information to a teammate in co-located joint activities. To make the approach generalize across many task settings, we base it on existing task formulations from multi-robot systems. Unlike existing work in human-robot collaboration, it does not rely on a specific conversational structure such as turn-taking, allowing for it to be used in dynamic tasks. We hypothesize that providing feedback in the form of coordinating social communication will lead to improved team performance and subjective evaluation of the task itself by users and also allow the robot to provide guidance for the user's behavior.

The rest of the paper is organized as follows. First, we provide some background information on human-human and human-robot collaboration and related existing work. Next, we provide an overview of the communication planning approach and detail the three types of verbalizations used by the robot. Next, the details of an initial specific task to which we have applied the communication planning approach is presented. Finally, an experimental validation of the approach in this task setting is presented as an initial proof of concept, demonstrating improved performance and subjective evaluation when people work with the communicating robot.

## 2. BACKGROUND

We present a brief review of relevant literature from the social sciences about human collaboration, as well as from human-computer and human-robot interaction as relates to this work.


People have a natural ability to collaborate comprised of a complex series of behaviors that attributes agency to the people and things around them. We recognize our collaborator’s actions using visual observation via the mirror neuron system [16] and theorize about the mental models of others including their beliefs, desires, and intentions [12]. Finally, we issue a variety of implicit and explicit social cues aimed at helping others understand their own actions [10, 2]. Many of these processes, with the exception perhaps of strategic planning, occur automatically without conscious thought on our part and are active even when the person is collaborating with a robot. Additionally, people rely heavily on speech for coordination. Robots that collaborate with people should take advantage of the capabilities of their human counterparts by providing appropriate signaling to improve their partner’s efficiency. By enabling a robot to produce human-like communication during a task, this work aims to support effective coordination between people and robots without requiring users to be trained how to use the robot beforehand.

Relevant prior research on human-machine collaboration includes an extensive body of work on top-down deliberative approaches for modeling and constructing shared plans for collaboration [6, 18]. There has also been extensive work on intent recognition relying on perspective taking or Theory of Mind-inspired models to allow a robot to recognize a person’s intentional behavior through observation [4], as well as to recognize and generate intentional face-to-face meeting initiation behavior [9], and to recognize helping and hindering social behavior via an MDP formalization utilizing inverse planning [21]. Other relevant work demonstrated robot learning of simple tasks through human tutelage and effective collaboration via turn-taking or biased pre-emptive action [1, 8]. Task scheduling systems have been developed to change the task planning of a robot teammate in response to actions taken by a human teammate [17, 22]. Another area of research has focused on identifying and evaluating promising coordination behaviors, particularly eye gaze cues [14], implicit and explicit verbal communication [20], and legible path planning [5]. There is also relevant work in learning demonstration on methods for teaching and instructing robots including collaborative settings [11, 3].

This 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 [17] and that anticipatory action can play an important role in increasing team fluency [8]. Other work has treated the collaborative process as a dialog, supporting verbal turn-taking and sub-task assignment [1]. Speech has been demonstrated to be an effective input modality for commanding a robot [17, 11].

### 3. APPROACH

Our task control and communication approach consists of three components: 1) the task control system, 2) the human activity model and recognizer, and 3) the communication planner and executive. Planning with a human partner in the environment is distinct from planning in multi-robot scenarios mainly in how communication takes place between teammates. Similarly, there are existing approaches to segmenting, recognizing, and modeling human activity in various contexts. We have developed a simplified methodology to serve these purposes in our experimental task, although



**Figure 1: System diagram of the collaborative communication approach (inset) and associated inputs and outputs.**

this work could be integrated with and perhaps improved by state-of-the-art recognition and task planning systems. The main contribution of this paper is in the area of planning and execution of communication actions during collaborative task execution in HRI.

#### 3.1 Task Formulation

We formulate the coordination problem in two parts, as follows. First, the robot represents the task as an MDP to plan its actions in the presence of noise. We define a Markov decision process (MDP) for the task  $M = S, A, T, R$ , where  $S$  is the finite state of the environment,  $A$  is the set of task actions the robot can execute,  $T$  is a function giving a probability distribution over states for executing a given action in a given state, and  $R$  is a reward for each state. We assume that the robot can perform the task and has a policy  $\pi(s) = a$  that allows it to perform the task. In this case, the transition function captures the robot’s uncertainty due to environmental sources, such as sensor or motor noise, and also due to changes the human collaborator might make in the environment. This formulation has been successfully used previously in human-robot collaboration [15] to coordinate the task-actions of the robot without any communicative feedback. In the future, we would like to investigate the use of multiagent policies with joint action mapping for both agents to further improve the robot’s feedback.

#### 3.2 Role-Based Activity Model

In addition to ensuring that the robot can jointly perform the task in the presence of a human partner, we want the robot to provide verbal feedback to its partner. We focus on communicating intended action via the mechanism of *roles*. People use explicit and implicit roles in team activity and other organized behavior. For example, if a group of people is tasked with assembling a wooden box, they might assign a sawing role, a hammering role, and a painting role that different people assume during the course of the activity. Interestingly, we could find no well-defined, formal definition of roles or how to partition a given activity into a set of roles [19]. In our pilot experiments we determined that people tended to assign others roles associated with common work objects and locations, although some of these consisted of multiple discrete activities.

In the evaluation task, described later, the act of herding a sheep (finding a stray sheep, collecting it, and bringing it to the pen) was commonly grouped into a single activity or role (“go herd sheep”). To capture these properties we model roles using a set of assignable policies,  $\Pi = \pi_1, \pi_2, \dots$ . This set of policies is domain-dependent and not assumed to be optimal or otherwise sufficient, on its own, for solving the task. Rather, the policies in the set are a means of quantifying each agent’s likely behavior over the course of the task, and of grouping similar actions according to the roles people typically assign in that type of task. If the robot and the person have differing capabilities, the set of possible policies for each may be different. For example, if the robot cannot perform all of the actions the person can, the robot activity model may consist of a subset of  $\Pi$ . In the following explanation we will assume that the robot and person can both perform all of the task actions and thus that the robot policy set,  $\Pi_{robot}$  is identical to that of the person,  $\Pi_{user}$ . We maintain a distribution,  $D$ , over  $\Pi$ , based on the likelihood that the user is executing each policy and updated by observing state transitions over time, inferring user actions, and reweighting policies that agree.

To accurately infer the role,  $r \in \Pi$ , that best describes the user’s action selection preferences, we assume we have access to some means of action recognition. Specifically, we want to know when an agent (person or robot) takes an action  $a \in A$  at time  $t$ . Human action segmentation and recognition is a challenging open problem with relevant existing work from computer vision [7]. We have written a heuristic action recognition system that makes use of a factored representation of the state of the task as a set of discrete random variables,  $S = \{X_1, X_2, X_3, \dots, X_n\}$ . In order to execute its policy in a fully observable environment, the robot must take sensor input and determine the value of each random variable so that it can determine the current state. Given this capability, we observe that some of the  $X_i$  can be directly associated with actions that the agents can take, while others are consequences of the environment. For example, if a switch must be pressed during the course of a task, then we might have an action  $a \in A_{task}$  called *FlipSwitch* and an associated random variable  $X_{switch} = \langle Off, On \rangle$ , describing the state of the switch. In our heuristic action recognition system, we determine the subset of  $X_{recognizable} \subseteq \{X_1, X_2, X_3, \dots, X_n\}$  that are associated with state changes that are attributable only to the person or the robot. On each update, if the current inferred value of one of these variables is not equal to its previous value,  $X_i^{(t-1)} \neq X_i^t$ , we assume that one of the agents made the change. To determine which agent, we select the agent closest to the object at that point in time.

This heuristic method works well in simple scenarios with straightforward mechanics, when the sets of random variables are carefully selected, but does not transfer well to all spatiotemporal actions that a person and robot might undertake. The overall system could benefit from integration with a more principled state of the art action recognition method; this is beyond the scope of this paper.

Given the stream of recognized actions, the system next attempts to recognize the role the user is employing at the current time. This process is detailed in Algorithm 1. We first discard all recognized actions as being taken by the robot, leaving only the person’s actions. For each  $a_{rec} \in A_{task}$ , we perform a Bayesian update of a multinomial distribution over the set of role policies,  $p(R = r_i | A = a_{rec}) =$

$\alpha * p(A = a_{rec} | R = r_i) * p(R = r_i)$ . The likelihood of the action being taken given a certain role  $r_i$  is selected based on whether the policy would have executed the action in the previous state, with a specified weight, as follows:

$$p(A = a_{rec} | R = r_i) = \begin{cases} 0.95, & \text{if } r_i(s_{prev}) \equiv a_{rec} \\ 0.05, & \text{otherwise} \end{cases}$$

After normalization, the inferred user role is selected by  $\pi_{user} = \text{argmax}_r p(r | a_{rec})$ . This estimate of the user’s role is thus based on their action selections at previous times during the task and is dependent on individual actions taking around the same amount of time. If this is not the case, the likelihood of the action given the role could be expanded to include more than one past state, with an appropriate discount factor. In our validation, the system maintains a multinomial distribution over the set of user roles based on the likelihood that the user is executing each policy. On each update, policies are reweighted based on their agreement or disagreement with the recognized action.

---

#### Algorithm 1 User role recognition algorithm

---

**Require:** Recognized human action  $a_{rec} \in A_{task}$ , previous state  $s_{prev} \in S_{task}$

```

for all roles  $r = \{\pi_1 \dots \pi_n\}$  do
  if  $r(s_{prev}) = a_{rec}$  then
     $p(r | a_{rec}) = 0.95$ 
  else
     $p(r | a_{rec}) = 0.05$ 
  end if
   $p_r \leftarrow p_r * p(r | a_{rec})$ 
end for
return  $\text{arg max}_r p(r)$ 

```

---

## 4. PLANNING ROBOT FEEDBACK

Based on this recognized action and the robot’s own policy, three types of verbal feedback are generated: *self-narrative*, *role-allocative*, and *empathetic*. A wide range of verbal communication could be produced to improve team coordination. These feedback types were selected as they provide a balance of information to the person about what the robot is doing and what the robot expects them to do. We further validated this selection by conducting a pilot study in which pairs of people performed a task. Based on their transcribed verbalizations, these three types of communication were widely used across different teams and are well-supported by our task model. Additional types of communication, such as providing key bits of salient information, which could occur at varying levels of granularity, as well as qualitative commentary about aspects of the task environment were also used by people but are not supported by the communication model.

### 4.1 Self-Narrative Feedback

The simplest form of feedback to produce is *self-narrative feedback*, where the robot tells the user what it is going to do, e.g., “I’ll take care of X”. We define a communicative intent for each policy available to the robot,  $C_{narrative} = \text{RobotDoing}(r_i), \forall r_i \in \Pi$ , and select a single active communicative intent based on the role the robot is executing. To convey these communicative intents to the person, the

robot has two options. It could say a phrase indicating its role assignment, i.e., telling a user what role it will execute during the task, or it could implicitly communicate its role assignment by providing feedback about which actions it is executing at a given point in time, leaving the person to infer the robot’s role (policy) from this sequence of action selections. The former case requires a mapping of roles to a set of verbal feedback phrases,  $V$ , that refer to the correct role,  $f_{narrate-role} : \Pi_{robot} \rightarrow V$ , while the latter case requires a mapping of every action to one or more verbal feedback phrases  $f_{narrate-actions} : A_{task} \rightarrow V$ . Ideally,  $V$  will contain multiple options for each policy or action so that the robot does not use the same redundant phrasing, which could lead to it being perceived as less intelligent, an issue to be avoided in HRI.

## 4.2 Role-allocative Feedback

To generate *role-allocative feedback* for the user, such as “you handle X”, the expected next action of the user given by our recognized policy  $\pi_{user}$  is compared to the robot’s next action as given by the robot’s policy,  $\pi_{robot}$ . This comparison is dependent on the type of model employed. If we have access to only a single-agent model of the task, like a typical MDP, we have no specific information about the value of joint action allocation about the best allocation for each agent. Therefore, we have developed a role-allocation method using a single-agent task that specifies a user’s role based on the principle of avoiding conflicting action. We assume a list of discrete actions in  $A_{task}$  that would conflict if executed by the robot and person at the same time. Given the policies of each agent, we check to see if they are expected to execute the same action given the current state. If this action is in the set of conflicting actions, we suggest an alternate role from  $\Pi_{roles}$  for the user to do that is not conflicting, as detailed in Algorithm 2.

This results in a communicative intent that assigns the user to a given role,  $C_{role-allocative} = PersonRole\langle r_i \rangle, \forall r_i \in \Pi_{roles}$ . The phrases for role-allocative feedback have similar requirements to those in the narrative feedback case. Specifically, we require a function mapping roles (or alternatively actions) to a set of verbal feedback phrases that the robot can say  $f_{role-allocative} : \Pi \rightarrow V$ . It should be noted that, at this step, the robot could proactively perform the better action instead of recommending that the user do it. For our initial experiments, we do not consider this option, but assume the robot’s policy is stationary.

## 4.3 Empathetic Feedback

Finally, to generate *empathetic feedback*, we monitor for specific state transitions in the task model. When a transition occurs that has been marked as requiring empathetic feedback, we generate an appropriate communicative intent in the form of either a positive or negative expression (e.g., “Oh no” or “Great”). Our current approach maps specific transitions to communicative intents, e.g.,  $f(s, s_{prev}) = c \in C_{empathetic}$ . By mapping specific state transitions to instance of empathetic feedback, these may also include more specific expressions of empathy, such as a description of what exactly went right or wrong. A general approach to providing empathetic feedback consists of monitoring the value of states over time and responding to large changes in value with a threshold for value changes that matches people’s policies for producing empathetic feedback. When a

---

### Algorithm 2 Alternative role-allocation suggestions

---

**Require:** Person and robot policies  $\pi_{user}, \pi_{robot}$ , current state  $s \in S_{task}$ ,  $ConflictingAction(a_{user}, a_{robot})$

```

 $a_{user} \leftarrow \pi_{user}(s)$ 
 $a_{robot} \leftarrow \pi_{robot}(s)$ 
if  $a_{user} = a_{robot}$  and  $ConflictingAction(a_{robot}, a_{user})$ 
then
  for all roles  $r = \{\pi_1 \dots \pi_n\}$  do
    alternatives  $\leftarrow \{\}$ 
    if  $(r(s) \neq a_{robot})$  or
      not  $ConflictingAction(r(s), a_{robot})$  then
        append  $r$  to alternatives
    end if
  end for
  return alternatives
else
  return  $\{\}$ 
end if

```

---

large value change is detected, the robot can offer the generalized positive or negative empathy response as described above. In this approach, the threshold specified is domain-dependent, although it could be estimated by observing people’s use of empathetic verbal cues during human-human or human-robot task performance. Additionally, more specific feedback may lead people to ascribe greater intelligence to the robot as it identifies why something went wrong, for instance leading people to conclude the robot has a higher-level understanding of the task than it does.

## 4.4 Communication Executive

To avoid repetitive speech, appropriate phrases for each type of verbal feedback were determined by having people perform the task in a small pilot experiment. For each action, we stored a set of phrases and randomly selected one phrase when executing the communication. We used baseline phrases as collected from the pilot study, with no attempt to make them more or less polite. The phrases, stored as text strings, are passed to a text-to-speech engine and played in real-time during the task as triggered by the communication planning system. A simple speech executive node was developed to allow multiplexing the various types of communicative feedback without interrupting individual phrases. In future work, we will investigate adapting the types of feedback and phrases spoken in response adherence and performance metrics, allowing for more personalized feedback for individual users.

## 5. EVALUATION

To test a system for planning coordinating communication in the context of human-robot task scenarios, the first requirement is a robot and accompanying low-level control systems that can reliably perform tasks in an environment co-located with a person. Manipulation of real objects is a challenging problem, and existing off-the-shelf software for performing manipulation tasks is typically slow and requires tuning for specific environmental factors. Additionally, since the robot must monitor its teammate’s activity, doing a task that involves real object manipulation also presents difficulties in perceiving the person. To overcome these challenges, we developed a task simulation engine and implemented a task to be performed by two-agent teams.

The task simulator models the behavior of multiple virtual agents and objects over time. The agents are projected onto the floor via an overhead projection system. The system tracks people using environmental sensing and localizes the robot using on-board sensing in the shared space. By updating the simulated positions of agents and objects, robots and people in the room are able to interact with the virtual world by navigating around the room and positioning themselves on top of the virtual objects. The virtual objects are projected on the floor of the room with multiple calibrated projectors and people are tracked with Microsoft Kinects RGB-D sensors using point cloud processing and an unscented Kalman filter. Due to the size of the room and difficulty of fusing the output of multiple skeletal pose estimates, the robot is assumed to only have the  $(x, y)$  position of the person in the room. The robot used in this scenario is the Pioneer 2AT (Figure 3); it was selected due to its effective navigation of the environment using an on-board Hokuyo laser rangefinder for localization as well as its small size, which reduces the possibility of harm in an accidental collision.


### 5.0.1 Synthetic Collaborative Task

This experimental setup has several advantages for conducting tasks compared with using physical objects and/or confederate experimenters. First, it allows for many repeatable, dynamic agents that move autonomously in a directed or pseudo-random manner. It also allows for calibrating the velocity, shape, and behavior dynamics of the simulated agents to make a given task easier or more difficult as desired. Finally, since it does not depend on physical objects, several different tasks can be conducted quickly by switching the task controller. The simulator abstracts away physical tasks that are not the focus of this research, such as object manipulation, allowing us to focus on collaborative behavior. This is consistent with the notion of research tasks described by [13] as a systematic abstraction of a real-world task. Despite this abstraction, we are able to preserve some of the complexity of real-world environments, such as partial observability, noise, and occlusion by augmenting or filtering the robot’s sensor data. This represents a reasonable trade-off of environmental realism for repeatable HRI experiments, tunable task difficulty, and expanded task possibilities. In future work we plan on applying the approach in a real-world environment outside of the task simulator.

## 5.1 Pseudo-herding Task

The task scenario implemented for evaluation in the augmented reality environment was a pseudo-herding task featuring a group of sheep that appear from the boundary of the room over time with a specified arrival rate. The sheep move about exhibiting Brownian motion with simple avoidance rules and must be collected and brought to a central holding pen. Only one sheep can be collected by an agent at a time, hence it is pseudo-herding. To collect a sheep, the agent must move to a position overlapping the sheep. People are shown their tracked location by projecting a ring around the last tracked position. After a sheep is collected, it follows the person until it is dragged inside the circular pen, at which point it stays inside.

In addition to the pen, there are two timed game-play elements: a lock and a light. When the lock timer reaches zero, the pen becomes unlocked and any captured sheep es-



**Figure 2: The experimental setup featuring a person and the Pioneer collaborating on the pseudo-herding task.**

cape and begin to wander again, effectively reverting any progress made. When the light timer reaches zero, the visibility of the projection is modified to simulate lights going out in the room. This prevents users from easily finding wandering sheep. The timer elements are rendered with a progress bar and color change to clearly indicate when they are about to expire. The task has a number of interesting properties, such as dynamic moving objects; it is divisible and loosely-coupled, supporting different completion strategies. The goal of the game, as introduced to participants, is to collect all the sheep into the pen area as quickly as possible, requiring maintenance of the lock and light elements.

## 5.2 Application of Verbal Feedback Approach

Applying the approach to the pseudo-herding task required the development of the task model, robot policies, person policies, and verbal phrases specific to the herding task. The state is comprised of a set of independent random variables  $S = \{X_{lock}, X_{light}, X_{has-sheep}\}$ . For the herding task we define a relatively simple state space consisting of three features  $X_{lock} = \{high, low, unlocked\}$ ,  $X_{light} = \{high, low, off\}$ , and  $X_{has-sheep} = \{true, false\}$ . Because navigation is used as a proxy for manipulation in the augmented reality task environment, we define a navigation action for the robot to activate each of the work items (lock, light, and pen). Since all the sheep are equally valuable, we define a single navigation action for sheep retrieval in which the robot collects the nearest sheep. Thus the possible actions are *GotoLock*, *GotoLight*, *GotoPen*, *CollectSheep*. The robot’s task behavior during the task consists of a set of policies mapping states to action  $\Pi : s \leftarrow a$ . Note that in the augmented reality environment the robot and person have the same capability and thus share the same action set,  $A = A_{robot} = A_{user}$ .

To apply the proposed communication approach, we first seek to define a set of policies that covers the set of roles that a team of people might define to complete the task. Fitting with our assumption that roles are correlated with distinct


work objects we have defined roles associated with each work object as well as their possible pairwise combinations:  $\Pi = (\pi_{lock}, \pi_{light}, \pi_{sheep}, \pi_{lock-light}, \pi_{lock-sheep}, \pi_{light-sheep})$ . In single object policies, the action in every state is the navigation action associated with the object, i.e.,  $\forall s \in S, \pi_{lock}(s) = GotoLock$ . For policies with multiple objects, action preferences are given by the following object priorities  $lock > light > sheep$ , thus the  $\pi_{light-sheep}$  policy will make the robot go to the light if it has turned off, otherwise it will collect sheep and take them back to the pen. No three-object roles were defined since those would consist of a single agent trying to do all parts of the task. This set of policies is used to both control the robot and to model the person’s behavior as they have equal capability in this case.

The mapping of roles to policies, as well as the phrases used by the robot, were informed by a pilot experiment in which two people performed the task together, where we found people roughly partitioning the two timed elements and sharing the duty of collecting sheep. For the robot’s speech, selected transcribed phrases and used Google Speech API for the robot’s voice. Since there are three possible speech channels, a multiplexing and queuing system was implemented to prevent a type of verbal feedback from interrupting the one already playing. This was sufficient for realizing the verbal feedback as each type of phrase is short in duration and only generated periodically when the task state or user model changes. This system selected randomly among all possible alternatives for the particular feedback type (e.g., one of several ways of the robot saying it is going to go trigger the lock) in order to prevent repeated phrasing.

## 6. STUDY DESIGN

Although the method presented for generating coordinating verbal communication is designed to be applicable in a wide range of task settings, we present the following user study as an initial proof of concept designed to evaluate the approach in a co-located human-robot task collaboration as compared to a non-communicating robot. The approach was evaluated by measuring the time to complete the pseudo-herding task in the augmented reality environment as well as collecting subjective measures of the robot’s performance via surveys. The communicating robot was implemented as described above and compared to a control, a non-communicating robot, i.e., a robot that performed the task in an identical manner without issuing any verbal feedback. This control was selected since few existing methods for high-level task planning integrate human feedback, making a task-only robot a likely scenario.

A within-subjects experiment was conducted, in which each participant performed the task twice, once with the communicating robot, and once with the non-communicating one, with order of presentation counterbalanced across subjects. A within-subjects design was selected since the overall interaction is quite short (usually about 2-5 minutes) and our pilot experiments found the times to finish the task varied a great deal among people. Also, since the task is unfamiliar to all users due to its unique setup, there is the potential for each person to learn the nuances of the task over time, and improve performance. Participants watched an experimenter explain and perform each of the task elements, including: herding a sheep and triggering the lock and light. Participants were told that the goal of the game was to herd all of the sheep as quickly as possible. The



**Figure 3: Time to completion and allocation counts of agent actions in communicating and silent conditions.**

random seeds used to generate the sheep behavior were initialized identically for each participant to keep the simulated sheep behavior as consistent as possible between runs. The task simulator was configured with 12 sheep. The speeds of the sheep and timing of the lock and light elements were tested in a pilot experiment. The task thus calibrated was possible but somewhat challenging for a single person to perform alone. The static robot policy used for the task was favored light- and sheep-collection, resulting in the robot’s first priority being turning on the light and collecting and returning sheep to the pen area while the light is on. The robot’s task behavior was identical for each condition with the on-board speakers either enabled or disabled for the talking and silent robots, respectively.

### 6.0.1 Data Collected

Data collected include a record of the underlying task simulator state, the tracked locations of the agents (people, robot, sheep and virtual objects), an audio recording of each person from a wireless headset microphone, and video of the interaction from overhead and side-view cameras. After completion of the task participants were asked to take a survey asking demographic questions as well as a series of 27 questions about the robot as a teammate, rated on a 7-point Likert scale from “strongly agree” to “strongly disagree”.

### 6.0.2 Participants

A total of sixteen participants (5 female, 11 male) were recruited from campus sources. Participants’ ages ranged from 17-29, with a mean of 21. Most had completed at least some undergraduate college education. One participant was an outlier, not able to complete the task due to a misunderstanding of the rules; this participant’s data were discarded.

### 6.0.3 Hypotheses and Outcome Measures

The following hypotheses were formulated:

**Hypothesis 1:** Participants will prefer the communicating robot and consider it a better teammate.

**Hypothesis 2:** The communicating robot will improve team task performance.

## 7. RESULTS

For the objective results, the time to complete the task to the nearest second was recorded for each trial by watching the recorded video, marking the start of the task as signaled by the experimenter, and the end as noted by the freezing of the projected display upon retrieval of the final sheep. Trials with the communicating robot were completed on average 17 seconds faster ( $M = 100.5$  s,  $SD = 21$  s) than the ones with the non-communicating robot ( $M = 117.6$  s,  $SD = 40$  s),  $p < 0.03$ . As our conditions were repeated within subjects, we then performed a post hoc analysis by fitting a linear mixed effects regression and comparing to a baseline single mean model. In comparing the two models, we find incorporating the communication factor has a significant impact on the duration ( $p < 0.04$ ), confirming Hypothesis 2.

In the pseudo-herding task, in the quantifiable terms of state value, the biggest mistake that can be made is to let the timed lock lapse, releasing all collected sheep and forcing the team to start over. We intentionally selected a static robot policy in which the lock was ignored by the robot, requiring the participant to manage it instead. Across all 30 trials, the lock was allowed to lapse in 2 of 15 trials in each condition. Because these costly errors were roughly evenly distributed across conditions and were relatively rare with ( $M = 0.4$ ,  $SD = 1.1$ ), it is unlikely that the objective performance improvements are due to an avoidance of mistakes; it is more likely due to the clarification of responsibilities provided by the verbal feedback and minimizing the user’s context switching between lock and light (Figure 4).

Distinguishing the impact of each type of feedback is challenging as they are issued while the robot and the person are moving about the room, making subtle effects difficult to discern from planned, task-directed actions. The most common feedback issued was self-narrative as the robot narrated its actions whenever they differed from what it was doing previously. Since the role-allocative feedback is very specific, with post hoc analysis we can compare the robot’s suggested action with the action ultimately undertaken by the person. Over the course of the 15 trials in which the robot used verbal communication, and therefore was capable of assigning roles, the robot made a total of 12 role-allocation suggestions to users. To assess users’ adherence to the robot’s requests, we assessed whether a matching action for the participant was recognized within a short time interval. In total, participants adhered to all 12 requests within 10 seconds and often much faster. This adherence rate is likely due to the low cost of completing the robot’s requests as well as the novelty of the interaction. Still, anecdotally, there are multiple instances of participants abandoning a prior plan and altering course in response to the robot’s suggestion.

The low overall number of role-allocation suggestions is likely due to most participants allowing the robot to take full responsibility for the light. Also, since the robot attempts to track the person’s active policy using a stream of recognized actions (currently from the task simulation), it only accounts for completed actions when determining the user’s most likely policy. This has the effect of ignoring additional information provided by the person’s trajectory, i.e., motion in the direction of a given object, and relies on the user reaching a steady state of task performance where they perform the same types of actions repeatedly. Anecdotally the refresh rate of the various communication models often appears to result in a narrative phrase and role-allocative

Table 1: Post-experiment survey results

Question	0 (strongly disagree) - 6 (strongly agree)	
	M	SD
The things the robot said made sense.	5.27	1.6
The talking robot was a better teammate than the silent robot.	5	1.4
The task was more fun with the robot than if I had done it alone.	5.4	0.8
The talking robot was more fun.	5.3	1.2
The robot’s talking helped me understand what it was going to do next.	4.9	1.9
The things the robot said helped me decide what to do.	3.8	2.0
I tried to do what the robot told me to do.	3.6	2.4

phrase happening immediately after one another, such as “I’ll go get the light. Can you take care of the lock?”, often resulting in the user expressing agreement. The survey data reveal that all but two people agreed that the talking robot was a better teammate (see Table 1 for a summary). The communicating robot had similarly high scores for being more fun and for its verbal feedback making sense during the interaction. Questions asking how participants adjusted their behavior in response to the communicating robot are less consistently agreed with, perhaps due to people taking credit for their own actions and the team’s success or failure.

## 8. CONCLUSION

In this paper we presented an approach for robot production of social communication during human-robot task collaboration to improve in situ decision-making and team performance. The problem of generating communication from a robot to a human teammate for the purposes of coordinating joint activity is formulated as a planning problem compatible with a class of pairwise, loosely-coupled tasks and supporting the notion of role assignment common to many human collaboration scenarios. An approach was developed integrating recognition of a person’s intentional role-driven behavior using spatiotemporal information with planning of the robot’s verbal feedback using a Markov decision process representation of the task environment. The approach uses sets of policies to capture the action selection preferences for different roles and to reason about the compatibility of the robot’s planned actions with the inferred action of a human teammate. A validation in a synthetic task environment with an autonomous robot collaborating with a single person was conducted as a proof-of-concept, in which the communicating robot was shown to significantly improve the task completion time of the team.

In future work we plan several enhancements to the communication planning approach and a series of additional validations aimed at determining the generalizability of the approach to different task settings. Further evaluations will be conducted including implementation in a real-world service robot task environment with actual, as opposed to simu-

lated, object manipulation. This will require extension to partially-observable environments and human activity monitoring with actual objects. We also plan on using a humanoid to allow for co-verbal gestures to augment the existing verbal feedback. Another enhancement includes investigating methods for adapting and personalizing the robot's task and communication preferences to individual users based on their own verbal feedback as well as their adherence to the robot's requests. Additionally, as the communication planning system requires a set of policies and associated speech actions to function properly, we will be examining methods for automatically segmenting these roles, and associated verbal and visual references, from human demonstrations.

## 9. ACKNOWLEDGMENTS

This work is supported in part by ONR MURI program (N00014-09-1-1031), and National Science Foundation (NSF) grants CNS-0709296, IIS-1117279, IIS-0803565.

## 10. REFERENCES

- [1] C. Breazeal, C. D. Kidd, A. L. 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 2005). 2005 IEEE/RSJ International Conference on*, pages 708–713. IEEE, 2005.
- [2] G. Buccino, F. Binkofski, and L. Riggio. The mirror neuron system and action recognition. *Brain and language*, 89(2):370–376, 2004.
- [3] S. Chernova and M. Veloso. Interactive policy learning through confidence-based autonomy. *Journal of Artificial Intelligence Research*, 34(1):1, 2009.
- [4] C. Crick and B. Scassellati. Inferring narrative and intention from playground games. *Proceedings of the 7th IEEE International Conference on Development and Learning (ICDL'08)*, pages 13–18, August 2008.
- [5] A. D. Dragan, K. C. Lee, and S. S. Srinivasa. Legibility and predictability of robot motion. In *Human-Robot Interaction (HRI), 2013 8th ACM/IEEE International Conference on*, pages 301–308. IEEE, 2013.
- [6] B. Grosz and C. Sidner. Attention, intentions, and the structure of discourse. *Computational linguistics*, 12(3):175–204, 1986.
- [7] R. Hamid, S. Maddi, A. Johnson, A. Bobick, I. Essa, and C. Isbell. A novel sequence representation for unsupervised analysis of human activities. *Artificial Intelligence*, 173(14):1221–1244, 2009.
- [8] G. Hoffman and C. Breazeal. Cost-based anticipatory action selection for human-robot fluency. *Robotics, IEEE Transactions on*, 23(5):952–961, 2007.
- [9] R. Kelley, C. King, A. Tavakkoli, M. Nicolescu, M. Nicolescu, and G. Bebis. An architecture for understanding intent using a novel hidden markov formulation. *International Journal of Humanoid Robotics, Special Issue on Cognitive Humanoid Robots*, 5(2):1–22, 2008.
- [10] B. Keysar, D. Barr, J. Balin, and J. Brauner. Taking perspective in conversation: The role of mutual knowledge in comprehension. *Psychological Science*, 11(1):32, 2000.
- [11] T. Kollar, S. Tellex, D. Roy, and N. Roy. Toward understanding natural language directions. In *Human-Robot Interaction (HRI), 2010 5th ACM/IEEE International Conference on*, pages 259–266. IEEE, 2010.
- [12] C. Lockridge and S. Brennan. Addressees' needs influence speakers' early syntactic choices. *Psychonomic Bulletin & Review*, 9(3):550, 2002.
- [13] 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 Annual Meeting*, volume 42, pages 123–127. SAGE Publications, 1998.
- [14] B. Mutlu, A. Terrell, and C.-M. Huang. Coordination mechanisms in human-robot collaboration. In *Proceedings of the Workshop on Collaborative Manipulation, 8th ACM/IEEE International Conference on Human-Robot Interaction*, 2013.
- [15] S. Nikolaidis and J. Shah. Human-robot cross-training: computational formulation, modeling and evaluation of a human team training strategy. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction*, pages 33–40. IEEE Press, 2013.
- [16] M. Pickering and S. Garrod. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27(02):169–190, 2004.
- [17] J. Shah, J. Wiken, B. Williams, and C. Breazeal. Improved human-robot team performance using chaski, a human-inspired plan execution system. In *Proceedings of the 6th international conference on Human-robot interaction*, pages 29–36. ACM, 2011.
- [18] C. Sidner, C. Kidd, C. Lee, and N. Lesh. Where to look: a study of human-robot engagement. In *Proceedings of the 9th International Conference on Intelligent User Interfaces*, page 84. ACM, 2004.
- [19] R. B. Smith, R. Hixon, and B. Horan. Supporting flexible roles in a shared space. In *Collaborative Virtual Environments*, pages 160–176. Springer, 2001.
- [20] J. G. Trafton, N. L. Cassimatis, M. D. Bugajska, D. P. Brock, F. E. Mintz, and A. C. Schultz. Enabling effective human-robot interaction using perspective-taking in robots. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 35(4):460–470, 2005.
- [21] T. Ullman, C. Baker, O. Macindoe, O. Evans, N. Goodman, and J. Tenenbaum. Help or hinder: Bayesian models of social goal inference. *Advances in Neural Information Processing Systems (NIPS)*, 22, 2010.
- [22] R. Wilcox, S. Nikolaidis, and J. Shah. Optimization of temporal dynamics for adaptive human-robot interaction in assembly manufacturing. *Robotics*, page 441, 2013.