

# Effects of Robot Competency and Motion Legibility on Human Correction Feedback

Shuangge Wang  
Yale University  
New Haven, CT, USA  
shuangge.wang@yale.edu

Anjiabei Wang  
Yale University  
New Haven, CT, USA  
anjiabei.wang@yale.edu

Sofiya Goncharova  
Phillips Exeter Academy  
Exeter, NH, USA  
sgoncharova@exeter.edu

Brian Scassellati  
Yale University  
New Haven, CT, USA  
brian.scassellati@yale.edu

Tesca Fitzgerald  
Yale University  
New Haven, CT, USA  
tesca.fitzgerald@yale.edu

**Abstract**—As robot deployments become more commonplace, people are likely to take on the role of supervising robots (i.e., correcting their mistakes) rather than directly teaching them. Prior works on Learning from Corrections (LfC) have relied on three key assumptions to interpret human feedback: (1) people correct the robot only when there is significant task objective divergence; (2) people can accurately predict if a correction is necessary; and (3) people trade off precision and physical effort when giving corrections. In this work, we study how two key factors (robot competency and motion legibility) affect how people provide correction feedback and their implications on these existing assumptions. We conduct a user study ( $N = 60$ ) under an LfC setting where participants supervise and correct a robot performing pick-and-place tasks. We find that people are more sensitive to suboptimal behavior by a highly competent robot compared to an incompetent robot when the motions are legible ( $p = 0.0015$ ) and predictable ( $p = 0.0055$ ). In addition, people also tend to withhold necessary corrections ( $p < 0.0001$ ) when supervising an incompetent robot and are more prone to offering unnecessary ones ( $p = 0.0171$ ) when supervising a highly competent robot. We also find that physical effort positively correlates with correction precision, providing empirical evidence to support this common assumption. We also find that this correlation is significantly weaker for an incompetent robot with legible motions than an incompetent robot with predictable motions ( $p = 0.0075$ ). Our findings offer insights for accounting for competency and legibility when designing robot interaction behaviors and learning task objectives from corrections.

**Index Terms**—Interactive Robot Learning; Learning from Corrections; Kinesthetic Teaching

## I. INTRODUCTION

Imagine that you are developing robots that will be deployed in warehouses and tasked with packing merchandise for shipment. These robots are pre-trained with basic abilities such as recognizing boxes and manipulating a range of items. Despite this pre-deployment training, they will inevitably make mistakes with identifying and manipulating novel inventory that is specific to the warehouse where the robot is deployed. Furthermore, new inventory may be introduced over time. In order for robots to function effectively post-deployment, they should learn continuously from their human co-workers and collaborators as novel objects and task constraints arise.

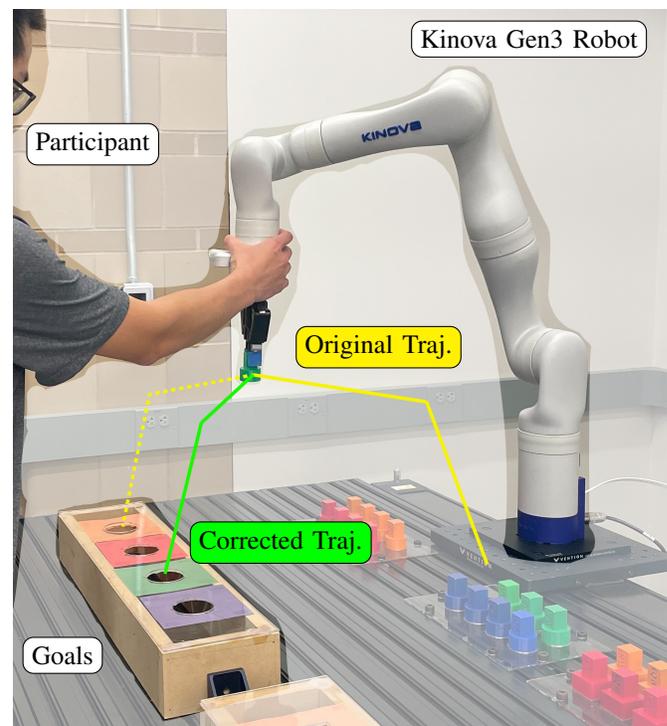


Fig. 1: The robot begins moving along the yellow trajectory to place a shape in one of the target holes. The participant intervenes, guiding the robot to a different target via the green trajectory, leaving the dashed portion of the robot's intended trajectory untraveled.

Traditionally, robots learn from people during training sessions that are distinct from its deployment/test sessions. In these training sessions, a person directly teaches the robot by providing numerous demonstrations of the desired behavior [1]–[5] or indicating preferences between numerous hypothetical robot behaviors [6]–[11].

There has been comparatively less work on how people might teach robots by supervising its behavior *during* deployment [12]–[15]. In a supervisory paradigm, a person observes

the robot’s behavior as it attempts to complete a task, intervening only when necessary to avoid an impending mistake or to train the robot to avoid future mistakes [16]–[30]. Compared to traditional training sessions, this paradigm leverages the robot’s autonomy whenever possible and avoids the need for continuous teaching effort by the human during long-horizon tasks [31]. However, supervisory feedback might not be as abundant [32]; the human might not be as attentive [33], [34]; and the quality of human feedback might not be as consistent [35], [36].

Yet, current learning techniques do not account for these differences when learning from supervisory feedback, which may impair the robot’s ability to learn accurate and precise task objectives. Existing learning algorithms are designed under the assumption that (1) people intervene and correct the robot only when it is about to make a mistake; (2) people can accurately predict when they do or don’t need to intervene; and (3) people trade off precision and physical effort when providing corrections. Prior works have yet to explore whether these assumptions are valid, or what other factors may influence how people correct a robot during long-duration tasks. From a learning perspective, it is important to identify and factor out these influences to facilitate effective learning for the underlying task information. Furthermore, to design robots and algorithms that people will adopt into their homes and workplaces, we need to understand better how a person reacts to these factors when supervising a robot.

Our work aims to study how people give feedback to a robot in a supervisory setting. We study two factors that we expect are particularly important in influencing how people correct the robot: robot competency (i.e., the robot’s prior task performance) and motion legibility (i.e., the interpretability of the robot’s motion toward its goal). We present an between-subject user study ( $N = 60$ ) where we examine how the robot’s competency and legibility conditions influence how participants supervise it through a series of 64 pick-and-place actions (Fig. 1). To the best of our knowledge, our work is the first to identify how these features of a robot’s behavior affect people’s correction feedback as they supervise a robot. We provide the following contributions:

- We design a learning from corrections study to measure how people supervise a robot’s behavior differently depending on its competency and motion legibility.
- We study how these factors influence the timing and perceived necessity of corrections.
- We examine the trade-off between precision and physical effort as people provide corrections.
- We propose recommendations for (1) designing robots to elicit better correction feedback post-deployment and (2) interpreting this feedback as training data.

## II. RELATED WORK

### A. Learning from Corrections

Research in interactive robot learning has explored multiple modalities through which people can provide training data and

feedback to a robot, such as rewards [37]–[39], demonstrations [1]–[5], preferences [6]–[11], physical corrections [16]–[25], implicit feedback [40]–[42], and natural language [43]–[46]. These modalities differ in the role that they prescribe to the human and the robot during the interaction [47]. Some modalities position humans in a more supervisory role than others, such as interventions [26]–[29], negative reinforcement [48]–[50], implicit feedback [40]–[42], and physical corrections [16]–[21], [23]–[25].

Correction feedback, in particular, involves a robot attempting to complete a task while supervised by a human teacher. The teacher can intervene and modify the robot’s motion kinesthetically, producing a *corrected* trajectory that is assumed to be more optimal with respect to the hidden task objectives [16]–[21], [23]–[25]. While prior works have proposed methods for Learning from Corrections (LfC), they rely on three important assumptions.

#### **Assumption 1: People Correct Only When There Is Significant Task Objective Divergence**

Prior works in learning from interventions [27]–[29] and corrections [22], [23] assume that people decide to intervene and correct the robot only when the robot is about to make a mistake. Quantitatively, we can represent this decision as a threshold for allowable divergence between the robot’s behavior and the “correct” goal for its task. As an example, a robot that moves away from an object it should be picking up and toward objects that it should be avoiding is *increasing* this divergence. Prior work assumes that this threshold is either consistent across all users [51], is user-specific [52], [53], or task-specific [54], but does not consider how the robot’s own behavior might influence this threshold.

#### **Assumption 2: People Can Accurately Predict If a Correction is Necessary**

A person’s decision to correct a robot’s motion can be viewed as a label for incorrect behavior; conversely, the lack of an intervention can be considered to be an endorsement [48], [49]. Current methods for learning from human-provided labels require that the human achieves optimal or near-optimal labelling accuracy [55], [56]. While previous works have attempted to design algorithms that learn from noisy human-provided labels [57], [58], the robot’s learning performance will still be affected by inaccurate labels.

#### **Assumption 3: People Trade Off Precision and Effort**

Prior works in interactive robot learning assume that the teacher is incentive-driven [52], [54], [59]–[61], providing feedback to maximize perceived reward and minimize perceived cost. Existing LfC models, in particular, assume that people aim to provide corrections that optimize task performance, while being biased toward corrections that require less physical effort to provide [16]–[21], [23]. Yet, no prior work has empirically validated this trade-off.

### B. Robot Competency

In response to these assumptions, we now consider additional factors that may influence how people correct a robot’s

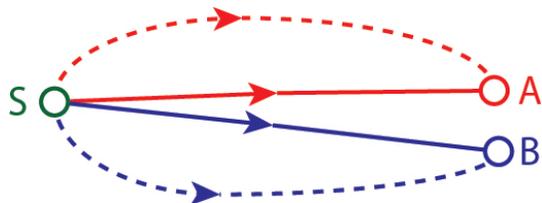


Fig. 2: A comparison between predictable (solid lines) and legible motions (dashed lines) starting from S to two potential goals, A and B.

behavior. Deployed robots will naturally exhibit different levels of competency across various tasks based on how well their training data aligns with their deployment environment [62] or constraints of its sensor, actuator, and computational capabilities [63], [64]. Prior works have shown that the robot’s competency is one of the most prominent factors in shaping people’s trust [65], expectation [66], and preferences [67] over the robot’s behaviors. DelPreto et al. [68] found that when a robot performed a task with lower accuracy, participants’ trust in the robot decreased, their perception of the robot’s intelligence decreased, and their workload increased. Hedlund et al. [69] showed that an incompetent robot led people to lower their trust in both the robot and in their own teaching ability. Paepcke et al. [66] found that setting low expectations for the robot’s ability at a task resulted in less disappointment and more positive appraisals from participants.

Overall, these prior works study how the robot’s competency affects a person’s subjective perception of it. Yet, they do not evaluate how this subjective perception influences how the person provides feedback to the robot. From a robot learning perspective, other works have focused on the effect of robot competency in a setting where people serve as direct teachers (i.e., in learning from demonstration) [68], [69]. Yet, this remains unexplored in *supervisory* paradigms such as LfC.

### C. Motion Legibility

While competency describes the robot’s *past* behavior, we now consider how the robot’s *current* behavior (i.e., its motion trajectory) influences a person’s perception of it while they supervise the robot. Typical motion planning algorithms optimize for short and efficient trajectories between the robot’s starting position and its goal position. This produces a trajectory that is *predictable*; i.e., the motion that a person would expect to see if they *already know* what goal position the robot is trying to reach [61]. To a person who does *not* know the robot’s goal, however, it may be difficult for them predict the robot’s goal in real-time as the robot moves. The aim of *legible* motion is to enable an observer to quickly predict the robot’s goal from its early motion [61], [70]. This typically results in the robot exaggerating its motion toward its goal, as shown in Fig. 2.

Prior works have verified that legible motion is more effective at conveying the robot’s intent to a human observer [61], [70], [71] and can elicit more informative feedback [72]–[74]. Furthermore, legible motion can improve a person’s perceived safety [75], comfort [76], trust [77], and positive affect of robots [78]. While prior works have focused on generating

legible motions [61], [70] and incorporating legibility while modeling human feedback in task learning [23], [79], its nuanced effects on the nature and quality of human feedback remain largely unexplored.

## III. RESEARCH QUESTIONS

Our work investigates the influence of the robot’s competency and motion legibility on how people correct its behavior. Our aim is to inform how we design (1) interactions for robots to elicit better correction feedback and (2) learning algorithms to interpret corrections for more accurate task models. We focus on the following three research questions (RQs) investigating how competency and legibility affect the validity of the assumptions described in Section II-A.

*RQ1: How do competency and legibility affect when people correct the robot?*

In response to **Assumption 1** (Section II-A), we consider how competency and legibility may aggravate or attenuate a person’s threshold for divergence (which informs their decision of when to correct the robot). We expect that a robot exhibiting legible motion will enable people to infer the goal of its motion earlier in its trajectory execution, allowing them to more quickly assess whether it will be aligned with the correct goal (and thus, whether a correction is necessary). We also expect that people will trust a highly competent robot more, thus raising their threshold for divergence and resulting in later corrections. We establish the following hypotheses:

**H1A:** When supervising robots with legible motion, people will correct the robot earlier in the trajectory (i.e., when there is a smaller task objective divergence).

**H1B:** When supervising a highly competent robot, people will correct the robot later in the trajectory (i.e., will have a higher tolerance for task objective divergence).

*RQ2: How do competency and legibility affect people’s accuracy in predicting robot success/failure?*

In response to **Assumption 2**, we expect that people will distrust an incompetent robot, and thus are more likely to predict that it will fail (even if it would have succeeded using its intended trajectory). Similarly, we expect that people will trust a highly competent robot, and thus are more likely to predict that a failing robot will eventually succeed (until it is too late for them to provide a correction). Since legible motion enables a person to predict the robot’s goal sooner, we also expect legible motions to improve predication accuracy. We establish the following hypotheses:

**H2A:** People are more likely to miss necessary corrections in high-competency conditions.

**H2B:** People are more likely to provide unnecessary corrections in low-competency conditions.

**H2C:** Legible motions will increase prediction accuracy, regardless of competency condition.

*RQ3: How do competency and legibility affect how people balance precision and effort in their corrections?*

In response to **Assumption 3**, we aim to confirm whether there is a consistent trade-off between the precision and physical effort that people expend as they correct the robot's motion. **H3:** People's corrections will exhibit a trade-off between precision and physical effort.

#### IV. METHODOLOGY

We conducted a user study with 60 participants recruited from our university community. Participants reported their age ( $M = 27.53$  years,  $SD = 9.49$  years) and gender (32 male, 27 female, 1 non-binary). Each participant supervised a robot arm performing a series of pick-and-place tasks. The task goal was for the robot to place each shape into the target hole with the corresponding color. Participants were instructed to interrupt the robot's motion and provide a correction whenever and however they saw fit to guide the robot toward successfully completing the task (Fig. 1).

To incentivize high-quality data (as would be expected of a person earnestly trying to train a robot collaborator), participants were told that the robot was learning from their feedback in real-time, and that they would receive additional compensation based on the number of successful robot trials. In reality, the robot followed pre-determined waypoints based on the participant's study condition (rather than learning in real-time), and participants received the maximum compensation (as if the robot had succeeded at every trial) to ensure that they were fairly compensated regardless of their study condition. We obtained approval for this study through our Institutional Review Board (IRB) and followed ethics protocol for debriefing participants on these hidden elements.

##### A. Experimental Design

Our experiment involved two independent variables: competency (consisting of two levels) and legibility (consisting of three levels). Our experiment was thus a between-subject  $2 \times 3$  user study, with 10 participants assigned to each condition. In each condition, participants supervised 64 task trials, divided into 4 sub-tasks (i.e., 4 different shapes) with 16 trials each.

Throughout this section, *actual* success/failure refers to whether the robot places a shape in the correct target. *Intended* success/failure refers to the result of the robot's planned trajectory without correction.

1) *Independent Variables:* In **low competency** conditions, the robot intended to succeed in only 25% of trials. Among the intended failures, 50% involved the robot placing the shape in a hole of the wrong color. The other 50% involved the robot missing the target holes entirely, with intended failure poses uniformly sampled from a neighborhood near the correct target. Fig. 4 illustrates the combined probability distribution over potential goals for an intended failure. In **high competency** conditions, the robot intended to succeed in 75% of trials, with intended failures also following the distribution shown in Fig. 4. The intended failures were distributed across colors and trials to minimize bias.

Motion legibility [61] consisted of 3 levels, shown in Fig. 3:

- **Predictable** motion was short and efficient; i.e., the default output of an RRT\* motion planner.
- **Legible** motion enabled the observer to quickly infer the robot's end-goal.
- **Illegible** motion obscured the robot's end goal by initially moving toward goals that the robot did not plan to visit.

##### B. Robot Control & Motion Planning

We performed the study using a Kinova Gen3 7-DoF robot arm equipped with a Robotiq 2F-85 gripper as its end effector (EEF). The arm was mounted on a horizontal linear actuator, allowing it to slide along the workbench to access each of the four sub-tasks. We computed Cartesian waypoints for every trajectory for each legibility level by optimizing their legibility score until convergence [61]. We then used RRT\* to plan the Cartesian waypoints into joint space trajectory [80]. We down-sampled and interpolated the planned trajectory for improved efficiency and smoothness. The trajectory was then executed via a velocity-based PID controller running at sub-1000 Hertz. To avoid bias toward certain colors, we uniformly randomized the order in which the robot picked up shapes in each sub-task.

We implemented admittance (inverse impedance) control [81], [82] to enable the robot to quickly follow any physical force exerted by the participant. We used the recursive Newton-Euler algorithm to compute the inverse dynamics and gravity compensation [83]. We collected joint encoder and torque sensor readings and control input at 10Hz and applied low pass filters to denoise the data. After the participant stopped applying force to the robot arm (i.e., after finishing the correction), the robot replanned and executed a trajectory from its new state to the nearest target goal, maintaining the EEF rotation as it was at the end of the correction. The participant may provide additional corrections during the robot's replanned motion; regardless, we focus our analyses on the *first* correction within each trial.

##### C. Procedures

1) *Pre-Interaction Phase:* Participants were pre-screened to meet the criteria of being at least 18 years old, fluent in English, able to stand for up to an hour, and not color blind. After the screening, the participant consented to the experiment and video recording. The experimenter followed a script to familiarize the participant with the robot, the pick-and-place task, and their role in correcting the robot.

2) *Tutorial Phase:* The researcher activated the robot's admittance mode to enable the participant to practice moving the robot around. Once the participant was ready to proceed, the experimenter played two example behaviors on the robot: one successful placing trial (requiring no correction) and one unsuccessful trial (where the experimenter demonstrated how to intervene and correct the robot). The participant then administered two similar trials.

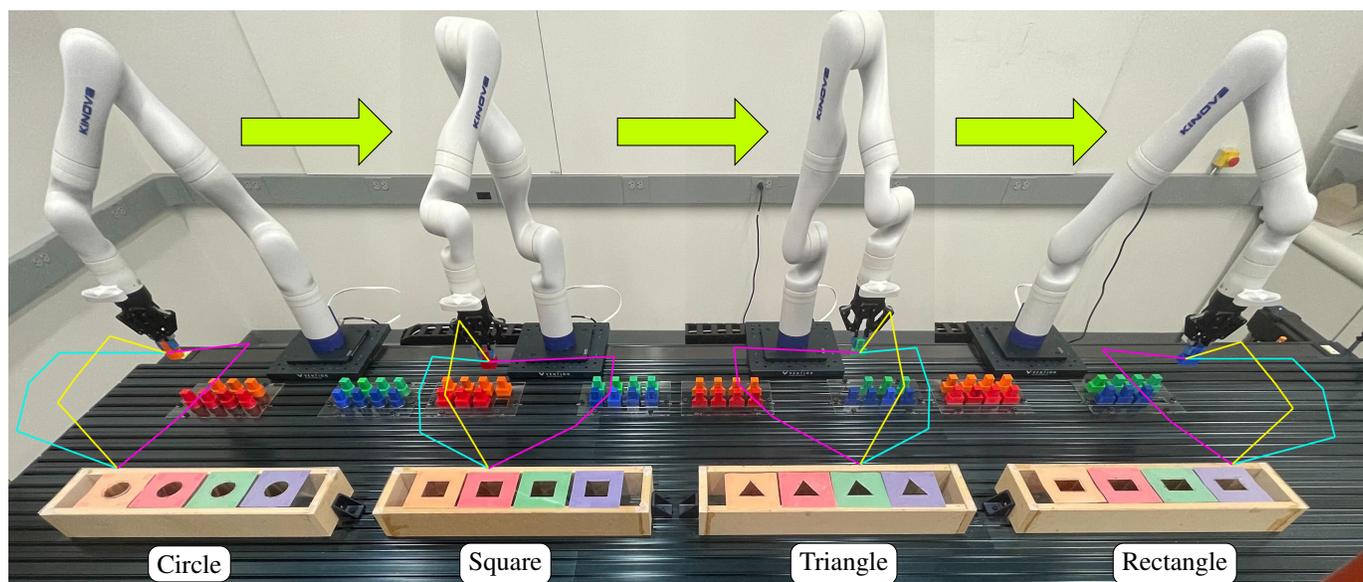


Fig. 3: This figure depicts the series of pick-and-place tasks performed by the robot. The tasks involved manipulating various shapes in sequence: circle, square, triangle, and rectangle. Each shape had 4 colors: orange, red, green, and blue. The colored lines illustrate the robot's trajectories, categorized by their legibility: cyan represents legible paths, yellow represents predictable paths, and purple represents illegible paths.

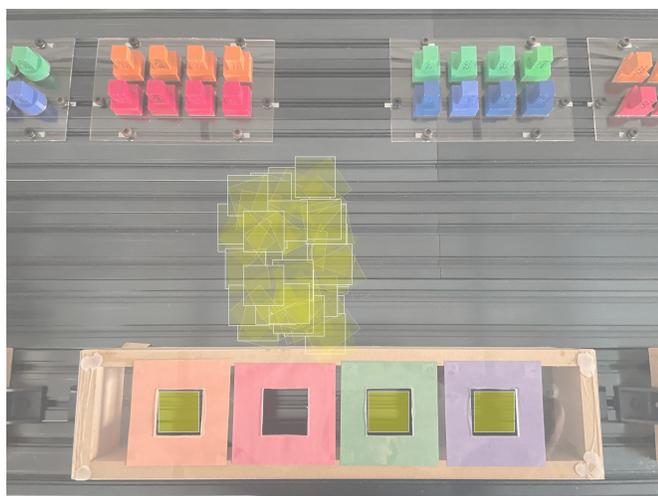


Fig. 4: In an intended-success trial, the robot will attempt to place a red square shape into the red target. In an intended-failure trial, the robot will attempt to place the red shape into the wrong-colored target or onto the workbench surface. The yellow-shaded squares represent the distribution of possible locations where the robot may attempt to place the shape.

3) *Main Experiment Phase*: After completing the tutorial, the experimenter left the participant alone to supervise and correct the robot for all 64 pick-and-place trials. The experimenter remained on the other side of a dividing curtain, monitoring the experiment for safety via a webcam and staying within reach of an emergency-stop button.

4) *Post-Interaction Phase*: After the experiment, the participant completed a survey on demographics, System Usability Scale [84], Perception of Agency [85], NASA Task Load

Index [86], and Trust in Automation [87] measures. Finally, the experimenter debriefed the participant about the deception, explaining that the robot was not actually learning from them in real-time and that they would receive the maximum compensation regardless of the robot's performance.

#### D. Measures

We now define the metrics we use to evaluate each RQ.

*RQ1: When do people correct the robot?:* Kullback–Leibler divergence (KLD) [88] has been widely adopted by learning researchers as a metric to model the human's perceived alignment between the robot's understanding of the task and the actual task constraints [72]–[74], where a smaller KLD indicates that the human believes the robot has more accurately learned the task. In addition, we included two other timing-related heuristic measures: the time elapsed prior to the correction and the proportion of the robot's intended trajectory that remained untraveled.

- **Task Objective Divergence**: We estimate the distribution of plausible motion goals supported by the robot's motion at each timestep [89]. When the participant intervenes to correct the robot's motion, we report the KLD between this distribution (based on the robot's motion thus far) and the actual, correct motion goal.
- **Time Until Correction**: The time (in seconds) between the start of the robot's trajectory and the first correction.
- **Proportion of Trajectory Untraveled**: The fraction of the trajectory that was left to be traveled at the time of correction (dashed trajectory in Fig. 1).

*RQ2: How well do people predict robot success/failure?:* Since the presence or absence of human intervention can be

construed as a binary label for correct behavior, researchers have tried to leverage this data to train robots to classify correct behavior. We therefore categorize human prediction outcomes within a confusion matrix [90]:

- **Missed Correction Rate:** The proportion of intended failures that the participant did not correct.
- **Unnecessary Correction Rate:** The proportion of intended successes that the participant corrected.
- **False Omission Rate:** The proportion of uncorrected trials that were intended failures.
- **False Correction Rate:** The proportion of corrected trials that were intended successes.

*RQ3: How do people trade-off precision and effort in their corrections?:* Prior works in LfC have modeled the task objective as a linear combination of relevant features [16]–[23], such as goal proximity, object avoidance, and motion smoothness and quantified physical effort using external torque applied by humans on the robot [16]–[23].

- **Precision:** A linear combination of EEF position and rotation features with respect to the correct goal at the end of the first correction (details in Appendix Sec. VIII-A).
- **Physical Effort:** The L2-norm of the time-integrated torque for the first correction.

## V. RESULTS

For **RQ1** and **RQ2**, we conducted a two-way Analysis of Variance (ANOVA) [91] with competency and legibility predicting each relevant measure in Section IV-D. The main effects were ignored if the interaction effect was significant. A post-hoc Tukey’s honestly significant difference (HSD) test was administered if the legibility main effect or interaction effect was significant. Although it is generally advisable to transform the data so that it follows a normal distribution, Blanca et al. have shown that the ANOVA and F-test are robust to non-normal data [92]. For **RQ3**, we examined the pairwise condition difference in correlation between task precision and physical effort. We measured statistical significance using a threshold of  $\alpha < .05$  and applied the Benjamini–Hochberg procedure for multiple testing corrections [93]. In all figures, asterisks (\*) are used to denote statistical significance: a single asterisk indicates  $p < 0.05$ ; double asterisks indicate  $p < 0.01$ ; and triple asterisks indicate  $p < 0.001$ .

By default, a robot’s motion planner will produce predictable (i.e., efficient) trajectories, which we considered as the baseline legibility. Hence, all pairwise comparisons included predictable vs. legible and predictable vs. illegible conditions.

### A. Analysis of RQ1

A two-way ANOVA examining the effect of competency and legibility on task objective divergence showed an interaction effect of ( $F(2, 1944) = 4.8242, p = 0.0081$ ). A Tukey test found that 1) in legible conditions ( $p = 0.0015$ ), people choose to correct the robot when the task objective divergence was smaller in high-competency conditions ( $M = 0.0047, SD = 0.0073$ ) than in low competency-conditions ( $M = 0.0075, SD = 0.0090$ ), 2) in predictable

( $p = 0.0055$ ) conditions, people choose to correct the robot when the task objective divergence was smaller in high-competency conditions ( $M = 0.0051, SD = 0.0064$ ) than in low-competency conditions ( $M = 0.0078, SD = 0.0092$ ). A two-way ANOVA examining the effect of competency and legibility on time until correction (Fig. 6) showed an interaction effect of ( $F(2, 1944) = 23.3877, p < 0.0001$ ). A Tukey’s test showed that 1) when motions were legible ( $p < 0.0001$ ), people correct a competent robot earlier in high-competency conditions ( $M = 13.5968, SD = 3.3740$ ) than in low-competency conditions ( $M = 14.7866, SD = 2.0448$ ), 2) when motions were predictable ( $p = 0.0020$ ), people correct a competent robot earlier in high-competency conditions ( $M = 13.4841, SD = 2.0831$ ) than in low competency-conditions ( $M = 14.2968, SD = 1.7970$ ). A two-way ANOVA examining the effect of proportion of trajectory to be traveled (Fig. 7) showed an interaction effect for competency and legibility of ( $F(2, 1944) = 34.7386, p < 0.0001$ ). A Tukey’s test showed that 1) when motions were legible ( $p < 0.0001$ ), robots were corrected later in the trajectory in low-competency ( $M = 0.0435, SD = 0.1136$ ) than high-competency conditions ( $M = 0.1610, SD = 0.2265$ ) 2) when motions were predictable ( $p < 0.0001$ ), robots were corrected later in the trajectory in in low-competency ( $M = 0.0457, SD = 0.0902$ ) than high-competency conditions ( $M = 0.1422, SD = 0.1849$ ).

### B. Analysis of RQ2

A two-way ANOVA examining the effect of competency and legibility on missed correction rate (Fig. 8) showed a competency main effect of ( $F(1, 54) = 17.4059, p < 0.0001$ ), indicating that, in low-competency conditions ( $M = 0.1127, SD = 0.1057$ ), people were more likely to withhold necessary corrections than in high-competency conditions ( $M = 0.0281, SD = 0.0322$ ). A two-way ANOVA examining the effect of competency and legibility on unnecessary correction rate (Fig. 9) showed a competency main effect of ( $F(1, 54) = 6.0503, p = 0.0171$ ), suggesting that, in high-competency conditions ( $M = 0.0983, SD = 0.1695$ ), people were more likely to offer unnecessary corrections than in low-competency conditions ( $M = 0.0202, SD = 0.0401$ ). From a robot-centric perspective, a two-way ANOVA examining the effect of competency and legibility on false omission rate (Fig. 10) showed a competency main effect of ( $F(1, 54) = 5.5637, p = 0.0220$ ), implying that, in high-competency conditions ( $M = 0.0708, SD = 0.0586$ ), failures were more likely when corrections were missed than in low-competency conditions ( $M = 0.0410, SD = 0.0404$ ). A two-way ANOVA on false omission rate (Fig. 11 in the Appendix) showed no significant effect of competency or legibility.

### C. Analysis of RQ3: Task and Effort Trade-Off

Two Shapiro-Wilk tests [94] showed that task precision ( $W = 0.8776, p < 0.001$ ) and physical effort ( $W = 0.7188, p < 0.001$ ) were not normally distributed. We, therefore, used the Spearman correlation [95] to study the relationships

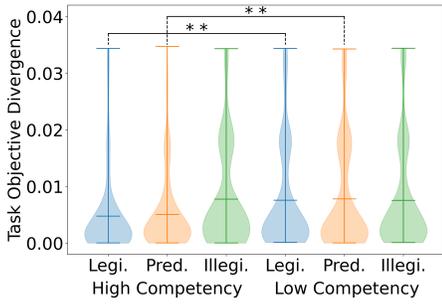


Fig. 5: Task Objective Divergence: The KLD between the estimated robot’s task belief distribution and the actual, correct motion goal.

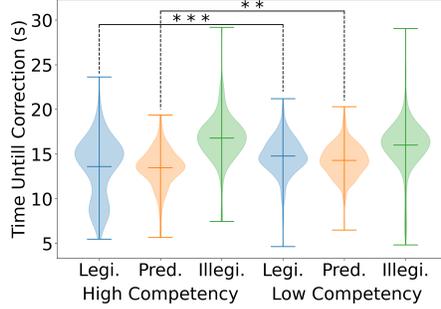


Fig. 6: Time Until Correction: The time (measured in seconds) between the start of the robot’s trajectory and the first correction.

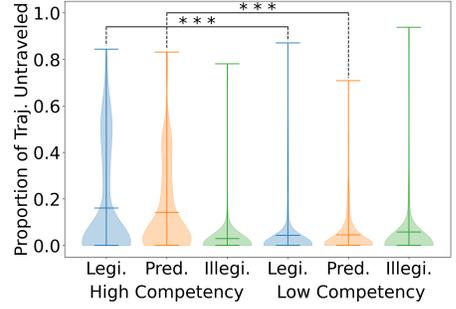


Fig. 7: Proportion of Trajectory Untraveled: The fraction of the trajectory that was left to be traveled at the time of correction.

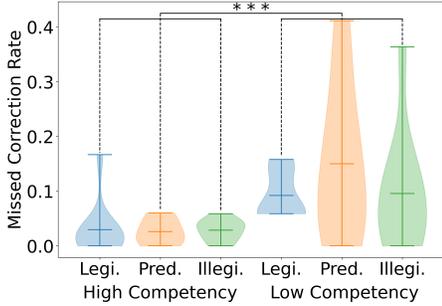


Fig. 8: Missed Correction Rate: The proportion of intended failures that the participant did not correct.

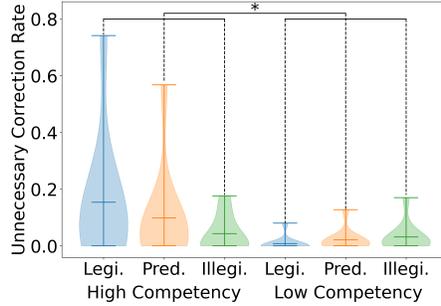


Fig. 9: Unnecessary Correction Rate: The proportion of intended successes that the participant corrected.

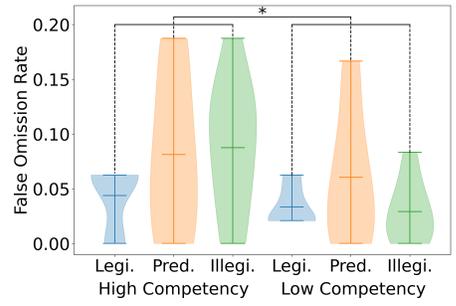


Fig. 10: False Omission Rate: The proportion of uncorrected trials that were intended failures.

between these two variables. Overall, the trend in each condition exhibited a weak positive correlation (Table I in the Appendix). After adopting a Fisher z-transformation on the correlations [96], [97], we conducted a pairwise Z-test to examine how competency and legibility affect these positive correlations (Table II in the Appendix). We found that, for an incompetent robot with legible motions, the positive correlation of task and effort was significantly weaker compared to an incompetent robot with predictable motions ( $p = 0.0075$ ).

## VI. DISCUSSION

Addressing **RQ1**, we find that, for legible and predictable motions, people chose to intervene as soon as they perceive a minor error in the robot’s behavior in high-competency conditions; in low-competency conditions, people were more tolerant of the robot deviating from optimal behavior. This which strongly supports the inverse of **H1B**; that is, people corrected the robot *earlier* in the trajectory when supervising a highly competent robot (with legible and predictable motions). We find no evidence supporting or disproving **H1A**.

Corrections are most informative when they reflect the constraints of the task. Therefore, it is advantageous for learning algorithms if people correct the robot only when it is about to make a mistake. To achieve this, robots should avoid deviating from optimal behaviors when highly competent. From a learning perspective, algorithms should not automatically assume that pre-correction trajectories are always informative

(particularly for an incompetent robot). Instead, they should primarily leverage the correction trajectory.

For **RQ2**, we find that competency significantly influences people’s ability to accurately predict robot success and failures. Our evidence does not support **H2A** and **H2B**, but strongly supports their inverse: when supervising an incompetent robot, people are more likely to miss necessary corrections; people are more likely to give corrections when not needed while supervising a highly competent robot.

We find no evidence supporting or disproving **H2C**; legibility does not significantly affect humans’ prediction accuracy. Note that, although both false omissions and false corrections are both incorrect labels, the former are more deleterious than the latter from a learning perspective; false corrections still reflect the correct task constraints, whereas false omissions can cause a robot to mistakenly interpret failures as correct behaviors. Therefore, to improve task learning, it is reasonable to design robot behaviors to minimize missed corrections, even if it also results in people providing more unnecessary ones.

For **RQ3**, we find that physical effort positively correlates with correction precision, providing empirical evidence to support the assumed trade-off between precision and physical effort (**H3**). However, when supervising an incompetent robot with legible motions, human corrections exhibited a much weaker correlation than when supervising a robot with predictable motions, hinting that learning researchers should rely less on the heuristics that greater effort suggests more precise task correction. A weaker correlation indicates that extra

effort is not necessarily converted to more precise corrections. Therefore, in this condition, solely relying on the task and effort trade-off to model human correction feedback is rather insufficient, and additional features should be explored.

#### A. Alternatives for Trust-Based Hypotheses

Several of our hypotheses (**H1B**, **H2A**, and **H2B**) are founded on the idea that people correct the robot based on their trust in it. Yet, we found strong evidence for the *inverse* of these hypotheses. We now propose an alternative explanation for these findings: that people develop high expectations for a highly competent robot, and intervene in the robot's behavior when there is any indication that the robot may not meet this high expectation. On the other hand, people may be giving an incompetent robot the "benefit of the doubt", waiting until the robot is clearly about to fail before correcting it. Participants being more strict and demanding of highly-competent robots is consistent with previous findings [66], [98].

#### B. Feedback Accuracy Over Time

To assess the potential effect of people "calibrating" to the robot's competency level, we repeated the analysis from Sec. V-B, excluding the first 4, 8, 16, and 32 trials. Further details are provided in Table III in the Appendix. Our analysis confirms the consistency of both human-centric findings across tasks. However, from a robot-centric perspective, no significant differences were observed in whether robots can reliably interpret uncorrected trials as failures later in the experiment, suggesting that this difference is most pronounced when humans are calibrating their expectations of the robot's competency. Additionally, we partitioned the data into the early and late halves of the experiment and conducted a three-way ANOVA on the relevant measures, with the third independent variable representing the data segment (early or late). Results revealed that participants were more likely to miss corrections during the early half of the experiment compared to the later half ( $F(1, 108) = 6.9635, p = 0.0095$ ).

#### C. Takeaway Messages

We summarize key takeaways for two audiences: interaction designers—who program robot interactions to elicit human feedback best aligned with LfC assumptions—and learning researchers—who develop algorithms to enhance learning outcomes, even when feedback diverges from these assumptions.

##### Implications for Interaction Designers:

- A highly competent robot should avoid deviative behavior, while explorative behavior is more acceptable for a robot that makes mistakes.
- Consider how occasionally making deliberate, low-stakes mistakes could reduce human workload from giving frequent, unnecessary corrections.
- To incentivize greater effort and better correction precision, consider making motions predictable rather than legible for an incompetent robot.

##### Implications for Learning Researchers:

- When designing algorithms for an incompetent robot that leverages pre-correction trajectories to learn task constraints, consider weighting them less (due to their potential for misalignment with the task goal) and assigning a higher weight to the correction trajectory itself.
- For active learning researchers, consider allowing more exploratory behavior for a low-competency robot, and less exploration for a highly competent robot.
- When interpreting humans' supervision of a highly competent robot, the lack of a correction should be less reliably interpreted as an endorsement of its behavior.

#### D. Limitations

Despite collecting over 1,950 correction samples during 3,840 pick-and-place trials, our study's overall sample size remains relatively small (60 participants), with only 10 participants per condition. This limited sample size may have reduced the statistical power for some hypotheses (**H1A**, **H2C**) and contributed to the lack of significance observed. For analyses on **RQ1** and **RQ3**, we assume each trial data point from each participant as independent to fully leverage the large volume of correction data. However, this assumption may have inflated the degrees of freedom in the statistical tests, as the data points could exhibit high within-participant correlations.

## VII. CONCLUSION

In this paper, we presented a user study exploring the effects of a robot's competency and motion legibility on how people supervise and correct its behavior. We found that when the robot followed predictable or legible motions, people were more sensitive to potential failures by a highly-competent robot compared to an incompetent robot. Additionally, people were more likely to withhold necessary corrections in low-competency conditions and were more prone to offering unnecessary ones in high-competency conditions. Finally, we empirically supported the assumed trade-off between task precision and human effort when giving corrections. However, we highlighted that this trade-off should be less reliably assumed when modeling humans supervising an incompetent robot with legible motions. These findings offer valuable insights for robot interaction designers and learning researchers working with LfC systems and, more broadly, in robot learning scenarios where humans act in a supervisory role.

## ACKNOWLEDGMENTS

This work is partially supported by the National Science Foundation (NSF IIS-2106690, IIS-1955653) and the Office of Naval Research (ONR N00014-24-1-2124). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF or ONR.

The authors would like to thank Haimanot Belachew, Dražen Brščić, Kate Candon, Jacky Chen, Nicholas C. Georgiou, Ulas Berk Karli, Alexander Lew, Jirachaya "Fern" Limprayoon, Anushka Potdar, Suba Ramesh, Sydney Thompson, the Yale StatLab, and the Cornell EmPRISE Lab for their feedback and contributions.

## REFERENCES

- [1] C. G. Atkeson and S. Schaal, "Robot learning from demonstration," in *ICML*, vol. 97, 1997, pp. 12–20.
- [2] S. Ekvall and D. Kragic, "Robot learning from demonstration: a task-level planning approach," *International Journal of Advanced Robotic Systems*, vol. 5, no. 3, p. 33, 2008.
- [3] S. Schaal, "Learning from demonstration," *Advances in neural information processing systems*, vol. 9, 1996.
- [4] L. Rozo, P. Jiménez, and C. Torras, "A robot learning from demonstration framework to perform force-based manipulation tasks," *Intelligent service robotics*, vol. 6, no. 1, pp. 33–51, 2013.
- [5] M. Arduengo, A. Colomé, J. Lobo-Prat, L. Sentis, and C. Torras, "Gaussian-process-based robot learning from demonstration," *Journal of Ambient Intelligence and Humanized Computing*, pp. 1–14, 2023.
- [6] R. Akrou, M. Schoenauer, and M. Sebag, "Preference-based policy learning," in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2011, Athens, Greece, September 5-9, 2011. Proceedings, Part I 11*. Springer, 2011, pp. 12–27.
- [7] —, "April: Active preference learning-based reinforcement learning," in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2012, Bristol, UK, September 24-28, 2012. Proceedings, Part II 23*. Springer, 2012, pp. 116–131.
- [8] T. Munzer, M. Toussaint, and M. Lopes, "Preference learning on the execution of collaborative human-robot tasks," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 879–885.
- [9] N. Wilde, A. Blidaru, S. L. Smith, and D. Kulić, "Improving user specifications for robot behavior through active preference learning: Framework and evaluation," *The International Journal of Robotics Research*, vol. 39, no. 6, pp. 651–667, 2020.
- [10] X. Wang, K. Lee, K. Hakhamaneshi, P. Abbeel, and M. Laskin, "Skill preferences: Learning to extract and execute robotic skills from human feedback," in *Conference on Robot Learning*. PMLR, 2022, pp. 1259–1268.
- [11] D. J. Hejna III and D. Sadigh, "Few-shot preference learning for human-in-the-loop rl," in *Conference on Robot Learning*. PMLR, 2023, pp. 2014–2025.
- [12] T. K. Faulkner, E. S. Short, and A. L. Thomaz, "Policy shaping with supervisory attention driven exploration," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 842–847.
- [13] D. L. McPherson, D. R. Scobee, J. Menke, A. Y. Yang, and S. S. Sastry, "Modeling supervisor safe sets for improving collaboration in human-robot teams," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 861–868.
- [14] D. Schreckenghost, T. Fong, and T. Milam, "Human supervision of robotic site surveys," in *AIP Conference Proceedings*, vol. 969, no. 1. American Institute of Physics, 2008, pp. 776–783.
- [15] S. Lang, M. Kleinhagenbrock, S. Hohenner, J. Fritsch, G. A. Fink, and G. Sagerer, "Providing the basis for human-robot-interaction: A multimodal attention system for a mobile robot," in *Proceedings of the 5th international conference on Multimodal interfaces*, 2003, pp. 28–35.
- [16] A. Bajcsy, D. P. Losey, M. K. O'Malley, and A. D. Dragan, "Learning robot objectives from physical human interaction," in *Conference on Robot Learning*. PMLR, 2017, pp. 217–226.
- [17] A. Bajcsy, D. P. Losey, M. K. O'Malley, and A. D. Dragan, "Learning from physical human corrections, one feature at a time," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 141–149.
- [18] D. P. Losey and M. K. O'Malley, "Including uncertainty when learning from human corrections," in *Conference on Robot Learning*. PMLR, 2018, pp. 123–132.
- [19] A. Bobu, D. R. Scobee, J. F. Fisac, S. S. Sastry, and A. D. Dragan, "Less is more: Rethinking probabilistic models of human behavior," in *Proceedings of the 2020 acm/ieee international conference on human-robot interaction*, 2020, pp. 429–437.
- [20] A. Bobu, A. Bajcsy, J. F. Fisac, S. Deglurkar, and A. D. Dragan, "Quantifying hypothesis space misspecification in learning from human-robot demonstrations and physical corrections," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 835–854, 2020.
- [21] D. P. Losey, A. Bajcsy, M. K. O'Malley, and A. D. Dragan, "Physical interaction as communication: Learning robot objectives online from human corrections," *The International Journal of Robotics Research*, vol. 41, no. 1, pp. 20–44, 2022.
- [22] E. Chisari, T. Welschehold, J. Boedecker, W. Burgard, and A. Valada, "Correct me if i am wrong: Interactive learning for robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3695–3702, 2022.
- [23] A. Wang and T. Fitzgerald, "Legibility-aware learning from corrections," 2024.
- [24] T. Fitzgerald, A. Goel, and A. Thomaz, "Modeling and learning constraints for creative tool use," *Frontiers in Robotics and AI*, vol. 8, p. 674292, 2021.
- [25] T. Fitzgerald, E. Short, A. Goel, and A. Thomaz, "Human-guided trajectory adaptation for tool transfer," in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS '19. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2019, p. 1350–1358.
- [26] A. Jevtić, A. Colomé, G. Alenya, and C. Torras, "Robot motion adaptation through user intervention and reinforcement learning," *Pattern Recognition Letters*, vol. 105, pp. 67–75, 2018.
- [27] J. Spencer, S. Choudhury, M. Barnes, M. Schmittle, M. Chiang, P. Ramadge, and S. Srinivasa, "Learning from interventions," in *Robotics: Science and Systems (RSS)*, 2020.
- [28] —, "Expert intervention learning: An online framework for robot learning from explicit and implicit human feedback," *Autonomous Robots*, pp. 1–15, 2022.
- [29] H. Liu, A. Chen, Y. Zhu, A. Swaminathan, A. Kolobov, and C.-A. Cheng, "Interactive robot learning from verbal correction," *arXiv preprint arXiv:2310.17555*, 2023.
- [30] S. Wang, B. Scassellati, and T. Fitzgerald, "Toward measuring the effect of robot competency on human kinesthetic feedback in long-term task learning," *4th Workshop on Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI)*, 2024.
- [31] T. Sheridan, "Human supervisory control of robot systems," in *Proceedings. 1986 IEEE International Conference on Robotics and Automation*, vol. 3. IEEE, 1986, pp. 808–812.
- [32] A. Mohtasib, G. Neumann, and H. Cuayáhuil, "A study on dense and sparse (visual) rewards in robot policy learning," in *Towards Autonomous Robotic Systems: 22nd Annual Conference, TAROS 2021, Lincoln, UK, September 8–10, 2021, Proceedings 22*. Springer, 2021, pp. 3–13.
- [33] S. Lemaignan, F. Garcia, A. Jacq, and P. Dillenbourg, "From real-time attention assessment to "with-me-ness" in human-robot interaction," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. Ieee, 2016, pp. 157–164.
- [34] Y. Chen, C. Liu, B. E. Shi, and M. Liu, "Robot navigation in crowds by graph convolutional networks with attention learned from human gaze," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2754–2761, 2020.
- [35] P. Valletta, R. Pérez-Dattari, and J. Kober, "Imitation learning with inconsistent demonstrations through uncertainty-based data manipulation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 3655–3661.
- [36] Z. Qian, M. You, H. Zhou, X. Xu, and B. He, "Robot learning from human demonstrations with inconsistent contexts," *Robotics and Autonomous Systems*, vol. 166, p. 104466, 2023.
- [37] W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The tamer framework," in *Proceedings of the fifth international conference on Knowledge capture*, 2009, pp. 9–16.
- [38] C. Daniel, M. Viering, J. Metz, O. Kroemer, and J. Peters, "Active reward learning," in *Robotics: Science and systems*, vol. 98, 2014.
- [39] S. Chernova and A. L. Thomaz, *Robot learning from human teachers*. Morgan & Claypool Publishers, 2014.
- [40] Y. Cui, Q. Zhang, B. Knox, A. Allievi, P. Stone, and S. Niekum, "The empathic framework for task learning from implicit human feedback," in *Conference on Robot Learning*. PMLR, 2021, pp. 604–626.
- [41] R. Loftin, B. Peng, J. MacGlashan, M. L. Littman, M. E. Taylor, J. Huang, and D. L. Roberts, "Learning something from nothing: Leveraging implicit human feedback strategies," in *The 23rd IEEE international symposium on robot and human interactive communication*. IEEE, 2014, pp. 607–612.
- [42] —, "Learning behaviors via human-delivered discrete feedback: modeling implicit feedback strategies to speed up learning," *Autonomous agents and multi-agent systems*, vol. 30, pp. 30–59, 2016.
- [43] S. Lauria, G. Bugmann, T. Kyriacou, J. Bos, and E. Klein, "Personal robot training via natural-language instructions," *IEEE Intelligent systems*, vol. 16, no. 3, pp. 38–45, 2001.

- [44] C. Matuszek, E. Herbst, L. Zettlemoyer, and D. Fox, "Learning to parse natural language commands to a robot control system," in *Experimental robotics: the 13th international symposium on experimental robotics*. Springer, 2013, pp. 403–415.
- [45] J. Thomason, S. Zhang, R. J. Mooney, and P. Stone, "Learning to interpret natural language commands through human-robot dialog," in *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [46] S. Stepputtis, J. Campbell, M. Phielipp, S. Lee, C. Baral, and H. Ben Amor, "Language-conditioned imitation learning for robot manipulation tasks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 13 139–13 150, 2020.
- [47] Y. Cui, P. Koppol, H. Admoni, S. Niekum, R. Simmons, A. Steinfeld, and T. Fitzgerald, "Understanding the relationship between interactions and outcomes in human-in-the-loop machine learning," in *International Joint Conference on Artificial Intelligence*, 2021.
- [48] A. L. Thomaz, C. Breazeal *et al.*, "Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance," in *Aaai*, vol. 6. Boston, MA, 2006, pp. 1000–1005.
- [49] A. L. Thomaz and C. Breazeal, "Teachable robots: Understanding human teaching behavior to build more effective robot learners," *Artificial Intelligence*, vol. 172, no. 6-7, pp. 716–737, 2008.
- [50] N. Navarro-Guerrero, R. J. Lowe, and S. Wermter, "Improving robot motor learning with negatively valenced reinforcement signals," *Frontiers in neurobotics*, vol. 11, p. 10, 2017.
- [51] T. Fitzgerald, P. Koppol, P. Callaghan, R. Q. J. H. Wong, R. Simmons, O. Kroemer, and H. Admoni, "Inquire: Interactive querying for user-aware informative reasoning," in *6th Annual Conference on Robot Learning*, 2022.
- [52] D. Ramachandran and E. Amir, "Bayesian inverse reinforcement learning," in *IJCAI*, vol. 7, 2007, pp. 2586–2591.
- [53] P. Carreno-Medrano, S. L. Smith, and D. Kulić, "Joint estimation of expertise and reward preferences from human demonstrations," *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 681–698, 2022.
- [54] B. D. Ziebart, A. L. Maas, J. A. Bagnell, A. K. Dey *et al.*, "Maximum entropy inverse reinforcement learning," in *Aaai*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.
- [55] M. Nourani, S. Kabir, S. Mohseni, and E. D. Ragan, "The effects of meaningful and meaningless explanations on trust and perceived system accuracy in intelligent systems," in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, vol. 7, 2019, pp. 97–105.
- [56] K. A. Tecimer, E. Tüzün, H. Dibeklioglu, and H. Erdogmus, "Detection and elimination of systematic labeling bias in code reviewer recommendation systems," in *Proceedings of the 25th International Conference on Evaluation and Assessment in Software Engineering*, 2021, pp. 181–190.
- [57] A. Gupta, A. Murali, D. P. Gandhi, and L. Pinto, "Robot learning in homes: Improving generalization and reducing dataset bias," *Advances in neural information processing systems*, vol. 31, 2018.
- [58] J. Wei, Z. Zhu, H. Cheng, T. Liu, G. Niu, and Y. Liu, "Learning with noisy labels revisited: A study using real-world human annotations," *arXiv preprint arXiv:2110.12088*, 2021.
- [59] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," in *International conference on machine learning*. PMLR, 2016, pp. 49–58.
- [60] M. Bloem and N. Bambos, "Infinite time horizon maximum causal entropy inverse reinforcement learning," in *53rd IEEE conference on decision and control*. IEEE, 2014, pp. 4911–4916.
- [61] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2013, pp. 301–308.
- [62] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: a survey," in *2020 IEEE symposium series on computational intelligence (SSCI)*. IEEE, 2020, pp. 737–744.
- [63] A. Afzal, C. Le Goues, M. Hilton, and C. S. Timperley, "A study on challenges of testing robotic systems," in *2020 IEEE 13th International Conference on Software Testing, Validation and Verification (ICST)*. IEEE, 2020, pp. 96–107.
- [64] S. M. Neuman, B. Plancher, B. P. Duisterhof, S. Krishnan, C. Banbury, M. Mazumder, S. Prakash, J. Jabbour, A. Faust, G. C. de Croon *et al.*, "Tiny robot learning: Challenges and directions for machine learning in resource-constrained robots," in *2022 IEEE 4th International Conference on Artificial Intelligence Circuits and Systems (AICAS)*. IEEE, 2022, pp. 296–299.
- [65] P. A. Hancock, D. R. Billings, K. E. Schaefer, J. Y. Chen, E. J. De Visser, and R. Parasuraman, "A meta-analysis of factors affecting trust in human-robot interaction," *Human factors*, vol. 53, no. 5, pp. 517–527, 2011.
- [66] S. Paepcke and L. Takayama, "Judging a bot by its cover: An experiment on expectation setting for personal robots," in *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2010, pp. 45–52.
- [67] M. M. Scheunemann, R. H. Cuijpers, and C. Salge, "Warmth and competence to predict human preference of robot behavior in physical human-robot interaction," in *2020 29th IEEE international conference on robot and human interactive communication (RO-MAN)*. IEEE, 2020, pp. 1340–1347.
- [68] J. DelPreto, J. I. Lipton, L. Sanneman, A. J. Fay, C. Fourie, C. Choi, and D. Rus, "Helping robots learn: a human-robot master-apprentice model using demonstrations via virtual reality teleoperation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 10 226–10 233.
- [69] E. Hedlund, M. Johnson, and M. Gombolay, "The effects of a robot's performance on human teachers for learning from demonstration tasks," in *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 2021, pp. 207–215.
- [70] A. Dragan and S. Srinivasa, "Generating legible motion," 2013.
- [71] F. Stulp, J. Grizou, B. Busch, and M. Lopes, "Facilitating intention prediction for humans by optimizing robot motions," in *2015 IEEE/RSJ international conference on intelligent robots and systems (iros)*. IEEE, 2015, pp. 1249–1255.
- [72] D. Sadigh, A. D. Dragan, S. Sastry, and S. A. Seshia, *Active preference-based learning of reward functions*, 2017.
- [73] H. Hu and J. F. Fisac, "Active uncertainty reduction for human-robot interaction: An implicit dual control approach," in *International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2022, pp. 385–401.
- [74] S. Wang, Y. Lyu, and J. M. Dolan, "Active probing and influencing human behaviors via autonomous agents," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1514–1521.
- [75] C. Lichtenthäler, T. Lorenzy, and A. Kirsch, "Influence of legibility on perceived safety in a virtual human-robot path crossing task," in *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 2012, pp. 676–681.
- [76] N. J. Hetherington, E. A. Croft, and H. M. Van der Loos, "Hey robot, which way are you going? nonverbal motion legibility cues for human-robot spatial interaction," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5010–5015, 2021.
- [77] A. D. Dragan, S. Bauman, J. Forlizzi, and S. S. Srinivasa, "Effects of robot motion on human-robot collaboration," in *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*, 2015, pp. 51–58.
- [78] S. Kim and W. Park, "A review on robot motion legibility for human-robot interaction," *Korean Society of Ergonomics Conference*, pp. 220–228, 2017.
- [79] B. Busch, J. Grizou, M. Lopes, and F. Stulp, "Learning legible motion from human-robot interactions," *International Journal of Social Robotics*, vol. 9, no. 5, pp. 765–779, 2017.
- [80] I. A. Şucan, M. Moll, and L. E. Kavraki, "The Open Motion Planning Library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, December 2012, <https://ompl.kavrakilab.org>.
- [81] M. J. Kim, F. Beck, C. Ott, and A. Albu-Schäffer, "Model-free friction observers for flexible joint robots with torque measurements," *IEEE Transactions on Robotics*, vol. 35, no. 6, pp. 1508–1515, 2019.
- [82] R. K. Jenamani, D. Stabile, Z. Liu, A. Anwar, K. Dimitropoulou, and T. Bhattacharjee, "Feel the bite: Robot-assisted inside-mouth bite transfer using robust mouth perception and physical interaction-aware control," in *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 2024, pp. 313–322.
- [83] J. Carpentier and N. Mansard, "Analytical derivatives of rigid body dynamics algorithms," in *Robotics: Science and Systems (RSS)*, 2018.
- [84] J. Brooke, "Sus: A quick and dirty usability scale," *Usability Evaluation in Industry/Taylor and Francis*, 1996.
- [85] J. G. Trafton, C. R. Frazier, K. Zish, B. J. Bio, and J. M. McCurry, "The perception of agency: Scale reduction and construct validity," in *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2023, pp. 936–942.

- [86] S. G. Hart, "Nasa-task load index (nasa-tlx); 20 years later," in *Proceedings of the human factors and ergonomics society annual meeting*, vol. 50, no. 9. Sage publications Sage CA: Los Angeles, CA, 2006, pp. 904–908.
- [87] M. Körber, "Theoretical considerations and development of a questionnaire to measure trust in automation," in *Proceedings of the 20th Congress of the International Ergonomics Association (IEA 2018) Volume VI: Transport Ergonomics and Human Factors (TEHF), Aerospace Human Factors and Ergonomics 20*. Springer, 2019, pp. 13–30.
- [88] S. Kullback and R. A. Leibler, "On information and sufficiency," *The annals of mathematical statistics*, vol. 22, no. 1, pp. 79–86, 1951.
- [89] C. Laidlaw and A. Dragan, "The boltzmann policy distribution: Accounting for systematic suboptimality in human models," *arXiv preprint arXiv:2204.10759*, 2022.
- [90] S. V. Stehman, "Selecting and interpreting measures of thematic classification accuracy," *Remote sensing of Environment*, vol. 62, no. 1, pp. 77–89, 1997.
- [91] R. A. Fisher, "Statistical methods for research workers," in *Breakthroughs in statistics: Methodology and distribution*. Springer, 1970, pp. 66–70.
- [92] M. J. Blanca Mena, R. Alarcón Postigo, J. Arnau Gras, R. Bono Cabré, and R. Bendayan, "Non-normal data: Is anova still a valid option?" *Psicothema*, 2017, vol. 29, num. 4, p. 552-557, 2017.
- [93] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the Royal statistical society: series B (Methodological)*, vol. 57, no. 1, pp. 289–300, 1995.
- [94] S. S. Shapiro and M. B. Wilk, "An analysis of variance test for normality (complete samples)," *Biometrika*, vol. 52, no. 3-4, pp. 591–611, 1965.
- [95] C. Spearman, "The proof and measurement of association between two things." 1961.
- [96] R. A. Fisher, "Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population," *Biometrika*, vol. 10, no. 4, pp. 507–521, 1915.
- [97] R. A. Fisher *et al.*, "014: On the" probable error" of a coefficient of correlation deduced from a small sample." 1921.
- [98] D. Belanche, L. V. Casaló, J. Schepers, and C. Flavián, "Examining the effects of robots' physical appearance, warmth, and competence in frontline services: The humanness-value-loyalty model," *Psychology & Marketing*, vol. 38, no. 12, pp. 2357–2376, 2021.