

Communicating Robot Conventions through Shared Autonomy

Ananth Jonnavittula and Dylan P. Losey

Abstract—When humans control robot arms these robots often need to infer the human’s desired task. Prior research on assistive teleoperation and shared autonomy explores how robots can determine the desired task based on the human’s joystick inputs. In order to perform this inference the robot relies on an internal mapping between joystick inputs and discrete tasks: e.g., pressing the joystick left indicates that the human wants a plate, while pressing the joystick right indicates a cup. This approach works well *after* the human understands how the robot interprets their inputs — but inexperienced users still have to learn these mappings through trial and error! Here we recognize that the robot’s mapping between tasks and inputs is a *convention*. There are multiple, equally efficient conventions that the robot could use: rather than passively waiting for the human, we introduce a shared autonomy approach where the robot *actively reveals* its chosen convention. Across repeated interactions the robot intervenes and exaggerates the arm’s motion to demonstrate more efficient inputs while also assisting for the current task. We compare this approach to a state-of-the-art baseline — where users must identify the convention by themselves — as well as written instructions. Our user study results indicate that modifying the robot’s behavior to reveal its convention outperforms the baselines and reduces the amount of time that humans spend controlling the robot. See videos of our user study here: <https://youtu.be/jROTVOp469I>

I. INTRODUCTION

Imagine teleoperating an assistive robot arm to reach for a notepad in a cluttered environment (like the one in Figure 1). You interact with a joystick, and the robot leverages these joystick inputs to infer which task you are trying to complete. Intuitively, you might press the joystick *straight* towards your desired object. Although this input makes sense to you, it could confuse the robot arm: there is a marker right next to the notepad, and the robot is not sure which of these two objects you really wanted to reach.

To interpret human inputs and predict their desired task robots use an inference or intent detection algorithm. At its heart, this inference algorithm is based on a robot-assigned *convention*: i.e., a mapping between high-level tasks and low-level inputs. For example, the robot could assume that the human will directly aim their joystick towards their desired goal. We refer to this model as a convention because there are *multiple, equally optimal* mappings from tasks to inputs. Returning to our example, another convention could be moving the joystick *up* to indicate the notepad, and *down* to indicate the marker — pressing up (or down) is no more challenging than aiming straight for the notepad.

Understanding the robot’s convention is key to communicating with the robot. Once the human knows how the

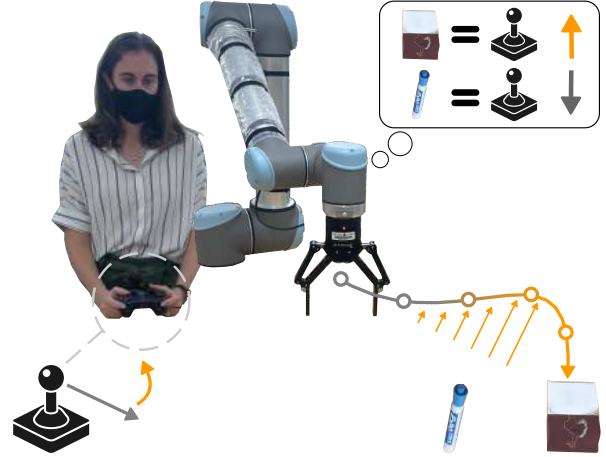


Fig. 1. Human interacting with a joystick to convey their desired task. The human initially presses their joystick straight towards the notepad. But the robot has in mind a convention for interpreting the human’s inputs — *up* for the notepad and *down* for the marker. To communicate this convention, we leverage shared autonomy to modify the robot’s trajectory and reveal more informative motions. Humans that gradually adapt their joystick inputs to match these motions will more concisely and accurately convey their intent.

robot interprets their inputs, the human can then follow this convention to seamlessly convey their intent. We specifically focus on *shared autonomy* settings — here the standard convention is for the human to move the robot arm directly towards their task [1]–[7]. At first glance, this makes inference challenging when the potential tasks are close together [8]: a small mistake could point the joystick towards the marker instead of the notepad. But this confusion is easily avoidable once the user understands the robot’s convention: the human can reliably convey their task by exploiting the convention and aiming to the right of the notepad, clearly avoiding the marker. This choice of action that unambiguously indicates the human’s task is an *exaggeration*.

In this paper we explore how robots can convey their built-in conventions to inexperienced human users. We recognize that — because the robot knows its own conventions — it also knows how humans *should* interact with these conventions to seamlessly communicate their task. Specifically in the context of shared autonomy, our insight is that:

Robots can reveal their conventions by guiding the human’s behavior towards more communicative inputs.

Robots that apply our insight leverage shared autonomy to actively demonstrate their conventions to the human. This process is shown in Figure 1: as the human moves directly towards the notepad, the robot intervenes to (a) help the human complete the task while simultaneously (b) guiding the human along a trajectory that would more clearly indicate the notepad. Our central hypothesis is that the human will

learn from this guidance: the next time they encounter this scenario, the human should update their joystick inputs to mimic the demonstrated behavior. If successful, our approach reduces the amount of time users spend interacting with the joystick to convey their desired task.

Overall we make the following contributions:

Formalizing Conventions in Shared Autonomy. We formulate the role of conventions in inferring the human’s desired task. We then enable robots to exploit these conventions and identify the exaggerated policy that the human *should* follow to indicate their task with fewer joystick inputs.

Communicating Conventions over Repeated Interaction. We leverage shared autonomy to guide users towards more communicative policies. Across repeated interactions our approach attempts to infer the human’s task and then suggests an improved way to indicate that task. We prove that — if the human mimics the robot — our approach is more efficient than letting the human find the convention on their own.

Comparison to Written Instructions. We test our resulting algorithm in scenarios where the human is reaching for objects on a cluttered table or performing continuous skills. We compare with other teaching modalities, including written, crowd-sourced descriptions of the conventions. Our results suggest that using shared autonomy to demonstrate conventions outperforms the alternatives, particularly when the conventions are complex or unintuitive.

II. RELATED WORK

Shared Autonomy. Over 13% of all American adults have some form of physical disability and need assistance during activities of daily living [9]. Robot arms can help these adults perform everyday tasks without relying on caregivers [10]. Rather than forcing the human to constantly teleoperate the robot arm, it is often beneficial to automate parts of the task [11]–[13]. Shared autonomy arbitrates between the human operator and autonomous assistance so that both agents control the robot’s motion. We specifically focus on shared autonomy paradigms where the robot is given the discrete set of candidate tasks *a priori*: during interaction, the robot tries to infer the human’s current task and then takes over to autonomously complete the motion [1]–[4]. Recent work on shared autonomy has focused on how the robot can gather information and learn new tasks from human interactions [5]–[7]. By contrast, our work explores how robots can *convey information* through shared autonomy in order to improve the human’s interactions.

Conventions. When there are multiple optimal solutions to a multi-agent problem, conventions determine which solution the team follows [14]–[16]. As an example, we follow a convention to drive on the right (or the left) side of the road. Within this paper, conventions refer to the mapping between the human’s desired task and their joystick inputs. Because the task space is discrete and the input space is continuous, there are an *infinite* number of possible conventions: returning to Figure 1, the human could press the joystick up, down, left, right, or any other angle to indicate the notepad. But

for this convention to function, the human must understand what the robot expects. Recent works have explored how robots can learn human conventions [15], how robots can avoid conventions [17], and how robots should respond to humans who know the robot’s convention [18]. Rather than causing the robot to adapt to the human, we study how the robot can drive *human adaptation* to the robot’s convention.

Algorithmic Teaching. To determine how the robot should teach conventions to the human we utilize recent work on algorithmic teaching (also called machine teaching) [19]. Within contexts where the robot is learning from demonstrations, prior work improves human teaching by providing *verbal or written teaching guidance* — i.e., heuristic instructions on how to teach [20]–[22]. However, within our shared autonomy context we propose to leverage the *robot’s motion* to teach conventions. Our approach draws from related research on legible motions for human-robot collaboration: here the robot purposely follows an exaggerated trajectory to make its intent clear to onlooking humans [23], [24]. We take inspiration from legible motions and machine teaching to establish a method by which the robot communicates its convention in shared autonomy settings.

Most related to our research are [25] and [26]. Both papers focus on shared autonomy, and use robot interventions to change or improve the human’s behavior. Like these papers, we will leverage shared autonomy to convey information from the robot to the human — but unlike these works, our goal is to teach the human how to teleoperate the robot.

III. PROBLEM STATEMENT

Let us return to our motivating example where a human is trying to teleoperate the robot arm to pick up a notepad. When the human starts interacting with the robot they know their task (i.e., the human knows that they want the notepad), but they do not know the robot’s convention, and therefore they do not know the most *efficient* way to communicate to the robot that they want the notepad. Conversely, the robot knows its convention — and the most efficient way for the human to indicate either task. However, the robot does not know whether the human wants the notepad or the marker.

Formalism. Accordingly, we are faced with an *asymmetry of information*. The human has a task (that the robot has to figure out) and the robot selects the convention (which the human cannot observe). We formulate this scenario as a cooperative inverse reinforcement learning (CIRL) game [16], i.e., a two-player Markov game where both human and robot receive the same reward: $\mathcal{M} = \langle \mathcal{S}, \{\mathcal{A}_H, \mathcal{A}_R\}, T, \{\Theta, R\}, \gamma \rangle$. Here $s \in \mathcal{S} \subseteq \mathbb{R}^n$ is the robot’s joint position, $a_H \in \mathcal{A}_H \subseteq \mathbb{R}^n$ is the human’s commanded velocity¹, and $a_R \in \mathcal{A}_R \subseteq \mathbb{R}^n$ is the robot’s assistance. Within shared autonomy settings the overall action is a combination of the human’s commanded action a_H and the robot’s assistance a_R , so that

¹The human pushes the joystick, and this input is mapped to a joint velocity. Our approach is not tied to any specific mapping: e.g., pressing the joystick right to move the robot’s end-effector along the x -axis [27].

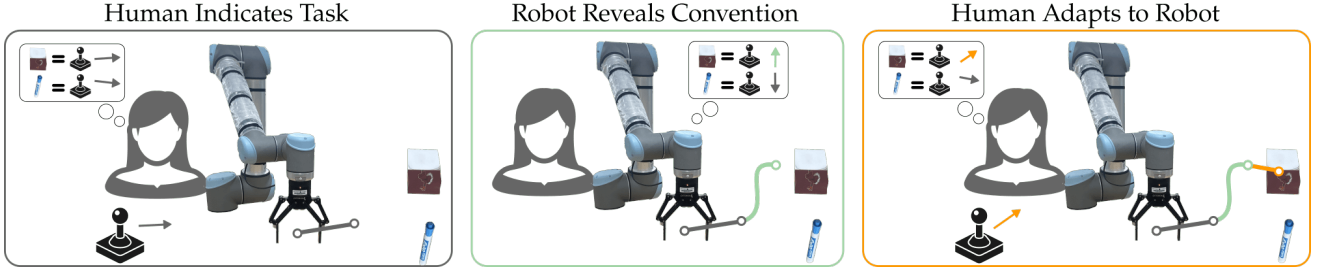


Fig. 2. Our proposed approach for communicating conventions through shared autonomy. (Left) At the start of each interaction the human uses their current understanding of the convention to try and indicate their desired task. (Middle) The robot infers which task is most likely and then intervenes to reveal its convention for that task. Here the robot moves up because pressing up on the joystick would increase its confidence in the notepad. We constrain the robot’s motion so that it still assists the human to complete the task. (Right) Our hypothesis is that the human will respond to the robot’s intervention by mimicking the robot’s behavior. Within this example, the human learns to press their joystick up to indicate that they want to reach the notepad.

the dynamics $T(s, a_{\mathcal{H}}, a_{\mathcal{R}})$ become:

$$s^{t+1} = s^t + \Delta t \cdot f(a_{\mathcal{H}}^t, a_{\mathcal{R}}^t) \quad (1)$$

One common instance of Equation (1) is linearly blending the human and robot actions [1], [2], [27]:

$$f(a_{\mathcal{H}}^t, a_{\mathcal{R}}^t) = \beta \cdot a_{\mathcal{H}}^t + (1 - \beta) \cdot a_{\mathcal{R}}^t \quad (2)$$

where $0 \leq \beta \leq 1$ arbitrates between the human and robot.

Continuing our CIRL formalism, Θ is the discrete set of candidate tasks that the human might want to complete (i.e., reaching for the notepad, picking up the marker, or opening a drawer), and $\theta \in \Theta$ is the human’s current task (which the robot is trying to infer). The reward function $R = \mathcal{S} \times \Theta \rightarrow \mathbb{R}$ depends on the human’s current task. To give an example from Figure 1, we could specify that $R(s, \theta) = 0$ when the robot reaches the notepad and $R(s, \theta) = -1$ at all other states. The scalar $\gamma \in [0, 1)$ in \mathcal{M} is a discount factor.

Conventions. Solving the CIRL game produces a pair of human and robot policies $(\pi_{\mathcal{H}}, \pi_{\mathcal{R}})$ that communicate the human’s task and maximize long-term reward [16]. However, choosing this policy pair becomes challenging when there are multiple, equally efficient ways to communicate the human’s goal [14]. Let $\pi_{\mathcal{H}}(a_{\mathcal{H}} | s, \theta)$ be the human’s policy: given the robot state s and the human’s task θ , this policy determines which joystick input (i.e., which commanded action $a_{\mathcal{H}}$) the human will provide. We focus on scenarios where there are a set of equally optimal human policies $\pi_{\mathcal{H}} \in \Pi_{\mathcal{H}}$. More formally, we consider settings where there are multiple solutions to the CIRL game. The choice of which policy (i.e., which solution) to use determines the team’s *convention*.

Consistent with prior work on shared autonomy, here the robot picks the convention [1]–[4]. In practice, this means that the robot selects some $\pi_{\mathcal{H}} \in \Pi_{\mathcal{H}}$ that it expects the human to follow, and the robot infers the task θ based on this model. Returning to our motivating example from Figure 1, the convention could be $\pi_{\mathcal{H}}(a_{\mathcal{H}} = up | s, \theta = notepad) = 1$. We reiterate that there are multiple optimal conventions, i.e., there are multiple policies that convey the human’s goal while minimizing the human’s effort. For instance, flipping the convention so that down indicates the notepad²: $\pi_{\mathcal{H}}(a_{\mathcal{H}} = down | s, \theta = notepad) = 1$.

²Although our examples involve only a single joystick input, conventions in complex environments may require a sequence of human commands.

Inference. The purpose of establishing a convention is to enable the robot to infer the human’s task θ . Recall that the human has a specific task that they want to accomplish, and the robot needs to infer that task. We denote the robot’s belief over the discrete set of candidate tasks as:

$$b^{t+1}(\theta) = P(\theta | (s^0, a_{\mathcal{H}}^0), (s^1, a_{\mathcal{H}}^1), \dots, (s^t, a_{\mathcal{H}}^t)) \quad (3)$$

This belief captures the likelihood of task $\theta \in \Theta$ given the history of robot states and human actions. Following [28], [29], we assume that the human’s inputs $a_{\mathcal{H}}$ are conditionally independent given s and θ . Applying Bayes’ rule:

$$b^{t+1}(\theta) \propto \pi_{\mathcal{H}}(a_{\mathcal{H}} | s, \theta) \cdot b^t(\theta) \quad (4)$$

Hence, the robot’s convention $\pi_{\mathcal{H}}$ (i.e., the robot’s chosen model of the human’s policy) determines how the robot interprets human inputs and infers θ . Within shared autonomy two common instantiations of $\pi_{\mathcal{H}}$ are the Boltzmann rational model [29] and the cosine similarity between the human’s commanded action and the optimal action for a given task. Both of these conventions expect the human to point their joystick directly towards their target [1]–[4], [6], [7], [12].

Robot. We want to develop an approach that works across arbitrary conventions. Hence, we leave the robot’s convention $\pi_{\mathcal{H}}$ as a general human model that maps between high-level tasks and low-level joystick inputs. Recalling that there are multiple solutions $(\pi_{\mathcal{H}}, \pi_{\mathcal{R}})$ to our CIRL game, the robot now executes the policy $\pi_{\mathcal{R}}$ that pairs with convention $\pi_{\mathcal{H}}$. Returning to our running example, let $\pi_{\mathcal{H}}(a_{\mathcal{H}} = up | s, \theta = notepad) = 1$. Accordingly, if the human presses their joystick up, the robot’s correct response is to autonomously guide its arm to the notepad. But for the robot to provide the right assistance, it must first understand what the human wants — and to do this, the robot must teach the human to follow its chosen convention.

IV. REVEALING ROBOT CONVENTIONS

Our proposed approach for revealing the robot’s convention is based on *shifting* the human’s behavior across repeated interactions (see Figure 2). The first time that they interact with the robot, the human leverages their *own convention* to communicate their task (e.g., pressing the joystick directly towards the notepad). We want to shift this input over time so that the human gradually understands and effectively

leverages the *robot's convention*. In this section we introduce a constrained optimization approach to **generate** actions that reveal the robot's convention. We then explore the conditions the human must satisfy to **adopt** this convention, and prove that demonstrating the convention is more efficient than waiting for the human to learn by themselves.

A. Generating Revealing and Assistive Actions

To reveal the robot's convention we modify the motion of the robot arm so that — if the human provides joystick inputs that match the demonstrated motion — the human will follow the robot's convention. Recall that b is the robot's belief over the discrete set of candidate tasks Θ , and let $\theta^* = \max_{\theta \in \Theta} b^t(\theta)$ be the human's most likely task at the current timestep t . Here we optimize for actions that reveal θ^* . Put another way, we seek the commanded human action that will most effectively increase the robot's confidence in θ^* . Within our formalism this action maximizes $b^{t+1}(\theta^*)$, the robot's belief in task θ^* at the next timestep.

So far we are describing a straightforward optimization. However, this is made more challenging by our shared autonomy setting. On the one hand, the robot should demonstrate informative actions to the human; on the other hand, the robot needs to assist the human and help them to correctly complete their task. We therefore *constrain* the robot's action to ensure that it still assists the human. Our resulting approach for generating revealing and assistive actions is:

$$\begin{aligned} a_{\mathcal{R}} = \arg \max_{a \in \mathcal{A}_{\mathcal{R}}} & \quad b^{t+1}(\theta^*) \\ \text{s.t.} & \quad V_{\theta^*}(s) - Q_{\theta^*}(s, a) \leq \epsilon \end{aligned} \quad (5)$$

Once we substitute in Equation (4) and simplify, we reach:

$$\begin{aligned} a_{\mathcal{R}} = \arg \max_{a \in \mathcal{A}_{\mathcal{R}}} & \quad \frac{\pi_{\mathcal{H}}(a \mid s, \theta^*)}{\sum_{\theta \in \Theta} \pi_{\mathcal{H}}(a \mid s, \theta)} \\ \text{s.t.} & \quad V_{\theta^*}(s) - Q_{\theta^*}(s, a) \leq \epsilon \end{aligned} \quad (6)$$

Here $\pi_{\mathcal{H}}$ is the convention that the robot wants to reveal to the human: the robot leverages this convention to identify actions that maximize the belief in θ^* . We define $Q_{\theta^*}(s, a)$ as the cumulative reward the robot will receive by taking action a in state s , and then optimally completing task θ^* afterwards (with no human assistance) [3]. Finally, $V_{\theta^*}(s) = \max_a Q_{\theta^*}(s, a)$ is the maximum expected reward the robot can achieve if it completes task θ^* autonomously.

When selecting the hyperparameter $\epsilon \geq 0$ in Equation (6) the designer chooses how much deviation from the optimal policy is allowable. If $\epsilon = 0$, the robot always takes assistive actions (and never reveals information to the human). By contrast, as $\epsilon \rightarrow \infty$ the robot only shows revealing actions, and does not consider whether these actions help the human complete the task. We note that our constrained optimization approach here is similar to [30]: but unlike [30], we explicitly encode task performance as a constraint.

Algorithm. Our overall approach is displayed in Figure 2 and Algorithm 1. At each timestep the robot infers the human's most likely task (based only on the human's inputs) and then solves for an assistive action $a_{\mathcal{R}}$ that reveals the robot's

Algorithm 1 Communicating Robot Conventions

- 1: **Input:** Robot's chosen human convention $\pi_{\mathcal{H}}$, discrete set of tasks Θ , and designer-specified hyperparameter ϵ
 - 2: **Precompute:** The Q -function for each task $\theta \in \Theta$
 - 3: **while** task not completed **do**
 - 4: Observe human command $a_{\mathcal{H}}^t$ and state s^t
 - 5: $\theta^* \leftarrow \max_{\theta \in \Theta} b^t(\theta)$
 - 6: $a_{\mathcal{R}}^t \leftarrow$ solution to Equation (6)
 - 7: $s^{t+1} \leftarrow s^t + \Delta t \cdot f(a_{\mathcal{H}}^t, a_{\mathcal{R}}^t)$
 - 8: **end while**
-

convention for that task. Finally, the robot blends the human and assistive actions and transitions to a new state. If the designer wants to make the robot more or less revealing, $\epsilon(t)$ can be changed based on human performance.

B. Driving Adaptation to Robot Conventions

Our approach reveals the robot's conventions to the human — but is this any more efficient than letting the human find these conventions for themselves? To answer this question we take the human's perspective, and write the setting as a *multi-arm bandit*. The human is interacting with a joystick, and can press this joystick in N different directions: these N discrete inputs become the arms of our bandit. If the human matches the robot's convention and pulls the correct arm (i.e., pushing *up* to indicate the notepad), the human is rewarded by the robot performing the task autonomously. Otherwise, the human has to continually intervene and correct the robot's motion, resulting in more human effort. Define $REG(k)$ as the number of *incorrect* joystick inputs up to interaction k .

Without Revealing Actions. If the robot does not actively reveal its convention then the human must explore the space of joystick inputs to find the most effective actions. No matter which policy the human uses to explore these inputs, prior work on multi-arm bandits [31], [32] has shown that — in expectation — the number of incorrect joystick inputs is at least logarithmic in time: $\mathbb{E}[REG(k)] \geq \Omega(\log k)$.

With Revealing Actions. Our approach to revealing conventions has the potential to lower this bound. But to be effective, the human must actively *explore* different joystick inputs and *learn* from the robot's response. Specifically, we assume that (a) the initial probability of each joystick input is nonzero and (b) the human update their inputs to match the robot's motion. Given these assumptions, the number of incorrect joystick inputs is constant in time: $\mathbb{E}[REG(k)] = C$. This result follows from Proposition 1 in [33] where the roles of the human and robot are reversed: the robot reveals the informative actions for a given task θ after a finite number of interactions, and thereafter the human mimics the robot's demonstrated convention for that task.

V. USER STUDY

Our analysis suggests that demonstrating the robot's convention will help users find and follow that convention more rapidly. But now we need to show that our approach works

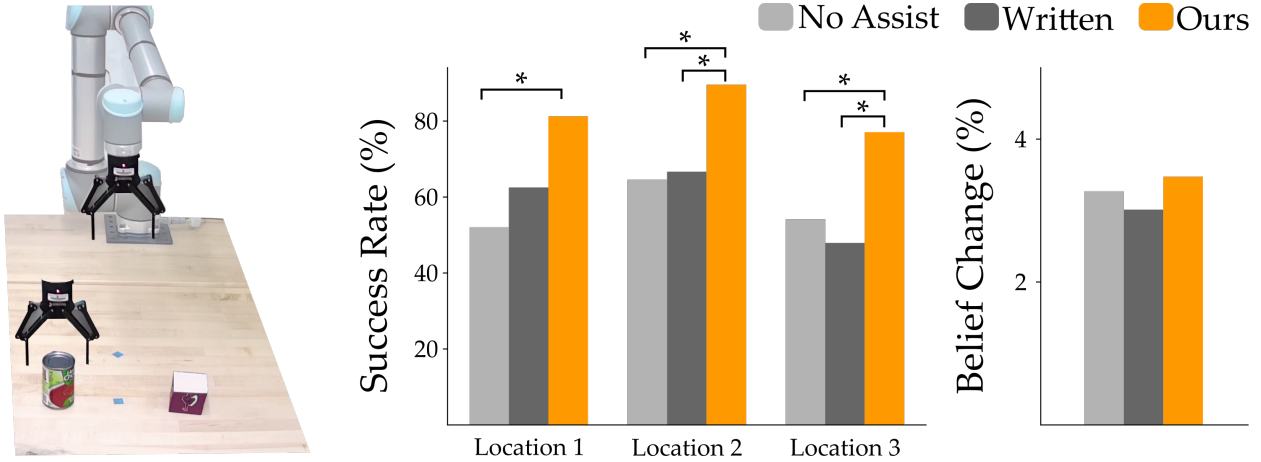


Fig. 3. Experimental setup and results for the first and second parts of our user study. (Left) Participants controlled the robot arm to reach for a soup can or notepad. The robot used a Boltzmann rational model to interpret the participants’ joystick inputs: under this convention, exaggerated motions more efficiently indicated the desired goal (e.g., pointing the joystick left for the soup can). We explored whether humans adapted to this convention over multiple interactions. (Center) Across three sets of object locations, participants who interacted with **Ours** were able to convey their desired goal more accurately and concisely. Here * denotes statistical significance ($p < .05$). (Right) In the second part of our user study we measured how teleoperation behavior changed before and after being exposed to **No Assist**, **Written**, or **Ours**. Participants did improve but the change was minimal. Because we collected the participants’ updated behavior in a setting where the robot did not provide assistance, it is possible that users did not feel the need to exaggerate.

in practice. Accordingly, we conducted an in-person user study where participants teleoperated a 6-DoF robot arm (Universal Robots UR10) to reach for objects and perform skills. The study was divided into *three parts* to explore how conventions impact human effort.

Independent Variables. Over the course of the study participants learned about the robot’s conventions: to prevent this from affecting our results, we used a *between subjects* design. Each participant only interacted with one of the following methods: **No Assist**, **Written**, and **Ours** (Algorithm 1).

In **No Assist** the robot used an existing shared autonomy approach to identify the human’s task [1]³. This robot never attempted to show its convention to the human: participants using **No Assist** had to learn the convention through trial and error based on whether the robot assisted for their desired task. In **Written** the robot also assisted the human without any exaggerations. However, here users were given written, crowd-sourced descriptions of the robot’s conventions. These descriptions were obtained from 15 Amazon Mechanical Turk workers with over a 99% HIT approval rating. As an example, one description told participants to point their joystick “left and down” to indicate the notepad. Finally, in **Ours** the robot revealed its convention by guiding the human towards more informative inputs.

Experimental Setup. Our user study was divided into the three parts that are described below. Since we followed a between subjects design, each participant completed every part with only one method.

In the first part users were tasked with learning the robot’s convention while reaching for either a soup can or a notepad (see Figure 3). The robot followed a Boltzmann rational convention [1]–[3]. The hyperparameter $\epsilon(t)$ was set to 0.04 at the start of the task and $\epsilon(t) \rightarrow 0$ as the robot got closer to

the goal. Although users could indicate their task by pointing the joystick directly towards their target, exaggerated inputs conveyed the human’s task more efficiently (i.e., the human needed fewer joystick inputs to indicate their task if they exaggerated). We tested three object locations: initially the soup can and notepad were located far apart, and then were gradually moved closer together. Participants interacted with each location three times while using their joystick to try and convey the intended goal.

The first part of our user study focused on adapting to a robot convention: in the second part, we tested whether humans would generalize that convention to new scenarios. The soup can and notepad were placed in a new, previously unseen location, and participants teleoperated the robot without any assistance. We measured how the human’s teleoperation behavior changed *before* and *after* the first part of our user study. If participants understood the convention that they had experienced in part one, we expected them to provide more informative teleoperation inputs in part two. Regardless of the method used by the participants, the robot only observed the human’s input actions and did not apply its policy blending algorithm $f(a_H, a_R)$ to assist.

Finally, in the third part of our user study we removed the Boltzmann rational model and introduced a less intuitive convention. The robot was able to perform a continuous skill (open a drawer) or reach for an object (soup can). To indicate the drawer, the human needed to input small right and left motions; to indicate the soup can, the human provided larger right and left motions. We tested this unintuitive convention to explore how each method performs when the mapping from joysticks to tasks is more complex. As in part one, participants interacted with the robot three times.

Dependent Measures. For the first part of the user study we calculated the percentage of users who successfully followed the robot’s convention (*Success Rate*). To be successful, participants had to (a) complete the intended task and (b)

³We note that this baseline is interchangeable with other shared autonomy approaches that infer the human’s intent.

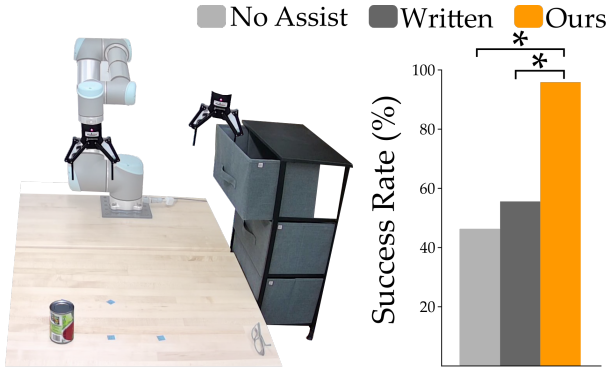


Fig. 4. Experimental setup and results from the third part of the user study. (Left) Here the robot had a complex, sinusoidal convention, and participants needed to match this convention to get the robot to open the drawer or reach the soup can. (Right) We found that participants who *see* the convention (with **Ours**) outperform participants who *read* the convention (with **Written**). Here * denotes statistical significance ($p < .05$).

provide fewer joystick inputs than the average across all users. For the second task we calculated how the robot’s belief in the human’s task changed before and after part one (*Belief Change*). We found the robot’s average confidence in the human’s true task during their initial interactions, and then subtracted this from the robot’s average confidence during the most recent interactions. Hence, this metric captures whether there was a change in how participants teleoperated the robot after experiencing the robot’s convention. Finally, in the third part of the user study we again leveraged *Success Rate* to calculate how many times the robot correctly and efficiently inferred the intended task.

Participants and Procedure. A total of 27 members of the Virginia Tech community participated in our user study (8 female, average age 25 ± 4 years). All participants provided informed written consent prior to the experiment.

Hypotheses. We tested three hypotheses:

- H1.** *Robots will reveal their conventions by guiding humans towards more communicative inputs.*
- H2.** *Over repeated interactions, humans will generalize these conventions to new scenarios.*
- H3.** *Revealing conventions through robot motion will outperform written descriptions.*

Results. The results from the first and second part of our user study are shown in Figure 3, and the results from the third part are shown in Figure 4.

In the first part of our user study we measured *Success Rate* to see whether participants would adapt to the robot’s conventions. Across all 27 participants, users who interacted with **Ours** were most likely to follow the robot’s convention. Post hoc comparisons between **Ours** and **No Assist** were statistically significant for all three object locations ($p < .05$). When comparing against **Written**, we found that *Success Rate* was significantly higher for locations 2 and 3 ($p < .05$). These results support **H1**, and suggest that our proposed approach encouraged the participants to exaggerate their joystick inputs and concisely convey their desired task.

The results from the second part of our user study were

not clear-cut. Although users did improve after working with the robot in all conditions, the gains in *Belief Change* were minimal (between 2% and 4%). This suggests that participants did not internalize the robot’s conventions or transfer those conventions to new scenarios. One possible explanation is that — because the robot was directly following the human’s commanded actions — participants saw no need to exaggerate and indicate their task. Our results from the second part of the user study do not support **H2**.

In the third part of the user study we again measured *Success Rate*, but now with the robot following a complex and unintuitive convention. As shown in Figure 4, participants who interacted with **Ours** were best able to match this convention: post hoc analysis confirms that the differences are statistically significant between **Ours** and **No Assist** ($p < .05$), and between **Ours** and **Written** ($p < .05$). Users who interacted with either **No Assist** or **Written** had difficulty learning the robot’s unintuitive convention, perhaps because it was easier to *show* than to *tell*.

Looking at each part of the user study, we find support for **H1** and **H3**. Robots that leveraged our approach from Algorithm 1 were not only able to communicate their convention to the human, but they also communicated this convention more effectively than written instructions. On the other hand, our results do not support **H2**. None of the methods caused participants to transfer their learned conventions to a new scenario; however, we recognize that this may have been because the robot was not actively inferring their task or assisting the human in this scenario.

Discussion. Our results suggest that users who interact with **Ours** can be divided into two distinct groups. The majority of participants (8 out of the 9 users) adapted to the robot’s convention immediately — these participants only needed one or two interactions with our approach to understand the mapping from joystick to task. At the other end of the spectrum we also had one participant who *never* adapted to the robot’s convention (1 of the 9 users). This participant was often confused by **Ours**, especially when the robot intervened to modify their motion. Instead of mimicking the robot’s behavior, the user applied the opposite input in order to cancel out any autonomous guidance and move directly to the goal. Moving forward, we believe that we can reach this second group by combining both **Written** and **Ours**. We hypothesize that the written instructions will provide the context these users need to understand why the robot is altering their motion.

VI. CONCLUSION

When humans teleoperate robots there are many conventions the robot can leverage to map joystick inputs to discrete tasks. We have enabled robots to actively communicate their chosen convention. Specifically, we leveraged constrained optimization to identify robot actions that *reveal* how humans should convey the current task while simultaneously *assisting* for that task. Our user study results suggest that this shared autonomy approach reduces the number of joystick inputs that humans need to indicate their task.

REFERENCES

- [1] 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.
- [2] 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.
- [3] 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.
- [4] H. J. Jeon, D. P. Losey, and D. Sadigh, "Shared autonomy with learned latent actions," in *Robotics: Science and Systems*, 2020.
- [5] A. Jonnavittula and D. P. Losey, "Learning to share autonomy across repeated interaction," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.
- [6] M. Zurek, A. Bobu, D. S. Brown, and A. D. Dragan, "Situational confidence assistance for lifelong shared autonomy," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.
- [7] C. Brooks and D. Szafr, "Balanced information gathering and goal-oriented actions in shared autonomy," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2019, pp. 85–94.
- [8] M. Fontaine and S. Nikolaidis, "A quality diversity approach to automatically generating human-robot interaction scenarios in shared autonomy," in *Robotics: Science and Systems*, 2021.
- [9] D. M. Taylor, *Americans With Disabilities: 2014*. US Census Bureau, 2018.
- [10] B. D. Argall, "Autonomy in rehabilitation robotics: An intersection," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 441–463, 2018.
- [11] 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.
- [12] D. Gopinath, S. Jain, and B. D. Argall, "Human-in-the-loop optimization of shared autonomy in assistive robotics," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 247–254, 2016.
- [13] D. Park, Y. Hoshii, H. P. Mahajan, H. K. Kim, Z. Erickson, W. A. Rogers, and C. C. Kemp, "Active robot-assisted feeding with a general-purpose mobile manipulator: Design, evaluation, and lessons learned," *Robotics and Autonomous Systems*, vol. 124, p. 103344, 2020.
- [14] C. Boutilier, "Sequential optimality and coordination in multiagent systems," in *International Joint Conference on Artificial Intelligence*, vol. 99, 1999, pp. 478–485.
- [15] A. Shih, A. Sawhney, J. Kondic, S. Ermon, and D. Sadigh, "On the critical role of conventions in adaptive human-ai collaboration," in *International Conference on Learning Representations*, 2021.
- [16] D. Hadfield-Menell, S. J. Russell, P. Abbeel, and A. Dragan, "Co-operative inverse reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [17] H. Hu, A. Lerer, A. Peysakhovich, and J. Foerster, "Other-Play for zero-shot coordination," in *International Conference on Machine Learning*, 2020, pp. 4399–4410.
- [18] S. Milli and A. D. Dragan, "Literal or pedagogic human? Analyzing human model misspecification in objective learning," in *Uncertainty in Artificial Intelligence*, 2020, pp. 925–934.
- [19] X. Zhu, "Machine teaching: An inverse problem to machine learning and an approach toward optimal education," in *AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.
- [20] M. Cakmak and M. Lopes, "Algorithmic and human teaching of sequential decision tasks," in *AAAI*, 2012.
- [21] M. Cakmak and A. L. Thomaz, "Eliciting good teaching from humans for machine learners," *Artificial Intelligence*, pp. 198–215, 2014.
- [22] A. Sena and M. Howard, "Quantifying teaching behavior in robot learning from demonstration," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 54–72, 2020.
- [23] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2013, pp. 301–308.
- [24] A. D. Dragan, S. Bauman, J. Forlizzi, and S. S. Srinivasa, "Effects of robot motion on human-robot collaboration," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2015, pp. 51–58.
- [25] S. Nikolaidis, Y. X. Zhu, D. Hsu, and S. Srinivasa, "Human-robot mutual adaptation in shared autonomy," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2017, pp. 294–302.
- [26] J. Bragg and E. Brunskill, "Fake it till you make it: Learning-compatible performance support," in *Uncertainty in Artificial Intelligence*, 2020, pp. 915–924.
- [27] B. A. Newman, R. M. Aronson, S. S. Srinivasa, K. Kitani, and H. Admoni, "Harmonic: A multimodal dataset of assistive human-robot collaboration," *arXiv preprint arXiv:1807.11154*, 2018.
- [28] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, and J. Peters, "An algorithmic perspective on imitation learning," *Foundations and Trends in Robotics*, vol. 7, no. 1-2, pp. 1–179, 2018.
- [29] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *AAAI Conference on Artificial Intelligence*, vol. 8, 2008, pp. 1433–1438.
- [30] A. Dragan and S. Srinivasa, "Generating legible motion," in *Robotics: Science and Systems*, 2013.
- [31] A. Slivkins, "Introduction to multi-armed bandits," *arXiv preprint arXiv:1904.07272*, 2019.
- [32] S. Mannor and J. N. Tsitsiklis, "The sample complexity of exploration in the multi-armed bandit problem," *Journal of Machine Learning Research*, vol. 5, pp. 623–648, 2004.
- [33] L. Chan, D. Hadfield-Menell, S. Srinivasa, and A. Dragan, "The assistive multi-armed bandit," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2019, pp. 354–363.