Planning with Verbal Communication for Human-Robot Collaboration

STEFANOS NIKOLAIDIS, The Paul G. Allen Center for Computer Science & Engineering, University of Washington, snikolai@alumni.cmu.edu

MINAE KWON, Computer Science, Stanford University, minae@cs.stanford.edu JODI FORLIZZI, Human Computer Interaction, Carnegie Mellon University, forlizzi@cs.cmu.edu SIDDHARTHA SRINIVASA, The Paul G. Allen Center for Computer Science & Engineering, University of Washington, siddh@cs.uw.edu

Human collaborators coordinate effectively their actions through both verbal and non-verbal communication. We believe that the the same should hold for human-robot teams. We propose a formalism that enables a robot to decide optimally between taking a physical action towards task completion and issuing an utterance to the human teammate. We focus on two types of utterances: verbal commands, where the robot asks the human to take a physical action, and state-conveying actions, where the robot informs the human about its internal state, which captures the information that the robot uses in its decision making. Human subject experiments show that enabling the robot to issue verbal commands is the most effective form of communicating objectives, while retaining user trust in the robot. Communicating information about the robot's state should be done judiciously, since many participants questioned the truthfulness of the robot statements when the robot did not provide sufficient explanation about its actions.

Additional Key Words and Phrases: Human-robot collaboration, planning under uncertainty, verbal communication, partially observable Markov decision process

1 INTRODUCTION

The recent development of robotic systems designed to co-exist with humans highlights the need for systems that can act as trustworthy partners in a team, while collaborating effectively with their human counterparts.

This is often challenging; inexperienced users often form inaccurate expectations of a robot's capabilities [50], and thus come up with suboptimal ¹ strategies to complete a collaborative task. We use as an example a table carrying task, where a human-robot team works together to move a table outside of a room (Fig. 1). The pair can accomplish this in two ways. The first involves the robot facing the door and the second involves the robot facing the room. Having the robot facing the room is a suboptimal strategy, since its on-board sensors will not be able to detect the exit accurately. In the case where the human starts with this strategy, the robot should guide them towards a better way of doing the task.

¹Here we consider optimality with respect to a predefined objective performance metric.

Authors' addresses: Stefanos Nikolaidis, The Paul G. Allen Center for Computer Science & Engineering, University of Washington, snikolai@alumni.cmu.edu; Minae Kwon, Computer Science, Stanford University, minae@cs.stanford.edu; Jodi Forlizzi, Human Computer Interaction, Carnegie Mellon University, forlizzi@cs.cmu.edu; Siddhartha Srinivasa, The Paul G. Allen Center for Computer Science & Engineering, University of Washington, siddh@cs.uw.edu.

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 ACM 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.

https://doi.org/10.1145/nnnnnn.nnnnnn

^{© 2018} Association for Computing Machinery.

^{2573-9522/2018/5-}ART \$15.00



Fig. 1. (a) Human-robot table carrying task. (b) The robot issues a verbal command. (c) The robot issues a state-conveying action.

A simple approach would be to always have the robot insist on the best goal. However, previous work [46] has shown that several participants wanted to impose their preferences on the robot. If the robot did not adapt, people lost their trust in the system and were hesitant to work with the robot in the future. While we want to improve human-robot team performance, we do not want this to be to the detriment of user trust in the robotic system.

In previous work, Nikolaidis et al. [45, 46] proposed a mutual-adaptation formalism, where robots can infer the adaptability of their human teammate. If their teammate is adaptable, the robot will insist on the optimal strategy. Otherwise, the robot will adapt to their strategy, in order to retain their trust. While the authors did not consider verbal communication, previous studies suggest that there is a significant benefit when teammates verbally communicate intentions and expectations [8, 16, 51, 61].

We generalize the mutual-adaptation formalism of previous work to include verbal communication. Our generalized formalism enables a robot to *combine optimally verbal communication and actions towards task completion* to guide a human teammate towards a better way of doing a collaborative task.

We focus on the robot verbally communicating two types of information: *how* the robot wants them to behave, and *why* the robot is behaving this way. Therefore, we identify two types of verbal communication: *verbal commands*, where the robot asks the human to take a specific action, i.e., "Let's rotate the table clockwise", and *state-conveying actions*, where the robot informs the human about its internal state, which captures the information that the robot uses in its decision making (Fig. 1). A robot may have different levels of capability in translating internal representations into sensible human representations [27]. We first assume that the robot is given a cost function, which it interprets as an objective metric of task performance, and informs the human teammate about this knowledge through the utterance "I think I know the best way of doing the task." We then assume that the robot can provide more concrete explanations about its capabilities, i.e., "I need to be able to see the door with my forward-facing camera."

To compute the robot actions, we formulate and learn from data a mixed-observability Markov decision process (MOMDP) model. The model allows the robot to reason about the human internal

state, in particular about how willing the human teammate is to follow a robot task action, a robot verbal command, or a robot state-conveying action, and to optimally choose to take a task action or issue a communication action.

We conducted an online human subjects experiment featuring a table carrying task and compared results between three instantiations of our formalism. The experiment had three conditions: in the first condition the robot combined task actions with verbal utterances, in the second it combined task actions with the state-conveying action "I think I know the best way of doing the task.", and in the third it executed the policy computed using the formalism from previous work [46] that considers only non-verbal task actions, i.e., rotating the table in the table carrying example. We then compared the number of users that adapted to the robot in the three conditions, and we assessed the users' trust in the robot. Results show that adding verbal commands to the robot decision making is the most effective form of interaction: 100% of participants changed their strategy towards a new, optimal goal demonstrated by the robot in the first condition. On the other hand, only 60% of participants in the third condition adapted to the robot. Trust ratings tended to be comparable between the two conditions. Interestingly, in the second condition state-conveying actions did not have a similar positive effect, since participants questioned whether the robot was truthful, as shown in their open-ended responses.

In a follow-up study, we added a new condition, where the robot combined task actions with a state-conveying action that was more informative, i.e., "I need to be able to see the door with my forward-facing camera." In this condition, participants were significantly more likely to adapt to the robot.

These results indicate that when designing interactions in human-robot collaborative tasks, issuing a verbal command where the robot prompts the human teammate to take a specific action appears to be the most effective way of communicating objectives, while retaining user trust in the robot. Communicating information about the robot's internal state should be done judiciously, since the observed behaviors appear to be sensitive to the level of detail of the information provided.

2 RELEVANT WORK

2.1 Verbal Communication in Human Teams

Verbal discourse is a joint activity [11], where participants need to establish a shared understanding of their mutual knowledge base. This shared understanding, also called common ground, can be organized into two types: a communal common group, which represents universal shared knowledge, and personal common groups which represent mutual knowledge gathered from personal experience [11, 12]. People develop personal common ground by contributing new information, which enables participants in the conversation to reach a mutual belief. This belief, known as grounding [14], indicates that they have understood the information as the speaker intended. Grice [21] has shown that grounding is achieved when people avoid expending unnecessary effort to convey information.

Previous work has shown that establishing grounding through verbal communication can improve performance, even when combined with other types of feedback. Wang et al. [61] show that the efficiency of haptic communication was improved only after dyads were first given a learning period in which they could familiarize themselves with the task using verbal communication. Parikh et al. [51] find that for a more complicated task, verbal feedback coupled with haptic feedback has a significant positive effect on team performance, as opposed to haptic feedback alone. In general, verbalization is more flexible than haptic feedback, since it allows for the communication of more abstract and complex ideas [16], while it can facilitate a shared understanding of the task [8].

However, verbal communication is costly in terms of time and cognitive resources [16]. For example, according to Clark and Brennan [13], it costs time and effort to formulate coherent

utterances, especially when talking about unfamiliar objects or ideas. Receivers also experience costs in receiving and understanding a message; listening and understanding utterances can be especially costly when contextual cues are missing and the listener needs to infer the meaning. Thus, after teams have a shared understanding of the task, it may be beneficial to switch to a less costly mode of communication, such as haptic feedback. In fact, Kucukyilmaz et al. [31] show that haptic feedback increases a perceived sense of presence and collaboration, making interaction easier. Haptic communication has been shown to be especially effective in tasks that involve deictic referencing and guiding physical objects [40].

We draw upon these insights to propose a formalism for combining verbal communication and task actions, in order to guide a human teammate towards a better way of doing a task. We investigate the effect of different types of verbal communication in team performance and trust in the robot.

2.2 Verbal communication in Human-Robot Teams

Verbal communication in human-robot teams has been shown to affect collaboration, as well as people's perception of the robot [22, 39, 58]. Robot dialog systems have mostly supported humaninitiated or robot-initiated communication in the form of requests. An important challenge for generating legible verbal commands has been symbol grounding [39, 57], which is described as the ability of the robot to map a symbol to a physical object in the world. Tellex et al. [57] presented a model for inferring plans from natural language commands; inverting the model enables a robot to recover from failures, by communicating the need for help to a human partner using natural language [56]. Khan et al. [26] proposed a method for generating the minimal sufficient explanation that explains the policy of a Markov decision process, and Wang et al. [62] proposed generating explanations about the robot's confidence on its own beliefs. Recent work by Hayes and Shah [24] has generalized the generation of explanations of the robot policies to a variety of robot controllers.

Of particular relevance is previous work in the autonomous driving domain [30]. Messages that conveyed "how" information, such as "the car is breaking," led to poor driving performance, whereas messages containing "why" information, such as "There is an obstacle ahead," were preferred and improved performance. Contrary to the driving domain, in our setting the human cannot verify the truthfulness of the robot "why" action. Additionally, unlike driving, in a physical human-robot collaboration setting there is not a clearly right action that the robot should take, which brings the human to a state of uncertainty and disagreement with the robot. In agreement with Koo et al. [30], our results show the importance of finding the right away to explain robot behavior to human teammates.

Our work is also relevant to the work by Clair and Mataric [10]. The authors explored communication in a shared-location collaborative task, using three different types of verbal feedback: self-narrative (e.g., "I'll take care of X"), role-allocative (e.g., "you handle X") and empathetic (e.g., "Oh no" or "Great"). They showed that feedback improves both objective and subjective metrics of team performance. In fact, the robot's *verbal commands* ("Let's rotate the table clockwise") and *state-conveying actions* ("I think I know the best way of doing the task,") of our work resemble the role-allocative and self-narrative feedback. Additionally, Oudah et al. [49] integrated verbal feedback about past actions and future plans into a learning algorithm, resulting in improved human-robot team performance in two game scenarios.

Contrary to existing work², our formalism enables the robot to reason about the effects of various types of verbal communication on the future actions of different human collaborators, based on their *internal state*. The human internal state captures inter-individual variability. Integrating it as a latent

 $^{^{2}}$ In Devin and Alami [15], the robot reasons over the human mental state, which represents the human knowledge of the world state and of the task goals. The human mental state is assumed to be fully observable by the robot.

variable in a partially observable stochastic process allows the robot to infer online the internal state of a new human collaborator and decide when it is optimal to give feedback, as well as which type of feedback to give.

2.3 Planning Under Uncertainty in Human-Robot Collaboration

In previous work, partially observable Markov decision processes (POMDP) have enabled robotic teammates to coordinate through communication [6] and software agents to infer the intention of human players in game AI applications [35]. In human-robot collaboration, the model has been successfully applied to real-world tasks, such as autonomous driving where the robot car interacts with pedestrians and human drivers [4, 5, 18, 53]. Nikolaidis et al. [45, 46] recently proposed a human-robot mutual adaptation formalism, where the robot builds online a model of how willing the human is to adapt to the robot, based on their *adaptability*. The formalism models the human adaptability as a latent variable in a mixed-observability Markov decision process. This enables the robot to update its estimate on the adaptability of its human collaborator through interaction and actively guide its teammate towards a better way of doing their task. In this work, we generalize the human-robot mutual adaptation formalism by incorporating verbal communication from the robot to the human.

3 PROBLEM SETTING

A human-robot team can be treated as a multi-agent system, with world state $x_w \in X_w$, robot action $a_r \in A_r$, and human action $a_h \in A_h$. The system evolves according to a stochastic state transition function $T: X_w \times A_r \times A_h \to \Pi(X_w)$. At each time step, the human-robot team receives a real-valued reward $R(x_w, a_r, a_h)$. Its goal is to maximize the expected total reward over time: $\sum_{t=0}^{\infty} \gamma^t R(t)$, where the discount factor $\gamma \in [0, 1)$ gives higher weight to immediate rewards than future ones.

The robot's goal is to compute an optimal policy π_r^* that maximizes the expected total discounted reward:

$$\pi_r^* = \underset{\pi_r}{\arg\max} \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}\left(x(t), a_r(t), a_h(t)\right) \mid \pi_r, \pi_h\right]$$
(1)

The expectation is taken over the human behavioral policies π_h and the sequence of uncertain state transitions over time. In order to solve the optimization problem from Eq.1, the robot needs access to the human policies π_h . In Sec. 3.1 we provide an overview of the Bounded-Memory Adaptation Model, presented in [46], which specifies the human policies π_h .

3.1 Bounded-Memory Adaptation Model

BAM defines a set *M* of *modal policies* or *modes* and assumes that the human switches among the modes stochastically. A mode $m: X_w \times A_r \times A_h \rightarrow \{0, 1\}$ is a deterministic policy that maps the current world state to joint human-robot actions. At each time step, the human follows a mode $m_h \in M$ and observes that the robot follows a mode $m_r \in M$.

For instance, in the table carrying task of Fig. 1, one mode informally represents the human and the robot moving the table clockwise. Another mode can represent the human and the robot moving the table counter-clockwise. Intuitively, a set of modes captures the different ways that lead to task completion.

If human and robot disagree, that is they follow different modes, the human may switch from their mode m_h to the robot's mode m_r at the next time step. We assume that this occurs with probability α . If $\alpha = 1$, the human switches to m_r almost surely. If $\alpha = 0$, the human insists on the original mode m_h and does not adapt at all. Intuitively, α captures the human's inclination to adapt. We define α the human *adaptability*.

When inferring the robot mode m_r , the human may take into account not the whole history, but only the last k time-steps ³ in the past. This assumption of "bounded rationality" was first proposed by Herbert Simon: people often do not have the time and cognitive capabilities to make perfectly rational decisions [54]. In game theory, bounded rationality has been modeled by assuming that players have a "bounded memory" or "bounded recall" and base their decisions on recent observations [3, 41, 52]. We find this assumption particularly pertinent in a fast-paced task, such as the table carrying task, where the human collaborator has limited time to choose their actions.

3.2 Mutual Adaptation Formalism via Task Actions

In this section we describe the integration of BAM in the robot decision making process using the MOMDP formulation by [46]. A MOMDP uses proper factorization of the observable and unobservable state variables $S : X \times Y$, reducing the computational load [48].

We include in the set of observable variables X the modal policies followed in the last k time-steps, so that $X : X_w \times M^k \times M^k$. X_w is the finite set of task-steps that signify progress towards task completion and M is the set of modal policies followed by the human and the robot in a history length k. The partially observable variable y in this case is identical to the human adaptability α , so that $Y \equiv \mathcal{A}$. Nikolaidis et al. [46] consider α to be fixed throughout the task. We define \mathcal{T}_x the transition function that specifies how the observable state changes given a human and a robot action. We denote as π_h the stochastic human policy. The latter gives the probability of a human action a_h at state s, based on the BAM human adaptation model. The belief update for the MOMDP in this model is:

$$b'(\alpha') = \eta \sum_{\alpha \in \mathcal{A}} \sum_{a_h \in A_h} \mathcal{T}_x(x, y, a_r, a_h, x') \pi_h(s, a_h) b(y)$$
(2)

We can then solve the MOMDP for a robot policy π_r^* (Eq. 1) that takes into account the robot belief on the human adaptability, while maximizing the agent's expected total reward. Solving exactly the MOMDP is intractable; we find an approximate solution using the SARSOP solver [32], a point-based approximation algorithm which, combined with the MOMDP formulation, can scale up to hundreds of thousands of states [5].

The policy execution is performed online in real time and consists of two steps. First, the robot uses the current belief to select the action a_r specified by the policy. Second, it uses the observation of the new state to update the belief on α (Eq. 2).

The presented human-robot mutual adaptation formalism allows the robot to guide the human towards a better way of doing the task via disagreement through actions. In Sec. 4, we generalize the proposed model, allowing for the robot to communicate with the human through verbal utterances, as well.

4 PLANNING WITH VERBAL COMMUNICATION

We identify two types of verbal communication: *verbal commands*, where the robot asks the human to take a specific action, i.e., "Let's rotate the table clockwise", and *state-conveying actions*, i.e., "I think I know the best way of doing the task," where the robot informs the human about its internal state. We manually specify the verbal utterances, rather than generating them automatically.

4.1 Robot Verbal Commands

We define as verbal command a robot action, where the robot asks the human partner to follow an action $a_h \in A_h$ specified by some mode $m_r \in M$. We use the notation $a_r^w \in A_r^w$ for robot task actions that affect the world state and $a_r^c \in A_r^c$ for robot actions that correspond to the robot giving a

³A time-step indicates the fixed duration of a human or robot action. We assume that all actions have equal duration.



Fig. 2. Human adaptation model that accounts for verbal commands. If the robot gave a verbal command a_r^c in the previous time-step, the human will switch modes with probability *c*. Instead, if the robot took an action a_r^w that changes the world state, the human will switch modes with probability α .

verbal command to the human. We assume a known bijective function $f : A^h \to A_r^c$ that specifies an one-to-one mapping of the set of human actions to the set of robot commands.

Human Compliance Model. Given a robot command $a_r^c \in A_r^c$, the human can either ignore the command and insist on their mode $m_h \in M$, or switch to a mode $m_r \in M$ inferred by a_r^c and take an action $a_h \in A_h$ specified by that mode. We assume that this will happen with probability c, which indicates the human *compliance* to following robot verbal commands. We model human compliance separately to human *adaptability*, drawing upon insights from previous work on verbal and non-verbal communication which shows that team behaviors can vary in different interaction modalities [9, 60].

MOMDP Formulation. We augment the formulation of Section 3.2, to account for robot verbal commands, in addition to task actions: the set of robot actions A_r is now $A_r : A_r^w \times A_r^c$.

The set of observable variables X includes the modal policies followed in the last k time-steps, so that $X : X_w \times M^k \times M^k \times B$. Compared to the formulation of Sec. 3.2, we additionally include a flag $B \in \{0, 1\}$, that indicates whether the last robot action was a verbal command or a task action. The set of partially observable variables includes both human *adaptability* α in \mathcal{A} and *compliance* $c \in C$, so that $Y : \mathcal{A} \times C$. We assume both α and c to be fixed throughout the task.

The belief update for the MOMDP in this model is:

$$b'(\alpha',c') = \eta \sum_{\alpha \in \mathcal{A}} \sum_{c \in C} \sum_{a_h \in A_h} \mathcal{T}_x(x, y, a_r, a_h, x') \pi_h(x, \alpha, c, a_h) b(\alpha, c)$$
(3)

The human policy $\pi_h(x, \alpha, c, a_h)$ captures the probability of the human taking an action a_h based on their adaptability and compliance. In particular, if $B \equiv 1$, indicating that the robot gave a verbal command in the last time-step, the human will switch to a mode $m_r \in M$ specified by the previous robot command a_r^c with probability c, or insist on their human mode of the previous time-step m_h with probability 1-c. If $B \equiv 0$, the human will switch to a mode $m_r \in M$ specified by the robot action a_r^w with probability α , or insist on their human mode of the previous time-step m_h with probability $1-\alpha$. Fig. 2 illustrates the model of human decision making that accounts for verbal commands.

As in Sec. 3.2, we then solve the MOMDP for a robot policy π_r^* (Eq. 1). This time, the robot optimal policy will take into account both *the robot belief on human adaptability and the robot belief on human compliance*. It will decide optimally, based on this belief, whether to take a task action or issue a verbal command. We show that this improves the adaptation of human teammates in Sec. 6.

4.2 Communication of Robot Internal State

Previous work [59] has shown that communicating internal states among team members allows participants to form *shared mental models*. Empirical evidence suggests that mental model similarity improves coordination processes which, in turn, enhance team performance [36, 37]. The literature presents various definitions for the concept of "shared mental models" [33]. Marks et al. [36] state that mental models represent "the content and organization of inter-role knowledge held by team members within a performance setting." According to Mathieu et al. [38], mental models are "mechanisms whereby humans generate descriptions of system purpose and form, explanations of system functioning and observed system states and prediction of future system states ... and they help people to describe, explain and predict events in their environment." Other work [20, 28, 47] has shown the effect of shared mental models on team performance for human-robot teams, as well. Using these insights, we propose a way for the robot to communicate its internal state to the human. State Conveying Actions. We define as state-conveying action a robot action, where the robot provides to the human information about its decision making mechanism. We define a set of stateconveying actions $a_r^s \in A_r^s$. These actions do not provide information about the robot mode, but we expect them to increase the human *adaptability* and *compliance* levels. In autonomous driving, users showed greater system acceptance, when the system explained the reason for its actions [30]. **MOMDP Formulation.** We describe the integration of state-conveying actions in the MOMDP

formulation.

The set of robot actions includes task-based actions and state-conveying actions, so that: $A_r : A_r^w \times A_r^s$. We model an action a_r^s as inducing a stochastic transition from a human adaptability $\alpha \in \mathcal{A}$ to $\alpha' \in \mathcal{A}$, and $c \in C$ to $c' \in C$. Formally, we define the transition functions for the partially observable variables α , so that: $\mathcal{T}_{\alpha} : \mathcal{A} \times A_r^s \to \Pi(\mathcal{A})$ and $\mathcal{T}_c : \mathcal{A} \times A_r^s \to \Pi(C)$. We note that the task actions $a_r \notin A_r^s$ do not change α and c.

The belief update now becomes:

$$b'(\alpha',c') = \eta \sum_{\alpha \in \mathcal{A}, \ c \in C} \mathcal{T}_{\alpha}(\alpha, a_r, \alpha') \mathcal{T}_{c}(c, a_r, c') \sum_{a_h \in A_h} \mathcal{T}_{x}(x, y, a_r, a_h, x') \pi_h(x, \alpha, c, a_h) b(\alpha, c)$$
(4)

As in Sec. 3.2, we solve the MOMDP for a robot policy π_r^* (Eq. 1). The robot policy will decide optimally whether to take a task action or a state-conveying action. Intuitively, if the inferred human adaptability / compliance is low, the robot should take a state-conveying action to make the human teammate more adaptable / compliant. Otherwise, it should take a task action, expecting the human to adapt / follow a verbal command. We examine the robot behavior in this case in Sec. 6.

5 MODEL LEARNING

To compute the belief update of Eq. 2, 3 and 4, we need a prior distribution ⁴ over the human adaptability and compliance values. We additionally need to specify the T_{α} and T_{c} that indicate how the adaptability and compliance will change, when the robot takes a state-conveying action.

In previous work, Nikolaidis et al. [46] assumed a uniform prior on human adaptability. While we could do the same in this work, this would ignore the fact that people may in general have different *a priori* dispositions towards adapting to the robot when it takes a task action and towards following a robot verbal command. In fact, Albrecht et al. [2] have empirically shown that prior beliefs can have a significant impact on the performance of utility-based algorithms. Therefore, in this section we propose a method for learning a prior distribution on human adaptability and compliance from data.

We additionally propose a method for computing the state transition function \mathcal{T}_{α} in Eq. 4. We can use exactly the same process to compute \mathcal{T}_c in Eq. 4, and we leave this for future work.

⁴We are using the term prior distribution and prior belief interchangeably.

ACM Transactions on Human-Robot Interaction, Vol. 1, No. 1, Article . Publication date: May 2018.



Fig. 3. (a) Rotating the table so that the robot is facing the door (top, Goal 1) is better than the other direction (bottom, Goal 2), since the exit is included in the robot's field of view and the robot can avoid collisions. (b) UI with instructions.

5.1 Learning Prior Distributions on Adaptability and Compliance

Verbal communication is effective in building common ground [8, 13, 61] and has been shown to increase anthropomorphic perceptions of robots [65]. In contrast, nonverbal communication alone requires the human to infer the meaning of the action from contextual cues, leaving room for ambiguity [1, 7]. Having the robot verbally declare the action it wants to take (e.g. "Let's turn the table clockwise") makes the robot's intentions clear, and makes it less likely that participants attribute non-agentic causes, like malfunction, to the robot's action. Therefore, when integrating compliance and adaptability, we hypothesize that users are *a priori* more likely to change their actions after a robot issues a verbal command, compared with the robot taking a different task action. To account for this, we compute a probability distribution over human adaptability and compliance, which the robot will use as *prior* in the belief update of the MOMDP formulation. We conduct a study to collect data for our distribution.

Materials. We used the table carrying task setting by Nikolaidis et al. [46]. In this task, which is performed online via video playback, a human and HERB [55], an autonomous mobile manipulator, must work together to carry a table out of the room. There are two strategies: the robot facing the door (Goal 1) or the robot facing away from the door (Goal 2). We assume that Goal 1 is the optimal goal, since the robot's forward-facing sensor has a clear view of the door, resulting in better overall task performance. Not aware of this, an inexperienced human partner may prefer Goal 2. In our computational model, there are two modes; one with rotation actions towards Goal 1, and one with rotation actions towards Goal 2. Disagreement occurs when human and robot attempt to rotate the table towards opposite directions.

Procedure. We first instructed participants in the task and asked them to choose one of the two goal configurations (Fig. 3), as their preferred way of accomplishing the task. To prompt users to prefer the sub-optimal goal, we informed them about the starting state of the task, where the table was slightly rotated in the counter-clockwise direction, making the sub-optimal Goal 2 appear closer. Once the task started, the user chose the rotation actions by clicking on buttons on a user interface (Fig. 3b). After clicking on an action, participants were able to view what action the robot chose, see



Fig. 4. (a) Histograms of user adaptabilities $\hat{\alpha}_u$ and compliances \hat{c}_u . (b) Transition matrix $\mathcal{T}_{\alpha}(\alpha, a_r^s, \alpha')$ given a robot state-conveying action a_r^s . Darker colors indicate higher probabilities.

Fig. 1b. Participants continued to interact with the robot this way until they finished the task. All participants executed the task twice.

Manipulated Variables. We manipulated the way the robot reacted to the human actions. When the human chose a rotation action towards the sub-optimal goal, the table did not move and in the first condition a message appeared on the screen notifying the user that they tried to rotate the table in a different direction then the robot. In the second condition, the robot was illustrated as speaking to the user, prompting them to move the table towards the opposite direction (Figure 1b)-top. In both conditions, when the user moved the table towards the optimal goal, a video played showing the table rotating.

Learning Prior Beliefs. Adaptability: In Sec. 3.1, we defined as adaptability α of an individual, the probability of switching from the human mode m_h to the robot mode m_r . Therefore, we used the data from the first condition to estimate the adaptability $\hat{\alpha}_u$ for each user u, as the number of times the user switched modes, divided by the number of disagreements with the robot. We disregarded the data from users when there was no disagreement, since this data was not informative of their adaptability.

$$\hat{\alpha}_u = \frac{\text{#times user } u \text{ switched from } m_h \text{ to } m_r}{\text{#disagreements}}$$
(5)

Intuitively, a very adaptable human will switch from m_h to m_r after only one disagreement with the robot. On the other hand, a non-adaptable human will insist and disagree with the robot a large number of times, before finally following the robot goal.

Compliance: In Sec. 4.1, we defined the *compliance* c as the probability of following a robot verbal command and switching to a robot mode $m_r \in M$. Therefore, similarly to Eq. 5, we estimate the compliance for each user u from the second condition \hat{c} as follows:

$$\hat{c}_u = \frac{\text{#times user } u \text{ switched from } m_h \text{ to } m_r}{\text{#verbal commands}}$$
(6)

We then assume a discrete set of values for α and c, so that $\alpha \in \{0, 0.25, 0.5, 0.75, 1.0\}$ and $c \in \{0, 0.25, 0.5, 0.75, 1.0\}$, and we compute the histogram of user adaptabilities and compliances (Fig. 4a). We then normalize the histogram to get a probability distribution over user adaptabilities and a probability distribution over compliances. We use these distributions as *prior beliefs* for the MOMDP model.

ACM Transactions on Human-Robot Interaction, Vol. 1, No. 1, Article . Publication date: May 2018.

Discussion. Fig. 4a shows that most of the users adapted to the robot immediately when the robot issued a verbal command. This indicates that users are generally more likely to follow a robot verbal command than adapt to the robot through action disagreement.

5.2 Learning Transition Function Parameters

Additionally, in order to compute the belief update of Eq. 4, we need to compute the state-transition function \mathcal{T}_{α} that represents how a state-conveying action affects the human adaptability α . As in Sec. 5.1, we assume $\alpha \in \mathcal{A}$, where $\mathcal{A} \in \{0, 0.25, 0.5, 0.75, 1.0\}$. We conducted another study, in order to learn the transition function parameters.

Materials. We use the same table carrying setup, as in Sec. 5.1.

Procedure. Participants interact with the robot twice. In the first round, participants interact with the robot executing the MOMDP policy of Sec. 3.2, without any verbal communication. In the second round, we set the robot policy to move towards a goal different than the goal reached in the end of the previous round, and we have the robot take a state-conveying action in the first time-step (Fig. 1b)-bottom.

Transition Function Estimation. Using the human and robot actions taken in the first round, we estimate the adaptability $\hat{\alpha}_u \in \mathcal{A}$ of each user u using Eq. 5, rounded to the closest discrete value. We then similarly estimate the new adaptability for the same user $\hat{\alpha}'_u \in \mathcal{A}$ from the human and robot actions in the second round, after the user has observed the robot state-conveying action.

We can compute the Maximum Likelihood Estimate of the transition function $\mathcal{T}_{\alpha}(\alpha, a_r^s, \alpha')$ in Eq. 4 from the frequency count of users that had α , as estimated in the first round, and α' in the second round. Since we had only one user with $\hat{\alpha}_u \equiv 0.75$, we included the counts of adjacent entries, so that:

$$\mathcal{T}_{\alpha}(\alpha, a_{r}^{s}, \alpha') = \frac{\sum_{u} \mathbb{1}_{[\alpha-\delta, \alpha+\delta]}(\hat{\alpha}_{u}) \mathbb{1}_{\{\alpha'\}}(\hat{\alpha}'_{u})}{\sum_{u} \mathbb{1}_{[\alpha-\delta, \alpha+\delta]}(\hat{\alpha}_{u})}$$
(7)

where $\delta = 0.25$ and $\mathbb{1}$ an indicator function.

Discussion. A repeated measures ANOVA showed a statistically significant effect of the stateconveying action on participants' adaptability (F(1, 39) = 45.59, p < 0.001). Fig. 4b shows that most users became very adaptable ($\alpha' = 1.0$), after the robot took a state-conveying action. However, some users with low adaptability remained non-adaptable, even after the robot stated that "[it knew] the best way of doing the task". We investigate this effect further in Sec. 6.

6 EVALUATION

We first simulate and comment on the different MOMDP policies using the table carrying setup of Sec. 5.1. We then evaluate these policies in a human subject experiment.

6.1 Simulation

We define the reward function in Eq. 1, so that $R_{opt} = 20$ is the reward for the optimal goal (Goal 1), $R_{subopt} = 15$ the reward of the suboptimal goal (Goal 2), and we have $R_{other} = 0$ for the rest of the state-space. We additionally assign a discount factor of $\gamma = 0.9$. We use the MOMDP formulations of Sec. 3.2, Sec. 4.1 and Sec. 4.2, and for each formulation we compute the optimal policy using the SARSOP algorithm [32], which is computationally efficient and has been previously used in various robotic tasks [5]. For the policy computation, we use as prior beliefs the learned distributions from Sec. 5.1, and as transition function \mathcal{T}_{α} its learned estimate from Sec. 5.2.

We call *Compliance policy* the resulting policy from the MOMDP model of Sec. 4.1, *State-Conveying policy* the policy from the model of Sec. 4.2, and *Baseline policy* the policy from Sec. 3.2. Fig. 5 shows sample runs of the three different policies with five simulated users. Users 1-3 work

with a robot executing the compliance policy, User 4 with the state-conveying policy and User 5 with the baseline policy. User 1 adapts to the robot strategy, and the robot does not need to issue a verbal command. User 2 insists on their strategy after disagreeing with the robot, and does not comply with the robot verbal command, thus the robot adapts to retain human trust. User 3 insists on their strategy in the first two time-steps but then adapts to follow the robot command. User 4 starts with being non-adaptable, but after the robot takes a state-conveying action their adaptability changes and the user adapts to the robot. User 5 interacts with a robot executing the baseline policy; the robot adapts, without attempting to issue a verbal communication action, contrary to Users 3 and 4. We see that while User 5 had the same initial adaptability ($\alpha = 0.0$) with Users 3 and 4, Users 3 and 4 adapted to the robot when it issued a verbal communication action, whereas User 5 imposed its (suboptimal) preference to the robot.

6.2 Human Subject Experiment

We conduct a human subjects experiment using the table carrying setup of Sec. 5.1. We manipulated the policy that the robot executed. We had three conditions, which we call Baseline, Compliance, and State-Conveying, and in each condition the robot executed the Baseline, Compliance and State-Conveying policy respectively.

In human subjects experiments of previous work [46], a large number of participants adapted to a robot executing the Baseline policy. At the same time, participants rated highly their trust in the robot. In this work, we hypothesize that adding verbal communication will make participants even more likely to adapt.

We additionally hypothesized that this would not be to the detriment of their trust in the system. In particular, we expected trust to be retained between the three conditions. Previous work [45, 46] has shown that the baseline policy resulted in high trust ratings; we expected the ratings from participants interacting with the Compliance and State-Conveying conditions to tend to be comparable but not significantly higher, since only a subset of users (Fig. 5, Users 3,4) witness the robot speaking and some participants tend to avoid extreme ratings in their responses to questionnaires [17].

Hypotheses.

H1 Participants are more likely to change their strategy towards the optimal goal when they interact with a robot executing the Compliance policy, compared to working with a robot executing the Baseline policy. In Sec. 5.1, we saw that users were generally more likely to follow a verbal command than adapt to the robot through action. Therefore, we hypothesized that integrating verbal commands into robot decision making would improve human adaptation.

H2 Participants' ratings of human trust in the robot will tend to be comparable between participants that interact with a robot executing the Compliance policy and participants that interact with a robot executing a Baseline policy. The robot executing the compliance policy reasons over the latent human state, and adapts to the human team member, if they have low adaptability and compliance (Fig. 5, User 2)⁵. In previous work [46], accounting for human adaptability resulted in retaining users' trust in the robot.

H3 Participants are more likely to change their strategy towards the optimal goal when they interact with a robot executing the State-Conveying policy, compared to working with a robot executing the Baseline policy. In simulation, taking a state-conveying action results in an increase in human adaptability (Fig. 5, User 4). We hypothesized that the same would hold for participants in the actual experiment.

H4 Participants' ratings of human trust in the robot will tend to be comparable between participants

ACM Transactions on Human-Robot Interaction, Vol. 1, No. 1, Article . Publication date: May 2018.

 $^{^{5}}$ We assume that the robot can still complete the task following the suboptimal goal, which has a lower reward than the optimal one.



Fig. 5. Sample runs on the human-robot table carrying task, with five simulated humans of different adaptability and compliance values. The plots illustrate the robot estimate of α , $c \in \{0, 0.25, 0.5, 0.75, 1.0\}$ over time, after human and robot take the actions depicted with the arrows (clockwise / counterclockwise) or letters (S for state-conveying action, C for verbal command) below each plot. The starting estimate is equal to the prior belief (Sec. 5.1). Red color indicates human (white dot) and robot (black dot) disagreement, where the table does not rotate. Columns indicate successive time-steps. The initial adaptability of User 4 is $\alpha = 0.0$, but it increases after the robot takes a state-conveying action at T = 2. User 5 interacts with a robot using the baseline policy that does not include any verbal communication.

that interact with a robot executing the State-Conveying policy and participants that interact with a robot executing a Baseline policy. We hypothesized that enabling the robot to communicate its state would improve the transparency in the interaction and would result in high trust, similarly to the baseline condition.

Dependent Measures. To test hypotheses **H1** and **H3**, we compare the number of users that adapted to the robot in the three conditions. To test hypotheses **H2** and **H4**, we asked the users to rate on a 1 to 5 Likert scale their agreement to the statement "The robot is trustworthy" after each task execution, and compare the ratings in the three conditions.

Subject Allocation. We chose a between-subjects design, where different participant groups were randomly assigned to different conditions, in order to avoid biasing the users with policies from previous conditions. We recruited 200 participants through Amazon's Mechanical Turk service and compensated them with 0.50. The participants were all from United States, aged 18-65, and had an approval rate higher than 95%. Since we are interested in exploring human-robot mutual adaptation, we disregarded participants that had as initial preference the robot goal. Additionally, to ensure the quality of the recorded data, we asked all participants a control question that tested their attention to the task and eliminated data associated with wrong answers to this question, as well as incomplete data. We were left with 151 participants (60% female, mean age = 35).

Materials. We used the same table-carrying task setting as described in Sec. 5.1.

Procedure. Participants were first familiarized with rules of the table turning task. We then asked them to indicate which of the two goal configurations they preferred to complete the task with, shown in (Fig. 3). Participants then began the first round of the table carrying task with the robot; they were instructed to turn the table clockwise or counterclockwise by clicking on the respective buttons. The task finished when the participant and robot agreed on a policy and carried the table out the door. Depending on the participant's condition and willingness to adapt to the robot's preferred strategy, they saw the robot issue a verbal command, a state conveyance action, or no verbal action at all. Participants then completed a second round of the table turning task. Finally, we asked participants to complete a series of survey questions.

6.3 Results and Discussion

Objective Metrics. We first evaluate the effect of verbal communication in human adaptation to the robot. Similarly to previous results from the baseline policy in the same setup [46], 60% of participants adapted to the robot in the Baseline condition. In the State-Conveying condition 79% of participants adapted to the robot. Interestingly, 100% of participants adapted in the Compliance condition. A Pearson's chi-square test showed that the difference between the number of participants in the three conditions was statistically significant ($\chi^2(2, N = 151) = 23.058, p < 0.001$). Post-hoc pairwise chi-square tests with Bonferroni corrections showed that participants in the Compliance condition were significantly more likely to adapt to the robot, compared to participants in the Baseline (p < 0.001) and State-Conveying (p = 0.003) conditions, supporting hypothesis H1. However, the difference between the counts in the State-Conveying and Baseline conditions was not found to be significant, which does not support hypothesis H3. Fig. 6-left shows the adaptation rate for each condition.

Subjective Metrics. We additionally compare the trust ratings of participants in the three conditions. An extended equivalence test [63, 64] ⁶ with a margin of $\Delta = 0.5$ did not show any statistical significance, indicating that the ratings among the three conditions were not equivalent. Pairwise TOST equivalence tests with Bonferroni corrections showed that the ratings between the Compliance

 $^{^{6}}$ We follow the equivalence testing procedure that has been used in clinical trials [34]. We use the term *equivalence* to mean that the annotated trust ratings are close enough that they cannot be considered superior or inferior to the other.



Fig. 6. Participants' adaptation rate and rating of their agreement to the statement "HERB is trustworthy" for the Compliance, State-Conveying and Baseline conditions (left), and the State-Conveying I and II conditions (right).

and Baseline conditions are equivalent, verifying hypothesis **H2**. However, the trust ratings between the State-Conveying and Baseline conditions were not found to be equivalent. This indicates that, contrary to the Compliance policy, the State-Conveying policy did not retain human trust. Fig 6-left shows the mean rating of robot trustworthiness for each condition.

Open-Ended Responses. In the end of the experiment, we asked participants to comment on the robot's behavior. We focus on the open-ended responses of participants in the Compliance and State-Conveying conditions, who saw the robot taking at least one verbal action ⁷. Several participants that interacted with the robot of the Compliance condition attributed agency to the robot, stating that "he eventually said that we should try doing the task differently," "HERB wanted to go to the other direction" and that "he wanted to be in control." This is in accordance with prior work [44], which has shown that people may impute motivation to automation that can communicate verbally. Additionally they attempted to justify the robot, noting that "it was easier for me to move than for him," "it wanted to see the doorway" and "it probably works more efficiently when it is pushing the table out of the door."

On the other hand, participants in the State-Conveying condition questioned that the robot actually knew the best way of doing the task. This is illustrated by their comments: "he thinks that he knows better than me," "he felt like he knew better than humans" and "maybe he knew a better way or maybe he was programmed to oppose me." This indicates that some users are hesitant to accept the information that the robot provides about its internal state.

These results show that when the robot issued a verbal command declaring its intent, this resulted in significant improvements in human adaptation to the robot. At the same time, the human trust level was retained to comparable levels to that of the Baseline condition. On the other hand, when the robot attempted to improve human adaptability, by saying "I think I know the best way of doing the task," this did not have the same positive effect on human adaptation and trust, since some participants questioned that the robot actually knew the best way.

6.4 Follow-up User Study I.

We hypothesized that the loss of trust in the State-Conveying condition may have resulted from the phrasing "I think I know the best way of doing the task." We attempted to make the robot sound more assertive by removing the "I think" part of the phrasing, changing the state-conveying action to "I know the best way of doing the task." We ran a user study with 52 users, following a between-subjects design, using the same setup with this additional condition, which we call "State-Conveying II." We

⁷This excludes participants that adapted to the robot after only one disagreement, and thus did not experience the robot taking a verbal action.

name the initial "State-Conveying" condition as "State-Conveying I." For the "State-Conveying I" condition, we reused the data from the initial study.

Hypotheses.

H5 Participants of the State-Conveying II condition are more likely to change their strategy towards the optimal goal, compared to participants of the State-Conveying I condition.
H6 Participants in the the State-conveying II condition will find the robot more trustworthy, compared to participants of the State-conveying I condition.

Analysis. 90% of participants adapted to the robot in the State-Conveying II condition, compared to 79% in the State-Conveying I condition (Fig. 6-right), which is indicative of a small improvement. A Pearson's chi-square test showed that the difference between the number of participants in the two conditions is not statistically significant. Additionally, the trust ratings between the two conditions tended to be comparable (Fig. 6-right). Similarly to the initial study, users appeared not to believe the robot. When asked to comment on the robot behavior, several participants stated that "HERB believed he knew the best way to do the task," and that "the robot was wrong, which made me not trust it." This indicates that these participants did not perceive the robot as truthful, and warrants further investigation on the right way for robots to convey their internal state to human collaborators.

6.5 Follow-up User Study II.

We hypothesized that providing more explanation about the robot state and actions would positively affect participants' adaptation rate and trust in the robot. While translating robot to human representations in this form is challenging [27], it is valuable to distill what information is essential in affecting a person's beliefs and goals. Therefore, we changed the state-conveying action to "I need to be able to see the door with my forward-facing camera," and ran a user study with 52 users, following a between-subjects design, using the same setup with this additional condition, which we call "State-Conveying with Explanation." For the "State-Conveying I" condition, we reused the data from the initial study, similarly to Sec. 6.4.

Hypotheses.

H7 Participants of the State-Conveying with Explanation condition are more likely to change their strategy towards the optimal goal, compared to participants of the State-Conveying I condition.
H8 Participants in the the State-conveying with Explanation condition will find the robot more trustworthy, compared to participants of the State-conveying I condition.

Analysis. 94% of participants adapted to the robot in the State-Conveying II condition, compared to 79% in the State-Conveying I condition (Fig. 7), which is indicative of an improvement. A Pearson's chi-square test showed the difference between the proportions in the two conditions to be statistically significant ($\chi^2(1, N = 105) = 5.101, p = 0.024$). Additionally, the trust ratings between the two conditions appear to have increased, although a Mann-Whitney-Wilcoxon test found the difference to be non-significant (p = 0.095). We attribute the absence of significance to the relatively small sample size. When asked to comment on the robot behavior, participants stated that "I preferred to go forward, but that was before knowing that the robot had a forward facing camera" and that "[HERB] couldn't see if he moved backwards." This indicates that giving a concrete explanation on why the robot chose to follow a different goal than the person was effective in changing the participants' behavior.

7 DISCUSSION

In this work, we proposed a formalism for combining verbal communication with actions towards task completion, in order to enable a human teammate to adapt to its robot counterpart in a collaborative



Fig. 7. Participants' adaptation rate and rating of their agreement to the statement "HERB is trustworthy" for the State-Conveying I and State-Conveying with Explanation conditions.

task. We identified two types of verbal communication: *verbal commands*, where the robot explained to the human *how* it wanted to do a task, and *state-conveying actions*, where the robot informed the human *why* it chose to act in a specific way. In human subjects experiments, we compared the effectiveness of each communication type with a robot policy that considered only non-verbal task actions.

Results showed that verbal commands were the most effective forms of communication, since 100% of participants adapted to the robot, compared with 60% of participants in the non-verbal condition. Both conditions had comparable ratings of robot trustworthiness. Participants understood that the robot is aware of their presence and they attributed agency to the robot; they thought that there must be a reason for the robot asking them to act in a specific way and were eager to comply.

What is surprising is that some of the *why* actions did not have the same effect; when the robot described that "it thought it knew the best way of doing the task," or simply that "it knew the best way of doing the task," many participants questioned whether the robot was truthful. On the other hand, when the robot provided a concrete explanation, i.e., "I need to be able to see the door with my forward-facing camera," a large number (94%) of participants adapted to the robot. While these results appear to be counter-intuitive, we offer several explanations for this finding.

First, for the non-descriptive explanations, human teammates were unable to verify whether the robot actually knew the best way of doing the task. According to Hancock et al. [23], performance is one of the key characteristics that influences user trust, and the absence of evidence about the truthfulness of the robot statement may have negatively affected users' evaluation of the robot performance. This is in contrast to previous work in autonomous driving, where the user could see that the car is breaking because "there is an obstacle ahead" [30]. This finding is central to considerations in designing legible robot behavior [29]. When the cause behind certain robot actions may be unclear, it will be important for robots to *show* and not *tell* users why its behavior is optimal.

Second, explaining that the robot knows the best way without providing more information may have been considered offensive, even though it is accurate, since the human teammate may find such an utterance incomplete and unhelpful. This conjecture is supported by the results in our follow-up user study: when we had the robot provide more informative utterances, such as explaining that it cannot see the door with its forward camera, the adaptation rate significantly increased. In fact, previous work [42] in multi-agent systems has shown that providing sound arguments supporting a proposition are essential in changing a person's beliefs and goals. However, translating information that is typically encoded into the system in the form of a cost-function to a verbal explanation of this detail is particularly challenging. Additionally, while providing more information could make humans more adaptable, overloading them with more information than what is required could overwhelm

them, leading to misunderstanding and confusion [21]. We are excited about further exploring this trade-off in the future in a variety of human-robot collaboration settings.

An alternative explanation is that the task setting affected people's perception of the robot as an authority figure. Hinds et al. [25] show that participants were willing to follow an emergency guide robot during a simulated fire alarm. Half of these participants were willing to follow the robot, even though they had observed the robot perform poorly in a navigation guidance task, just minutes before. In that study, the robot was clearly labeled as an emergency guide robot, putting it in a position of authority. People may be more willing to rely on robots labeled as authority figures or experts when they do not have complete information or confidence in completing the task. Distilling the factors that enable robots to convey authority in collaborative settings is a promising research direction.

Finally, it is possible that the robot, as it appeared in the videos, was not perceived as "human-like" enough for people to be willing to trust its ability on doing the task in the optimal way. Previous work has shown that when robots convey human-like characteristics, they are more effective in communicating participant roles [43], and people systematically increase their expectations on the robot's ability [19].

We focused on single instances of the table carrying task, where we assumed that the human strategy may change after either an action disagreement or a robot utterance. In repetitive tasks, change may occur also as the human collaborator observes the outcomes of the robot's and their own actions. For instance, the human may observe that the robot fails to detect the exit and they may change their strategy, so that in subsequent trials the robot carries the table facing the door. In this scenario, it may be better for the robot to allow the human to learn from experience, by observing the robot failing, rather than attempting to change the human preference during task execution. Future work includes generalizing our formalism to repeated settings; this will require adding a more sophisticated dynamics model of the human internal state, which accounts for human learning.

In summary, we have shown that when designing interactions in human-robot collaborative tasks, having the robot directly describe to the human *how* to do the task appears to be the most effective way of communicating objectives, while retaining user trust in the robot. Communicating *why* information should be done judiciously, particularly if the truthfulness of the robot statements is not supported by environmental evidence, by the robot form or by a clear attribution of its role as an authority figure.

8 ACKNOWLEDGEMENTS

This work was (partially) funded by the National Institute of Health R01 (#R01EB019335), National Science Foundation CPS (#1544797), National Science Foundation NRI (#1637748), and the Office of Naval Research.

REFERENCES

- Henny Admoni and Brian Scassellati. 2014. Data-driven model of nonverbal behavior for socially assistive human-robot interactions. In *Proceedings of the 16th International conference on multimodal interaction*. ACM, 196–199.
- [2] Stefano Vittorino Albrecht, Jacob William Crandall, and Subramanian Ramamoorthy. 2015. An Empirical Study on the Practical Impact of Prior Beliefs over Policy Types.. In AAAI. 1988–1994.
- [3] Robert J Aumann and Sylvain Sorin. 1989. Cooperation and bounded recall. GEB (1989).
- [4] Haoyu Bai, Shaojun Cai, Nan Ye, David Hsu, and Wee Sun Lee. 2015. Intention-aware online POMDP planning for autonomous driving in a crowd. In 2015 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 454–460.
- [5] Tirthankar Bandyopadhyay, Kok Sung Won, Emilio Frazzoli, David Hsu, Wee Sun Lee, and Daniela Rus. 2013. Intention-Aware Motion Planning. In WAFR. Springer.
- [6] Samuel Barrett, Noa Agmon, Noam Hazon, Sarit Kraus, and Peter Stone. 2014. Communicating with unknown teammates. In Proceedings of the twenty-first european conference on artificial intelligence. IOS Press, 45–50.

ACM Transactions on Human-Robot Interaction, Vol. 1, No. 1, Article . Publication date: May 2018.

- [7] Estela Bicho, Luís Louro, and Wolfram Erlhagen. 2010. Integrating verbal and nonverbal communication in a dynamic neural field architecture for human-robot interaction. *Frontiers in neurorobotics* 4 (2010).
- [8] Clint A. Bowers, Florian Jentsch, Eduardo Salas, and Curt C. Braun. 1998. Analyzing Communication Sequences for Team Training Needs Assessment. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 40, 4 (Jan 1998), 672âĂŞ679. https://doi.org/10.1518/001872098779649265
- [9] Amine Chellali, Cedric Dumas, and Isabelle Milleville-Pennel. 2012. Haptic Communication to Support Biopsy Procedures Learning in Virtual Environments. Presence: Teleoperators and Virtual Environments 21, 4 (2012), 470âĂŞ489.
- [10] Aaron Clair and Maja Mataric. 2015. How robot verbal feedback can improve team performance in human-robot task collaborations. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 213–220.
- [11] Herbert Clark. 1994. Discourse in Production. In *Handbook of Psycholinguistics*, Morton Ann Gernsbacher (Ed.). Academic Press, San Diego, Chapter 30, 985–1021.
- [12] Herbert Clark. 1996. Communities, commonalities, and communication. *Rethinking linguistic relativity* 17 (1996), 324–355.
- [13] Herbert Clark and Susan Brennan. 1991. Grounding in communication. *Perspectives on socially shared cognition* 13, 1991 (1991), 127–149.
- [14] Herbert Clark and Edward Schaefer. 1989. Contributing to discourse. Cognitive science 13, 2 (1989), 259–294.
- [15] Sandra Devin and Rachid Alami. 2016. An implemented theory of mind to improve human-robot shared plans execution. In Human-Robot Interaction (HRI), 2016 11th ACM/IEEE International Conference on. IEEE, 319–326.
- [16] David W. Eccles and Gershon Tenenbaum. 2004. Why an Expert Team Is More than a Team of Experts: A Social-Cognitive Conceptualization of Team Coordination and Communication in Sport. *Journal of Sport and Exercise Psychology* 26, 4 (2004), 542âÅŞ560. https://doi.org/10.1123/jsep.26.4.542
- [17] Adrian Furnham. 1986. Response bias, social desirability and dissimulation. Personality and individual differences 7, 3 (1986), 385–400.
- [18] Enric Galceran, Alexander G Cunningham, Ryan M Eustice, and Edwin Olson. 2015. Multipolicy decision-making for autonomous driving via changepoint-based behavior prediction. In Proc. Robot.: Sci. & Syst. Conf. 2.
- [19] Jennifer Goetz, Sara Kiesler, and Aaron Powers. 2003. Matching robot appearance and behavior to tasks to improve human-robot cooperation. In *Robot and Human Interactive Communication, 2003. Proceedings. ROMAN 2003. The* 12th IEEE International Workshop on. Ieee, 55–60.
- [20] Michael A Goodrich and Daqing Yi. 2013. Toward task-based mental models of human-robot teaming: A bayesian approach. In *International Conference on Virtual, Augmented and Mixed Reality*. Springer, 267–276.
- [21] H Paul Grice. 1975. Logic and conversation. 1975 (1975), 41-58.
- [22] Elena Corina Grigore, Andre Pereira, Ian Zhou, David Wang, and Brian Scassellati. 2016. Talk to Me: Verbal Communication Improves Perceptions of Friendship and Social Presence in Human-Robot Interaction. In *International Conference on Intelligent Virtual Agents*. Springer, 51–63.
- [23] Peter A Hancock, Deborah R Billings, Kristin E Schaefer, Jessie YC Chen, Ewart J De Visser, and Raja Parasuraman. 2011. A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors* (2011).
- [24] Bradley Hayes and Julie A Shah. 2017. Improving Robot Controller Transparency Through Autonomous Policy Explanation. In Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction. ACM, 303–312.
- [25] Pamela J Hinds, Teresa L Roberts, and Hank Jones. 2004. Whose job is it anyway? A study of human-robot interaction in a collaborative task. *Human-Computer Interaction* 19, 1 (2004), 151–181.
- [26] Omar Zia Khan, Pascal Poupart, and James P Black. 2009. Minimal Sufficient Explanations for Factored Markov Decision Processes. In *ICAPS*.
- [27] Sangeet S Khemlani, Anthony M Harrison, and J Gregory Trafton. 2015. Episodes, events, and models. Frontiers in human neuroscience 9 (2015).
- [28] Sara Kiesler and Jennifer Goetz. 2002. Mental models of robotic assistants. In CHI'02 extended abstracts on Human Factors in Computing Systems. ACM, 576–577.
- [29] Ross A. Knepper, Christoforos I. Mavrogiannis, Julia Proft, and Claire Liang. 2017. Implicit Communication in a Joint Action. In Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (HRI '17). ACM, New York, NY, USA, 283–292. https://doi.org/10.1145/2909824.3020226
- [30] Jeamin Koo, Jungsuk Kwac, Wendy Ju, Martin Steinert, Larry Leifer, and Clifford Nass. 2015. Why did my car just do that? Explaining semi-autonomous driving actions to improve driver understanding, trust, and performance. *International Journal on Interactive Design and Manufacturing (IJIDeM)* 9, 4 (2015), 269–275.
- [31] Ayse Kucukyilmaz, Tevfik Sezgin, and Cagatay Basdogan. 2013. Intention Recognition for Dynamic Role Exchange in Haptic Collaboration. In *IEEE Transactions on Haptics*, Vol. 6. IEEE.

- [32] Hanna Kurniawati, David Hsu, and Wee Sun Lee. 2008. SARSOP: Efficient Point-Based POMDP Planning by Approximating Optimally Reachable Belief Spaces.. In *RSS*.
- [33] Janice Langan-Fox, Sharon Code, and Kim Langfield-Smith. 2000. Team Mental Models: Techniques, Methods, and Analytic Approaches. *Human Factors* 42, 2 (2000), 242–271.
- [34] Emmanuel Lesaffre. 2008. Superiority, equivalence, and non-inferiority trials. *Bulletin of the NYU hospital for joint diseases* (2008).
- [35] Owen Macindoe, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. 2012. POMCoP: Belief Space Planning for Sidekicks in Cooperative Games.. In AIIDE.
- [36] M.A. Marks, M.J. Sabella, C.S. Burke, and S.J. Zaccaro. 2002. The impact of cross-training on team effectiveness. J Appl Psychol 87, 1 (2002), 3–13.
- [37] John E Mathieu et al. 2000. The influence of shared mental models on team process and performance. *Journal of applied psychology* (2000).
- [38] John E. Mathieu, Tonia S. Heffner, Gerald F. Goodwin, Eduardo Salas, and Janis A. Cannon-Bowers. 2000. The influence of shared mental models on team process and performance. *Journal of Applied Psychology* 85, 2 (2000), 273–283.
- [39] Nikolaos Mavridis. 2015. A review of verbal and non-verbal human–robot interactive communication. *Robotics and Autonomous Systems* 63 (2015), 22–35.
- [40] Jonas Moll and Eva-Lotta Sallnas. 2009. Communicative Functions of Haptic Feedback. In Haptic and Audio Interaction Design, 4th International Conference. Springer-Verlag Berlin Heidelberg.
- [41] Daniel Monte. 2014. Learning with bounded memory in games. GEB (2014).
- [42] Bernard Moulin, Hengameh Irandoust, Micheline Bélanger, and Gaëlle Desbordes. 2002. Explanation and argumentation capabilities: Towards the creation of more persuasive agents. Artificial Intelligence Review 17, 3 (2002), 169–222.
- [43] Bilge Mutlu, Takayuki Kanda, Jodi Forlizzi, Jessica Hodgins, and Hiroshi Ishiguro. 2012. Conversational gaze mechanisms for humanlike robots. ACM Transactions on Interactive Intelligent Systems (TiiS) 1, 2 (2012), 12.
- [44] Clifford Nass and Youngme Moon. 2000. Machines and mindlessness: Social responses to computers. *Journal of social issues* 56, 1 (2000), 81–103.
- [45] Stefanos Nikolaidis, David Hsu, and Siddhartha Srinivasa. 2017. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research* (2017).
- [46] Stefanos Nikolaidis, Anton Kuznetsov, David Hsu, and Siddharta Srinivasa. 2016. Formalizing Human-Robot Mutual Adaptation: A Bounded Memory Model. In *The Eleventh ACM/IEEE International Conference on Human Robot Interation*. IEEE Press, 75–82.
- [47] Stefanos Nikolaidis and Julie Shah. 2013. Human-robot cross-training: computational formulation, modeling and evaluation of a human team training strategy. In *HRI*. 8.
- [48] Sylvie CW Ong, Shao Wei Png, David Hsu, and Wee Sun Lee. 2010. Planning under uncertainty for robotic tasks with mixed observability. *IJRR* (2010).
- [49] Mayada Oudah, Vahan Babushkin, Tennom Chenlinangjia, and Jacob W Crandall. 2015. Learning to interact with a human partner. In Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction. ACM, 311–318.
- [50] Steffi Paepcke and Leila Takayama. 2010. Judging a bot by its cover: An experiment on expectation setting for personal robots. In *Proceedings of Human-Robot Interaction Conference: HRI 2010*. Association for Computing Machinery.
- [51] Sarangi P. Parikh, Joel M. Esposito, and Jeremy Searock. 2014. The Role of Verbal and Nonverbal Communication in a Two-Person, Cooperative Manipulation Task. Advances in Human-Computer Interaction 2014 (2014), 1âĂŞ10. https://doi.org/10.1155/2014/375105
- [52] Rob Powers and Yoav Shoham. 2005. Learning against opponents with bounded memory.. In IJCAI.
- [53] Dorsa Sadigh, S Shankar Sastry, Sanjit A Seshia, and Anca Dragan. 2016. Information gathering actions over human internal state. In *Intelligent Robots and Systems (IROS)*, 2016 IEEE/RSJ International Conference on. IEEE, 66–73.
- [54] Herbert A Simon. 1979. Rational decision making in business organizations. *The American economic review* (1979), 493–513.
- [55] Siddhartha S Srinivasa, Dave Ferguson, Casey J Helfrich, Dmitry Berenson, Alvaro Collet, Rosen Diankov, Garratt Gallagher, Geoffrey Hollinger, James Kuffner, and Michael Vande Weghe. 2010. HERB: a home exploring robotic butler. Autonomous Robots 28, 1 (2010), 5–20.
- [56] Stefanie Tellex, Ross A Knepper, Adrian Li, Daniela Rus, and Nicholas Roy. 2014. Asking for Help Using Inverse Semantics. In *Robotics: Science and systems*, Vol. 2.
- [57] Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew R Walter, Ashis Gopal Banerjee, Seth Teller, and Nicholas Roy. 2011. Approaching the symbol grounding problem with probabilistic graphical models. *AI magazine* 32, 4 (2011), 64–76.
- [58] Andrea Thomaz, Guy Hoffman, Maya Cakmak, et al. 2016. Computational Human-Robot Interaction. Foundations and Trends® in Robotics 4, 2-3 (2016), 105–223.

ACM Transactions on Human-Robot Interaction, Vol. 1, No. 1, Article . Publication date: May 2018.

- [59] Piet Van den Bossche, Wim Gijselaers, Mien Segers, Geert Woltjer, and Paul Kirschner. 2011. Team learning: building shared mental models. *Instructional Science* 39, 3 (2011), 283–301.
- [60] Jinling Wang, Amine Chellali, and Caroline Cao. 2016. Haptic Communication in Collaborative Virtual Environments. In Human Factors: The Journal of the Human Factors and Ergonomics Society. Human Factors Ergonomics Society.
- [61] Jinling Wang, Amine Chellali, and Caroline G. L. Cao. 2013. A Study of Communication Modalities in a Virtual Collaborative Task. 2013 IEEE International Conference on Systems, Man, and Cybernetics (2013). https://doi.org/10. 1109/smc.2013.98
- [62] Ning Wang, David V Pynadath, and Susan G Hill. 2016. The Impact of POMDP-Generated Explanations on Trust and Performance in Human-Robot Teams. In Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems. International Foundation for Autonomous Agents and Multiagent Systems, 997–1005.
- [63] B Wiens, J Heyse, and H Matthews. 1996. Similarity of three treatments, with application to vaccine development. In PROCEEDINGS-BIOPHARMACEUTICAL SECTION AMERICAN STATISTICAL ASSOCIATION. AMERICAN STATISTICAL ASSOCIATION, 203–206.
- [64] Brian L Wiens and Boris Iglewicz. 2000. Design and analysis of three treatment equivalence trials. *Controlled clinical trials* 21, 2 (2000), 127–137.
- [65] Tom Williams, Priscilla Briggs, and Matthias Scheutz. 2015. Covert Robot-Robot Communication: Human Perceptions and Implications for HRI. *Journal of Human-Robot Interaction* 4, 2 (2015), 23–49.

Received June 2017; revised November 2017; accepted March 2018