2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 14-18, 2024. Abu Dhabi, UAE

Aligning Learning with Communication in Shared Autonomy

Joshua Hoegerman, Shahabedin Sagheb, Benjamin A. Christie, and Dylan P. Losey

2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) | 979-8-3503-7770-5/24/$31.00 ©2024 IEEE | DOI: 10.1109/IROS58592.2024.10801897

***Abstract*— Assistive robot arms can help humans by partially automating their desired tasks. Consider an adult with motor impairments controlling an assistive robot arm to eat dinner. The robot can reduce the number of human inputs — and how precise those inputs need to be — by recognizing what the human wants (e.g., a fork) and assisting for that task (e.g., moving towards the fork). Prior research has largely focused on *learning* the human's task and providing meaningful assistance. But as the robot learns and assists, we also need to ensure that the human understands the robot's intent (e.g., does the human know the robot is reaching for a fork?). In this paper, we study the effects of *communicating* learned assistance from the robot back to the human operator. We do not focus on the specific interfaces used for communication. Instead, we develop experimental and theoretical models of a) how communication changes the way humans interact with assistive robot arms, and b) how robots can harness these changes to better align with the human's intent. We first conduct online and in-person user studies where participants operate robots that provide partial assistance, and we measure how the human's inputs change with and without communication. With communication, we find that humans are more likely to intervene when the robot incorrectly predicts their intent, and more likely to release control when the robot correctly understands their task. We then use these findings to modify an established robot learning algorithm so that the robot can correctly interpret the human's inputs when communication is present. Our results from a second in-person user study suggest that this combination of *communication and learning* outperforms assistive systems that isolate either learning or communication. See videos here: https://youtu.be/BET9yuVTVU4**

I. Introduction

More than 24 million American adults need external assistance when performing activities of daily living [1]. Assistive robot arms that *share autonomy* with humans have the potential to help address this challenge [2], [3]. In these shared autonomy settings the human controls the robot arm using an input device (e.g., a joystick) to indicate their intent, and the robot helps automate tasks on the human's behalf (e.g., picking up foods and feeding them to the operator).

To achieve seamless assistance, both the human operator and robot arm must be on the same page. Consider Figure 1, where a human is using a robot arm to manipulate kitchen items. The human wants the robot to pick up a fork, and so the human provides joystick inputs that guide the robot towards that goal. *For the robot to align with the human*, the robot must *learn* from these inputs to determine the human's intent and partially automate their task. Here the robot might correctly infer what the human wants (e.g., a fork) and then coordinate its own motions to help reach that goal (e.g., fixing any errors in the human's inputs to precisely pick up the fork). On the other hand — *for the human to align with the robot* — the robot needs to *communicate* its intended assistance back to the user. Without this communication the human does not know what to expect from the robot: is the robot going to help automate the motion to the fork, or does the robot think the human wants something else entirely?

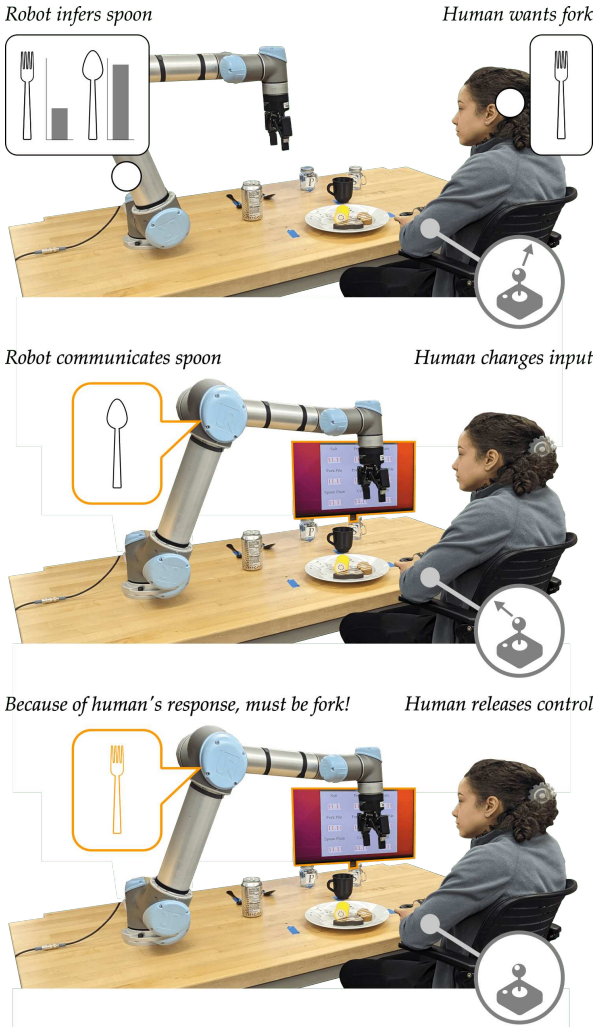

Fig. 1. Human sharing control with an assistive robot arm. (Top) The robot tries to infer the correct task from the human's joystick inputs. (Middle) We show that — when the robot communicates what it has inferred — the way humans provide inputs *changes*. (Bottom) If robots are aware of these changes, they can more accurately infer the human's goal.

Existing research on shared autonomy has largely separated learning and communication. On the one hand, methods such as [4]–[8] focus on inferring the human's task and partially automating the robot's motion, but do not consider

This work was supported by NSF Grants #2222468 and #2129201.
The authors are members of the Collaborative Robotics Lab (Collab), Dept. of Mechanical Engineering, Virginia Tech, Blacksburg, VA 24061.
e-mail: {jhoegerm, shahab, benc00, losey}@vt.edu

communication back to the human. On the other hand, approaches like [9]–[11] develop visual and haptic communication interfaces for shared autonomy, but do not modify the robot's learning algorithm. In this paper, we explore the intersection of learning and communication within shared autonomy settings. More specifically, we hypothesize that:

Humans will interact with shared autonomy systems differently *when those systems communicate their learning.*

This is important because — if humans do provide different inputs in the presence of communication — then the way the robot interprets and learns from human actions should also be modified. Accordingly, our paper has two main parts. First, in Section III-B we test our hypothesis and measure how communication can affect the way humans interact with assistive robot arms. Second, in Sections IV and V we harness the changes caused by communication to modify the robot's learning algorithm. In practice, this combination of learning and communication enables a) the robot to more seamlessly infer the human's task, and b) the human to more clearly indicate their intent. Returning to Figure 1, perhaps the human stops providing inputs because they observe from the robot's feedback that the fork is the robot's most likely goal. In response, our robot is able to confirm its prediction (i.e., because the human released control the robot must be correct), and complete the task more efficiently.

Overall, we make the following contributions:

Measuring the Effects of Communication. We consider shared autonomy settings where a human is operating a robot arm, and the robot updates the likelihood of each potential task based on the human's inputs. For these settings, we perform online and in-person user studies with and without robot communication. We find evidence that humans behave differently in the presence of communication.

Updating the Robot's Learning Rule. Our experimental results suggest that — when communication is present — humans are more likely to intervene if the robot has inferred the wrong task, and more likely to relinquish control if the robot is correct. We use these findings to modify the human model of an existing shared autonomy algorithm.

Combining Learning and Communication. We conduct another in-person user study with three conditions: learning (where the robot does not provide explicit feedback), communication (where the robot communicates its intent but does not adjust its learning rule), and our proposed approach. Our results suggest that the combination of learning and communication increases subjective and objective performance in shared autonomy settings.

II. Related Works

Below we discuss shared autonomy research that focuses on either learning (i.e., inferring the task and providing assistance) or communication (i.e., visual and haptic interfaces to convey the robot's internal state).

Learning in Shared Autonomy. Shared autonomy is a collaborative framework for human-robot interaction where the robot's behavior is a blend of the human's inputs and the robot's autonomous assistance [12]. The human's inputs convey the high-level task (e.g., grasping a fork), and the robot's inputs provide fine-grained corrections (e.g., coordinating the motion of the arm to reach that fork). Prior works develop algorithms to learn both the high-level task and low-level assistance method. For example, in [4]–[8], [13] the human's desired task is to reach a goal from a discrete set of options, and the robot infers this goal based on the human's inputs. As the robot becomes more confident in which goal the human wants, it can increasingly provide assistance to automate that task. Similarly, in [14]–[17] the robot builds an estimate of the task's reward function, and overrides any accidental or incorrect human inputs that would result in poor performance (e.g., preventing the human from moving the robot arm into a collision). Other methods such as [18]–[22] learn to assist the human by imitating their previous behaviors. For instance, if the human showed the robot how to pick up a fork in a past interaction, the robot leverages that data to help pick up forks during future interactions. Overall, each of these works provide a way for the robot to learn from and assist the human. However, they do not explicitly communicate what the robot has learned — hence, the user may not know what to expect from the system.

Communication in Shared Autonomy. Research outside of shared autonomy suggests that communicating robot learning has benefits for both the human and the robot. From the human's perspective, communication increases the user's acceptance and trust in the system [23]; from the robot's perspective, communication can result in more effective human teaching and accelerated robot learning [24]. Accordingly, recent works have started to apply communication strategies to shared autonomy [25].

In some scenarios, it is possible for the robot to *implicitly* convey what it has learned by exaggerating its motions [26]. However, for the robot to clearly indicate its latent state in everyday settings, *explicit* communication with visual, auditory, or haptic interfaces is often necessary. In [9] and [11] augmented reality headsets show the operator what the robot has learned about their high-level task (e.g., placing visual markers at the most likely goals) and how the robot plans to assist (e.g., displaying the robot's planned trajectory). Similarly, in [10] a wearable haptic interface notifies the human when the shared autonomy system is uncertain about their intent. Each of these interfaces has benefits and drawbacks, but our work seeks to focus on the addition of an application of a visual interface given its positive prior results and wider applicability for subject use. Our paper will build upon these related works by using explicit communication to convey the robot's inferred task back to the human. However, instead of focusing on the communication interface itself, we are interested in the effects of this communication on the human operator and assistive agent.

III. Effects of Communication

We consider shared autonomy settings where the human and robot collaborate in achieving a common goal. A key

aspect of shared autonomy is the ability of the robot to infer the human's goal (i.e., the task they are trying to complete). If the robot correctly infers the human's goal, it can complete the remaining task without requiring further human input. Alternatively, if the robot's inference is incorrect, the human must keep providing inputs towards their intended goal. However, it can be challenging for humans to determine what goal the robot has inferred without explicit communication.

In this section, we investigate how explicitly communicating the robot's belief about the human's goal affects their actions. We first introduce the policies of the human and the robot collaborator in the absence of communication. Then, we conduct a user study to understand the role of communication in shared-autonomy settings and determine how the users' actions change when communication is introduced. We aim to use these findings to improve the robot's inference of the human's goal and provide better assistance.

A. Shared Autonomy without Communication

We let $s \in \mathcal{S}$ be the environment state which includes the state of the robot, $a_{\mathcal{H}} \in \mathcal{A}$ and $a_{\mathcal{R}} \in \mathcal{A}$ be the human's and robot's actions respectively. The environment state transitions based on both the human and robot actions.

$$s^{t+1} = f\left(s^t, a_{\mathcal{H}}, a_{\mathcal{R}}\right) \tag{1}$$

We assume that the human chooses actions to minimize an internal cost-value function $Q^\star$:

$$a_{\mathcal{H}} \sim \pi^\star_{\mathcal{H}}\left(\circ \mid s, \theta, Q^\star\right) \tag{2}$$

Correspondingly, as the robot is trying to achieve the same goal as the human, it should take actions that minimize the human's cost-value function $Q^\star$. The robot does not directly observe the human's goal or their cost-value function. Instead, the robot selects actions according to an approximation Q of the cost-value function from prior work [4] where the robot's belief is not directly communicated to the human:

$$a_{\mathcal{R}} \sim \pi^\star_{\mathcal{R}}\left(\circ \mid s, b\left(\theta\right), Q\right) \tag{3}$$

where $b(\theta)$ is the robot's *belief* of the human's goal θ.

We suspect that the human's actions will change in the presence of communication. If the belief communicated by the robot aligns with the human's goal — will the human continue to provide actions that navigate the robot towards their goal or will they allow the robot to assume full control? On the other hand, if the robot's belief is incorrect — will the human *exaggerate* their corrective actions because they know that the robot's belief is incorrect? To evaluate how real users respond to robots that communicate their belief, we conducted two user studies with and without communication.

B. Shared Autonomy with Communication

We performed online and in-person user studies to gain insight into the effect of communication on shared autonomy. Participants collaborated with a robot to reach a goal while choosing how much input they think is enough for the robot to learn the task. In half of the interactions, the robot communicated its current belief of the user goal as a percentage using a digital interface. Our results from 25 online users and 10 in-person users show that people provide less input when the robot communicates its belief over the user goal with a noted subjective preference for the communicative system.

Experimental Setup. In the online study, participants taught a robot to reach a goal in multiple shared autonomy settings. In each setting, there were three objects with varying colors and the user's goal was to reach the green square (see Figure 2 (Left)). The position of these objects varied between settings. To simulate the settings we used an animated 2D environment with a top-down view. Online participants first watched the beginning of the robot arm's motion and then selected their choice of input to guide the robot toward the desired goal or to allow the robot to continue on its partially demonstrated path. In the in-person study, users commanded a robot arm using a Logitech F710 gamepad to perform a similar task of reaching a green cube.

We had five different settings for the objects in the online study and three settings in the in-person study. All participants interacted with the robot in each setting twice — with and without communication. In total, participants had six interactions in the in-person study and ten interactions in the online study. Each interaction ended when the robot reached the correct goal. The order of the settings was randomly counterbalanced across all users.

Independent Variables. For the online study, the users interacted with the robot in each setting across two variations. In one variation, the users had to infer the robot's intended goal through its animated motion (**Without Interface**). In the other variation, users were provided with the probabilities of the robot's belief over the goals (**With Interface**) in addition to their observation of the robot's motion.

For the in-person study, the robot used a state-of-the-art shared autonomy algorithm [4] to select its action $a_{\mathcal{R}}$ in each setting. In half of the interactions, the robot communicated its current belief $b(\theta)$ as percentages using a digital interface (**With Interface**), and for the other half, the users had to infer the robot's belief from its motion (**Without Interface**).

Dependent Variables. In both studies, we focused on how the user responses change when performing the shared autonomy tasks with and without communication. For the online study, we recorded whether the human chose to command the robot toward the desired goal or not. For the in-person study, we recorded the time that users spent using the gamepad and other joystick inputs (*Total Human Inputs*) as well as their subjective responses on a 7-point Likert scale for whether they preferred the settings with explicit communication or without the interface.

Participants. For the online study, we recruited 25 anonymous participants. We included an instruction and a qualifying question at the beginning of the survey for this study. For the in-person study, we recruited 10 participants from the Virginia Tech community (2 female, ages 23 ± 9 years). All participants provided informed consent as per university guidelines (IRB #20-755). To assist the participants in becoming familiar with the gamepad and the robot we provided

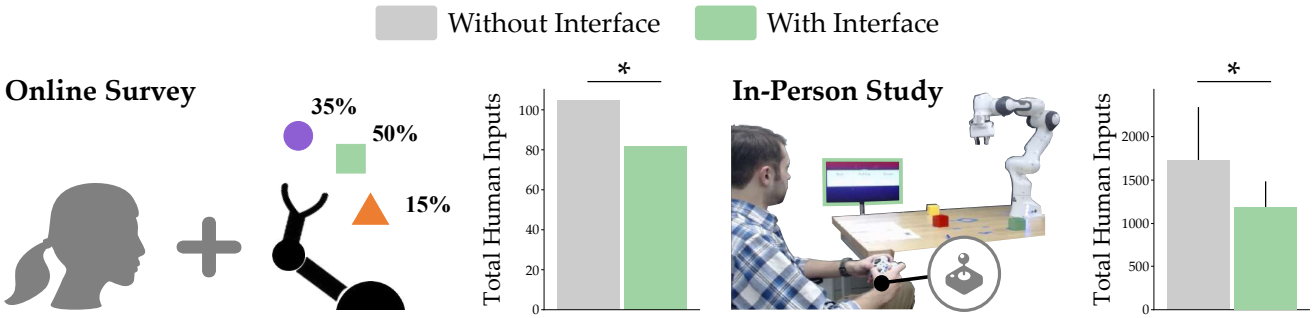

Fig. 2. Example settings and results from our user studies in Section III-B. Here we explored how communicating the robot's inferred distribution over a discrete set of tasks affected the human's inputs during shared autonomy. In all conditions, the robot used the same learning algorithm. (Left) Results from the online survey with and without a communication interface. Humans were more likely to release control to an assistive robot that conveyed its learned distribution over the tasks ($t(24) = 4.271, p < 0.005$). (Right) Corresponding results from our in-person study. Here humans required fewer inputs to guide the robot to their goal when the robot communicated its learning ($t(29) = 2.986, p < 0.005$). Overall, these results suggest that humans are more willing to yield control to a communicative system. An asterisk (*) denotes statistical significance.

practice time at the beginning of the interaction.

Hypothesis. We hypothesized that:

> **H1.** *When the robot communicates its belief over the goal, users will require less effort in commanding the robot to reach the desired goal.*
>
> **H2.** *Users will prefer using a shared autonomy system where the robot's belief is communicated.*

Results. Our results from the online and in-person user studies are summarized in Figure 2. To address **H1**, we evaluate the level of *effort* that users exhibited through the number of human inputs given through the gamepad. Here, there was a significant difference ($t(24) = 4.271, p < 0.005$) in the requisite human effort to reach the given goals. This result shows that when the robot communicates its belief, the user no longer has to provide the same level of effort for the robot to reach the goal, supporting **H1**.

For **H2**, we turn to our Likert-scale survey. We performed a Paired-Samples T-Test across polled user preferences for communication; these results were significant ($t(9) = 17.676, p < 0.001$). In our in-person user study, participants preferred interacting with a robot that communicated its belief of the human's intent.

IV. Harnessing Communication to Improve Learning

Our results from the first user study (Sec. III-B) demonstrate that humans behave differently in settings with communication than those without it. In this section, we leverage the human's response to the robot's communication in a novel shared-autonomy formalism. Instead of solely using communication to aid the human's guidance of the robot, we treat the human's feedback to the communication as an indication of the user's confidence in the robot.

We use this idea to present model human policies for both modalities: in the presence and absence of communication. The robot policy uses the appropriate human model to choose assistive actions that minimize the human's modality-specific cost-value function.

Human. The human takes actions that minimize their internal cost-value function Q. Following previous works [5], [26], we model the human as a nosily rational agent according to the Boltzmann distribution:

$$\pi_{\mathcal{H}}\left(a_{\mathcal{H}} \mid s, \theta\right) = \frac{\exp\left(\beta \cdot Q\left(s, a_{\mathcal{H}}, \theta\right)\right)}{\int \exp\left(\beta \cdot Q\left(s, a'_{\mathcal{H}}, \theta, \right)\right) da'_{\mathcal{H}}} \tag{4}$$

Here, $\pi_{\mathcal{H}}$ is a model of the human's true policy $\pi^{\star}_{\mathcal{H}}$. In the Boltzmann rational distribution, $\beta \in [0, \infty)$ is the rationality hyperparameter: as β approaches 0, the human is considered to be more irrational; their actions are essentially uniformly distributed. On the other hand, as β increases, the human's actions are increasingly optimal (i.e. "rational"). The robot does not have access to the human's policy; instead, it assumes an apriori model of the human. In continuous spaces, Equation 4 is intractable. Similar to [5], we tractably estimate the human's policy using the principle of maximum entropy: the probability of a goal decreases exponentially as its cost increases. This yields the following approximation:

$$\pi_{\mathcal{H}}\left(a_{\mathcal{H}} \mid s, \theta\right) \propto \exp\left(-\beta \cdot Q(s, a_{\mathcal{H}}, \theta)\right) \tag{5}$$

Firstly, in the absence of communication, we approximate the human's cost-value function as:

$$\mathcal{Q}(s, a_{\mathcal{H}}, \theta) = \text{dist}(a_{\mathcal{H}} + s, \theta) - \text{dist}(s, \theta) + \|a_{\mathcal{H}}\| \tag{6}$$

The first two terms measure the distance by which the human actions move the robot away from the human's goal, while the last term measures the magnitude of the human actions. Formally, Equation 6 is minimized when the human takes *low-effort* actions that minimize the distance between the robot and their goal and require the least effort to do so. However, in the absence of communication, humans cannot directly observe whether or not their actions have influenced the robot's belief to a state where they no longer need to provide input actions and thus, have no reliable basis on which to determine when they can minimize their effort.

In the absence of communication, the human must infer the robot's belief by observing the robot's actions. However, in many cases there can be uncertainty in determining the robot's goal — for example, if the spoon and the fork are

close to one another, how can the human reliably tell which goal the robot is moving towards? On the other hand, in the presence of communication, the human has a reliable prediction of the robot's future assistive actions given its belief and can respond to this communication *positively* by removing input or *negatively* by continuing to work against the robot. Our key insight is that when the robot's belief is communicated, human inputs can be interpreted as assurance or rebuttal of this communicated belief.

Therefore, we propose that the human's internal cost-value function in the presence of communication can be modeled by incorporating the robot's belief into the cost of the human's actions.

$$\mathcal{Q}(s, a_{\mathcal{H}}, \theta) = \text{dist}(a_{\mathcal{H}} + s, \theta) - \text{dist}(s, \theta) + b(\theta) \cdot \|a_{\mathcal{H}}\| \quad (7)$$

In the presence of communication, if the robot's belief is correct, then the human's cost is minimized by providing little effort in agreement with the robot's assistance. If the robot's belief is *incorrect*, then the human will provide inputs that contradict this belief. For example, in the case of ambiguous goals (i.e., the spoon and fork placed close together), with the presence of communication, the human will hold a definite answer for whether the robot is correct. This will result in either further adjustments to correct a misaligned belief or a submission of control seeing that they can minimize their effort by relying on the robot's assistance.

Robot. The robot updates its belief $b(\theta)$ based on the observed human actions. Let $P(\theta \mid s, a_{\mathcal{H}})$ denote the probability that the human is optimizing for the goal θ given the state s and human action $a_{\mathcal{H}}$. Using Bayes' theorem, the posterior probability is defined as:

$$P(\theta \mid s, a_{\mathcal{H}}) \propto P(a_{\mathcal{H}} \mid s, \theta) \cdot P(\theta) \quad (8)$$

Here, $P(\theta)$ is the prior of the robot's belief over the human's goal and $P(a_{\mathcal{H}} \mid s, \theta)$ is the likelihood function for the robot's prediction. Note that $P(a_{\mathcal{H}} \mid s, \theta)$ is equivalent to $\pi^{\star}_{\mathcal{H}}$, which we model as $\pi_{\mathcal{H}}$. Similar to Equation 5, we use the principle of maximum entropy to derive an equivalent form for Equation 8:

$$P(\theta \mid s, a_{\mathcal{H}}) \propto \exp\left(-\beta \cdot Q\left(s, a_{\mathcal{H}}, \theta\right)\right) \cdot P(\theta) \quad (9)$$

The robot takes actions $a_{\mathcal{R}}$ that minimize Equation 6 in the absence of communication and Equation 7 in the presence of communication according to:

$$a_{\mathcal{R}} = \sum_{\theta \in \Theta} P\left(\theta \mid s, a_{\mathcal{H}}\right) \cdot (\theta - s) \quad (10)$$

Since the robot's belief may be incorrect, the robot *blends* the human's commanded action with an assistive action:

$$a_{\mathcal{B}} = (1 - \alpha) \cdot a_{\mathcal{H}} + \alpha \cdot a_{\mathcal{R}} \quad (11)$$

The hyperparameter $\alpha \in [0, 1]$ is determined by a threshold according to the human's action such that when the robot displays the correct belief and the human surrenders control, the robot is allowed to take a higher level of control to assist.

For this, we transition α from a minimum value in the presence of human action to a maximum value when the robot is in full control.

$$\begin{cases} \alpha = \alpha + \textit{step}, \ \alpha \leq \alpha_{\max} & \text{if} \ \ \|a_{\mathcal{H}}\| \approx 0 \\ \alpha = \alpha - \textit{step}, \ \alpha \geq \alpha_{\min} & \text{if} \ \ \|a_{\mathcal{H}}\| \not\approx 0 \end{cases} \quad (12)$$

Here *step* is a hyperparameter chosen by the designer to control the rate at which the robot will increase its assistance proportionally to the number of timesteps that the user has allowed for complete robot assistance.

Altogether, Equations 7-12 form our method for selecting optimal actions in the presence of communication. These equations build upon existing shared autonomy approaches for inferring the human's goal and providing assistance [4]–[6]. But we have modified this existing learning framework to explicitly account for communication and the effect communication may have on the human's internal cost function Q^*. Without communication, earlier works such as [4] suggest that the human optimizes for their error and effort as shown in Equation 6. However, with communication, our experimental results from Section III-B indicate that humans are willing to increase their effort if the robot is wrong and

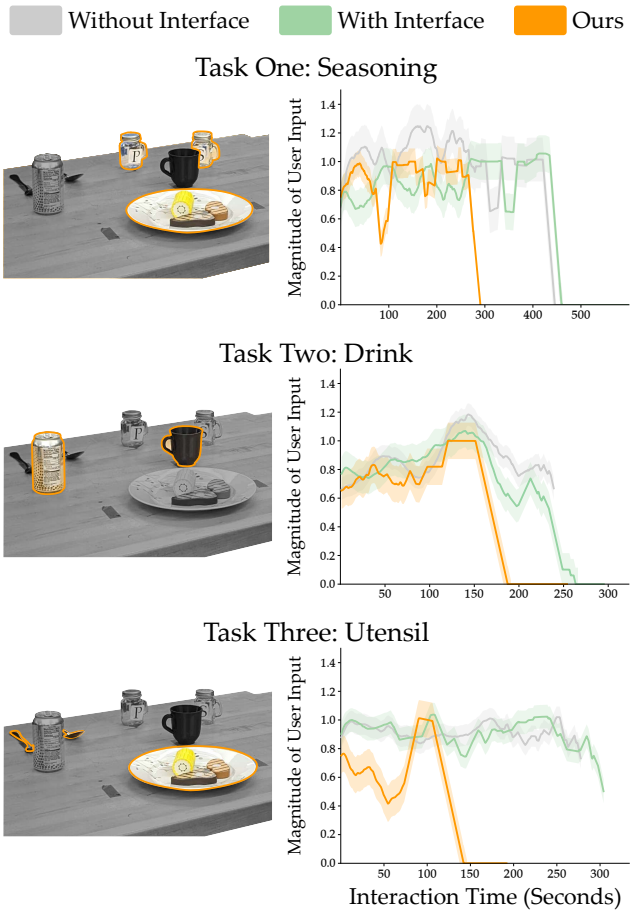

Fig. 3. Tasks and user inputs from the user study in Section V. (Left) The specific items that the human brought the robot to interact with for each task. (Right) The magnitude of the human's inputs over time averaged across all users. These results show that users completed the tasks more quickly with **Ours**, and overall needed fewer inputs to convey their intended goals to the robot.

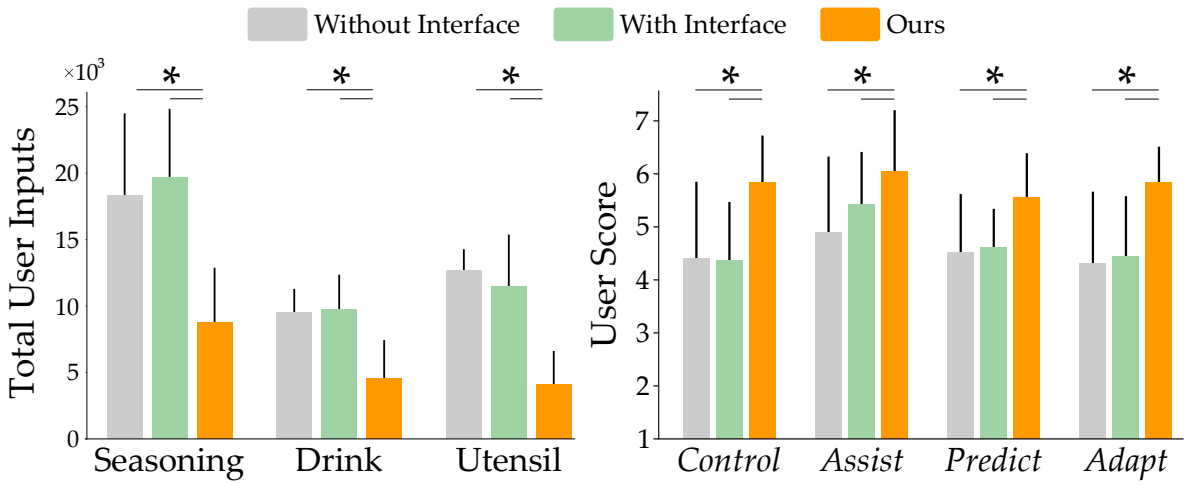

Fig. 4. Objective and subjective results from the user study in Section V. (Left) Total user inputs for **Seasoning**, **Drink**, and **Utensil** tasks. To count the number of inputs, the robot measured whether the human had pressed the joystick every 0.02 seconds. Across each task, users provided fewer inputs and relied on the robot's assistance more when using **Ours** ($p < 0.001$, $p < 0.001$, $p < 0.001$). These results support **H3**: Users spent less effort when using **Ours**. (Right) Subjective results for the three baselines. Across the four Likert-Scale items, users preferred **Our** method: they felt that they could easily *control* the system ($p < 0.001$), the robot provided effective *assistance* ($p < 0.005$), the robot better *predicted* their goal ($p < 0.001$), and the robot *adapted* more quickly to their actions ($p < 0.001$).

release control when the robot is correct. Using these findings we update our human model for Q in Equation 7. Up to this point, our modified learning rule is informed by experiments but has not yet been tested. Accordingly, in Section V we will compare our proposed method for aligning learning with communication against baselines that separately learn and communicate.

V. Testing the Combination of Learning and Communication

Lastly, we conduct an in-person user study to evaluate the performance of our proposed method in comparison to the state-of-the-art shared autonomy baseline [4] with and without a communicative interface. We wish to demonstrate that accounting for the knowledge of the robot's belief in the human's cost function, in addition to communicating the robot's belief, allows the robot to provide better assistance than simply communicating the robot's belief with the baseline shared autonomy approach.

Experimental Setup. Users were instructed to complete three tasks in a more complicated environment than the first in-person user study to highlight the utility of this approach:

1) **Seasoning:** Retrieve a salt or pepper shaker, bring it to a plate of food, and then return it to its base position.
2) **Drink:** Go to the can of soda, bring the can of soda to a mug, and return the can to its base position.
3) **Utensil:** Retrieve the spoon or fork and bring it to the relevant side of the plate.

Participants commanded the robot using a Logitech F310 gamepad to complete each of the three tasks using one of three methods: **Without Interface**, **With Interface**, and **Ours**. The order in which participants interacted with these methods was randomized to avoid any proficiency bias. Details of these tasks are illustrated in Figure 3 (Left).

Independent Variables. In each task, the robot starts with a uniform prior over the goals which is gradually updated according to the methods discussed in section IV. Participants performed each task three times — using the baseline shared autonomy approach **Without Interface**, using the same baseline **With Interface**, and using our method of feedback-enabled shared autonomy - **Ours** (which combines learning with communication).

Dependent Variables. We recorded the *Total User Inputs* to measure the amount of effort spent by the users in completing each task. We also recorded subjective *User Scores* through a 7-point Likert scale survey with four items — for how easy it was to *Control* the robot, how often they could tell when the robot *Assisted* them, whether the robot was able to *Predict* their goals, and if the robot *Adapted* to their actions.

Participants. A total of 12 participants from the Virginia Tech community took part in this study (2 female, ages 28.5 ± 6.5 years). Two of the twelve users had not interacted with robots before. Users provided written consent as per university guidelines (IRB #20-755).

Hypothesis. We hypothesized that for this study:

> **H3.** *The human will spend less effort in completing the tasks when using* Our *method.*
>
> **H4.** *Users will provide higher scores on the subjective metrics for* Our *method than the baselines.*

Results. The results of our user study are summarized in Figure 4. To address **H3**, we measured the number of user inputs across three separate tasks for each method. Here, a lower score is better where fewer inputs imply that the user is exhibiting less effort when completing the task. Paired-sample T-tests showed that participants used significantly fewer inputs when the robot used **Ours** for each task ($t(11) = 4.106, p < 0.001$, $t(11) = 5.806, p < 0.001$,

$t(11) = 9.636, p < 0.001$). Figure 3 shows the average magnitude of the user input over time for each task; these results further support **H3**. It was noted that given the apparent improvement from the addition of our algorithm, given how our study varied the order of appearance for each method, the users noted a lower user preference for the base interface option. It was noted that this was doubly apparent given the more complex environment given how, even with the interface, it didn't change the robot's struggle with the base environment. A struggle that has been noted across multiple different bodies of work in Shared Autonomy [2], [4].

Regarding **H4**, we present the subjective results from our Likert-scale survey in Figure 4 (right). A one-way ANOVA analysis of the users' responses showed a significant difference in the perceived *Control*, *Assistance*, *Prediction*, and *Adaptation* that the robot exhibited when using our method ($F(69) = 11.901, p < 0.001$, $F(69) = 6.368, p < 0.005$, $F(69) = 8.794, p < 0.001$, $F(69) = 13.345, p < 0.001$). Actions chosen by **Ours** were preferable to those selected by baselines; this supports **H4**.

VI. Conclusion

In this paper, we explored the effects of communicating learned assistance back to the human operator in shared autonomy. While previous research has focused on learning the human's task and providing assistance, we instead focused on harnessing the effect of the communication. We hypothesized that humans will interact with shared autonomy systems differently when those systems communicate their learning back to the human. Using the results from online and in-person user studies, we showed that humans are more likely to intervene when the robot incorrectly predicts their intent, and release control when the robot correctly understands their task. We used the insights from these results to modify the robot's learning algorithm: under our proposed approach, the robot adjusts its model of the human's cost function to account for how communication changes the human's input patterns. Finally, we compared our approach for combining learning and communication against shared autonomy baselines that separately handle learning or communication. In a user study with 12 in-person participants across three kitchen tasks, we found that our proposed approach for combining learning and communication increased the subjective and objective performance of the human-robot team.

References

[1] D. M. Taylor, "Americans with disabilities: 2014," *US Census Bureau*, pp. 1–32, 2018.

[2] B. D. Argall, "Autonomy in rehabilitation robotics: An intersection," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 441–463, 2018.

[3] T. Bhattacharjee, E. K. Gordon, R. Scalise, M. E. Cabrera, A. Caspi, M. Cakmak, and S. S. Srinivasa, "Is more autonomy always better? Exploring preferences of users with mobility impairments in robot-assisted feeding," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2020, pp. 181–190.

[4] S. Javdani, H. Admoni, S. Pellegrinelli, S. S. Srinivasa, and J. A. Bagnell, "Shared autonomy via hindsight optimization for teleoperation and teaming," *The International Journal of Robotics Research*, vol. 37, no. 7, pp. 717–742, 2018.

[5] S. Jain and B. Argall, "Probabilistic human intent recognition for shared autonomy in assistive robotics," *ACM Transactions on Human-Robot Interaction*, vol. 9, no. 1, pp. 1–23, 2019.

[6] A. D. Dragan and S. S. Srinivasa, "A policy-blending formalism for shared control," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 790–805, 2013.

[7] R. M. Aronson, T. Santini, T. C. Kübler, E. Kasneci, S. Srinivasa, and H. Admoni, "Eye-hand behavior in human-robot shared manipulation," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 4–13.

[8] C. Brooks and D. Szafir, "Balanced information gathering and goal-oriented actions in shared autonomy," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2019, pp. 85–94.

[9] M. Zolotas and Y. Demiris, "Towards explainable shared control using augmented reality," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019, pp. 3020–3026.

[10] J. F. Mullen, J. Mosier, S. Chakrabarti, A. Chen, T. White, and D. P. Losey, "Communicating inferred goals with passive augmented reality and active haptic feedback," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 8522–8529, 2021.

[11] C. Brooks and D. Szafir, "Visualization of intended assistance for acceptance of shared control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 11 425–11 430.

[12] D. P. Losey, C. G. McDonald, E. Battaglia, and M. K. O'Malley, "A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction," *Applied Mechanics Reviews*, vol. 70, no. 1, p. 010804, 2018.

[13] M. Fontaine and S. Nikolaidis, "A quality diversity approach to automatically generating human-robot interaction scenarios in shared autonomy," in *Robotics: Science and Systems*, 2020.

[14] A. Broad, I. Abraham, T. Murphey, and B. Argall, "Data-driven Koopman operators for model-based shared control of human–machine systems," *The International Journal of Robotics Research*, vol. 39, no. 9, pp. 1178–1195, 2020.

[15] S. Reddy, A. D. Dragan, and S. Levine, "Shared autonomy via deep reinforcement learning," in *Robotics: Science and Systems*, 2018.

[16] C. Schaff and M. R. Walter, "Residual policy learning for shared autonomy," in *Robotics: Science and Systems*, 2020.

[17] M. Hagenow, E. Senft, R. Radwin, M. Gleicher, B. Mutlu, and M. Zinn, "Corrective shared autonomy for addressing task variability," *IEEE Robotics and Automation Letters*, vol. 6, pp. 3720–3727, 2021.

[18] M. Zurek, A. Bobu, D. S. Brown, and A. D. Dragan, "Situational confidence assistance for lifelong shared autonomy," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 2783–2789.

[19] Y. Cui, S. Karamcheti, R. Palleti, N. Shivakumar, P. Liang, and D. Sadigh, "No, to the right: Online language corrections for robotic manipulation via shared autonomy," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2023, pp. 93–101.

[20] J. Z.-Y. He, Z. Erickson, D. S. Brown, A. Raghunathan, and A. Dragan, "Learning representations that enable generalization in assistive tasks," in *Conference on Robot Learning*, 2023, pp. 2105–2114.

[21] D. P. Losey, H. J. Jeon, M. Li, K. Srinivasan, A. Mandlekar, A. Garg, J. Bohg, and D. Sadigh, "Learning latent actions to control assistive robots," *Autonomous Robots*, vol. 46, no. 1, pp. 115–147, 2022.

[22] A. Jonnavittula, S. A. Mehta, and D. P. Losey, "SARI: Shared autonomy across repeated interaction," *ACM Transactions on Human-Robot Interaction*, 2024.

[23] E. Rosen, D. Whitney, E. Phillips, G. Chien, J. Tompkin, G. Konidaris, and S. Tellex, "Communicating and controlling robot arm motion intent through mixed-reality head-mounted displays," *The International Journal of Robotics Research*, pp. 1513–1526, 2019.

[24] S. Habibian, A. A. Valdivia, L. H. Blumenschein, and D. P. Losey, "A review of communicating robot learning during human-robot interaction," *arXiv preprint arXiv:2312.00948*, 2023.

[25] V. Alonso and P. De La Puente, "System transparency in shared autonomy: A mini review," *Frontiers in Neurorobotics*, 2018.

[26] A. Jonnavittula and D. P. Losey, "Communicating robot conventions through shared autonomy," in *IEEE International Conference on Robotics and Automation*, 2022, pp. 7423–7429.