RoboTeach: How Student Robots' Preexisting Proficiency and Learning Rate Affect Human Teachers Demonstrating Object Placement

Khaled Kassem
TU Wien
Vienna, Austria
khaled.k.kassem@tuwien.ac.at

Patrick Gietl
TU Wien
Vienna, Austria
pat.gietl@gmail.com

Florian Michahelles
TU Wien
Visual Computing & Human-centered Technology
Vienna, Austria
florian.michahelles@tuwien.ac.at

Andrii Matviienko
KTH Royal Institute of Technology
Stockholm, Sweden
andriim@kth.se

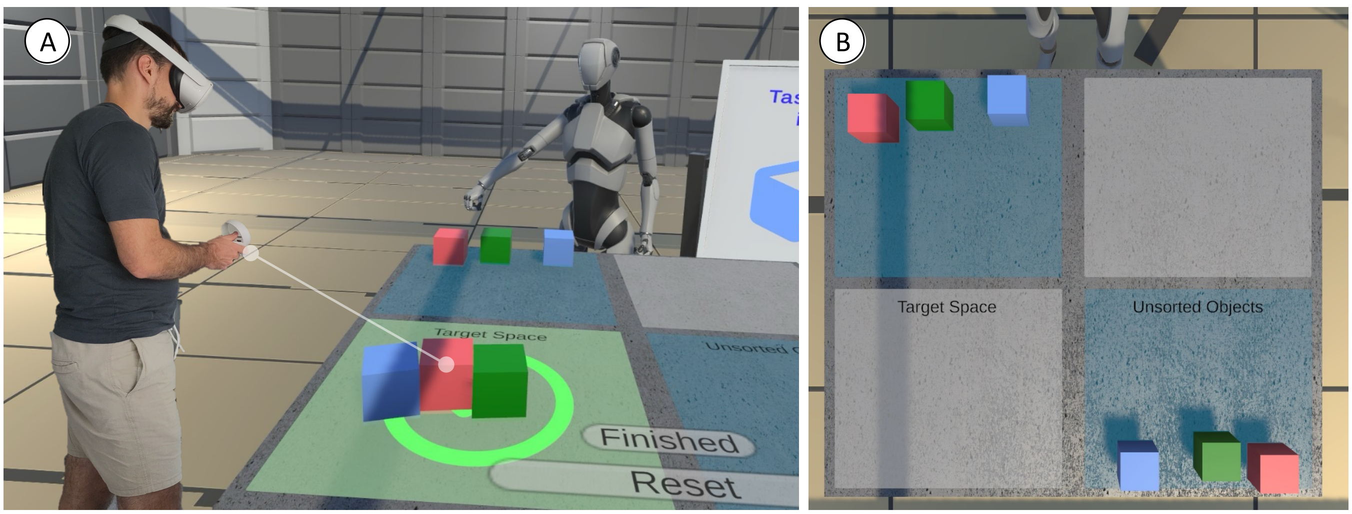

Figure 1: Teaching a robot the task of placing colored cubes in the right order and position: (A) a participant is teaching the robot how place a cube at the right location and (B) overview of the interaction surface from above.

Abstract

Social robots are employed as companions, helping in industrial and domestic environments. Adapting robots' capabilities to user needs can be achieved through teaching from human demonstrations. However, the influence of robots' preexisting proficiency and learning rate on human teachers' self-efficacy and perception of the robots is underexplored. In this paper, we simulated four robot performance types that combine: (1) preexisting proficiency (low/high) and (2) learning rate (slow/fast). We conducted a controlled lab experiment studying the impact of robots' performance type on teachers' self-efficacy, willingness to teach the robot, and perception of the robot (N=24), in which robots placed objects in suitable locations. Fast learners were perceived as more intelligent, anthropomorphic, and likable, and this caused higher teaching self-efficacy regardless of preexisting skills. Slow learners caused frustration while teaching. Moreover, participants stopped teaching robots with low preexisting skills sooner, regardless of the learning rate, indicating potential bias caused by expectations.

Keywords

teaching robots, object placement, learning rate, existing proficiency, self-efficacy, robot perception

CHI '25, Yokohama, Japan

ACM ISBN 979-8-4007-1394-1/25/04
https://doi.org/10.1145/3706598.3713113

ACM Reference Format:
Khaled Kassem, Patrick Gietl, Florian Michahelles, and Andrii Matviienko. 2025. RoboTeach: How Student Robots' Preexisting Proficiency and Learning Rate Affect Human Teachers Demonstrating Object Placement. In *CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3706598.3713113

1 Introduction

Commercial social robots are often employed as human companions and helpers, primarily in work and home contexts [2, 22, 34, 39]. These robots typically have limited proficiency and can complete a finite set of pre-programmed tasks. For example, a household robot assisting with cooking does not accommodate each potential user's individual preferences and taste palates. It becomes unfeasible to pre-program such robots to suit all users' needs, and successfully programmed conditions for a task may change with time and need to be updated [34]. To overcome this limitation, users must reprogram the robots and, therefore, "teach" how to perform customized tasks. However, since robot programming is not a trivial task [39], and not all users are equally tech-savvy, there is a need to introduce new methods to customize robots and accommodate various human needs [41], which we explore within the scope of this paper.

Previous work has explored different strategies to customize robots, using methods such as reprogramming, downloading new skills, reinforcement learning, and Learning from Demonstration (LfD) [2, 21, 34]. LfD (or imitation learning) involves "teaching" a robot based on how a human would accomplish a certain task. When teaching robot students, teachers experience different perceived self-efficacy and perceptions of the robot based on the robot learner's errors [33]. Compared to teaching robots, it has been shown that teachers' perception of, and trust in, the robots [2], and the human's impression of themselves [21] is largely influenced by the robots' learning process and its outcomes (e.g., errors). Yet, we do not fully understand how the perceived robots' preexisting proficiency at different tasks before teaching affects human teachers' efficacy while teaching robots new tasks or preferences. Moreover, while previous work has explored the impact of robot errors during the teaching process, the effect of the robot's learning rate on the human teacher is less understood.

In this paper, we investigate how different levels of robot preexisting proficiency and rate of learning a new skill influence human teachers' perception of the robot and sense of self-efficacy at teaching. For this, we designed an experiment based on two independent variables: (1) the robot's preexisting proficiency at a secondary task (low/high) and (2) the robot's learning rate of the task being taught (slow/fast). To indicate robot preexisting proficiency, robots demonstrated their existing skill of drawing to each participant before a teaching session. The learning rate demonstrated how slow or fast a robot can learn a new skill of object placement from participants. To systematically investigate how robot preexisting proficiency and learning rate influence human teachers, we conducted a controlled lab experiment in virtual reality environments ($N = 24$), in which participants had to teach a robot the correct placement of colored cubes on a work desk. Our results indicate that participants spent more time and made more attempts to teach slow-learning robots than fast-learning robots. Moreover, robots achieved lower proficiency with low existing skills than higher, and slow learners achieved lower proficiency than fast. Lastly, fast-learning robots with high existing skills reached higher proficiency levels than fast-learning robots with low existing skills.

Contribution Statement

This work studies the effect of a robot student's preexisting proficiency and learning rate on the human teacher while teaching the robot a new task. We contribute an empirical evaluation regarding how the teacher's self-efficacy and perception of the robot are influenced by the robot's shown initial proficiency at a different task and the robot's learning rate while teaching it a new task.

2 Related Work

In this section, we provide a summary of existing work on teaching robots. We focus on teaching robots through LfD, teaching environments, and the associated measures.

2.1 Humans Teaching Robots

Existing Human-Computer Interaction (HCI) and Human-Robot Interaction (HRI) work has shown that programming robots are not a trivial task [9, 26, 39]. Future robots employed in situations that need customization and skill uptake should, therefore, not rely on programming as the only method of enhancing the skill of the robots. Instead, it should be possible to transfer knowledge to the robot by any human who possesses this knowledge [26]. One possible solution to this problem is to have robots that are capable of learning from human input [1, 7, 25, 31, 38, 39], especially in contexts that rely on robots working in close proximity to humans [15, 22, 36].

Aliasghari et al. [2] examined the effect of a student robot's appearance and errors with multiple severities on teachers' trust in the robot and their future expectations. Mirnig et al. [33] similarly investigated the effects of robot errors on the perception of the robot. Other work examined the impact of a robot's learning method (download, reinforcement learning, human-in-the-loop reinforcement learning, and LfD) on the perception and perceived safety of the robot [34]. Calinon and Billard [7] explored how teachers' incremental refinement of the robot's skill by moving its arms manually provides the appropriate scaffolds to reproduce the action. Thomaz and Cakmak [45] explored different ways that a human partner can intuitively help the robot learn. Libera et al. [31] developed a system to edit the motions of a small humanoid robot by touching its body parts. Cakmak and Thomaz [6] investigated ways to guide humans in the teacher's role to teach more optimally.

Early work by Friedman [16] proposed a model where robots compare predicted outcomes with actual experiences to refine their perceptions and predictive abilities. More recent studies have also investigated how children perceive learning robots in educational settings. For example, Chandra et al. [8] found that children's perception of a robot's learning capabilities can affect their own learning gains in a handwriting task. The concept of dual learning has been explored by Kubota et al. [27], where robots simultaneously acquire perceptual and behavioral skills. This approach recognizes the interdependence of perception and action in robotic learning. These studies collectively highlight the complex relationship between robot perception, learning, and human interaction. Therefore, in this paper, we build on the practices from previous work that involve an iterative demonstration, i.e., learning by demonstration from the robot's perspective, which facilitates learning of object placement in the right place and order.

2.2 How Student Robots Affect Human Teachers

Previous work has shown that human teachers tend to adapt their teaching style to try to suit the robot's [43, 44] performance. Thomaz and Breazeal [44] have shown that transparency of the robot's internal state helps improve human guidance. Aliasghari et al. [2] showed that a robot's appearance affects perception of the robot, but not necessarily trust. However, even a small error could significantly reduce trust in a trainee robot performing a task regardless of the robot's appearance. Zafari et al. [50] showed that a student robot's method of communication affects the teacher's self-efficacy, as defined by Bandura [3]. Moorman et al. [34] found that the robot's learning method impacts the perceived anthropomorphism of the student robot. Robot failure in learning from demonstration tasks negatively impacts human teachers' trust, self-confidence, and impression of the robot [21]. Contrary to these findings, Mirnig et al. [33] showed in their work that robot errors do not affect perceived intelligence, but increase likeability instead.

While any learning algorithms is expected to make errors, previous work has not fully explored the impact of the robot student's learning rate on the human teacher. Further, the effect of expectations created by a robot student's preexisting proficiency on the human teacher's method of teaching and willingness to teach is not fully understood. This is the gap our work means to address.

2.3 Methods in HRI Studies on Humans Teaching Robots

Previous work examining the process of teaching robots takes place in different contexts. Studies have been conducted with physical robots in-situ [11], with interactive remote robots [2], as well as in Virtual Reality (VR) [13, 14, 32, 47]. The context of these user studies could be industrial [47], domestic (e.g. with cooking [2] or healthcare [34]), or more social [11, 23]. These studies were interested in a number of subjective and objective measures. Objective measures included robot performance [14]. Meanwhile, subjective measures included perception of the robot [2, 11, 34, 48], perceived safety [14, 32], user experience [47], and self-efficacy [50]. Researchers have investigated applying psychological and educational research methodologies to evaluate robot performance in classroom settings [46]. Typically non-experts naturally teach robots, revealing diverse teaching styles that can be categorized based on interaction types, testing patterns, and lesson organization [25]. Moreover, humans tend to use multiple teaching methods when instructing robots on social and moral norms, adapting their approach based on the robot's performance and task difficulty [10]. These findings contribute to a better understanding of human teaching patterns and offer insights for designing more effective human-robot teaching protocols. Although a teaching context might play an essential role in studying humans teaching robots, in this work, we focus on one abstract interaction context, in which a human and a robot are co-located in the same virtual space with a table between them and objects necessary to learn new skills. However, based on our knowledge, there is not much empirical evidence showing how student robots' preexisting proficiency and learning rate affect human teachers, which we explore within the scope of this paper.

3 Evaluation

To evaluate how a student robot's preexisting skill level and learning rate influence human perception of the robot, the self-efficacy of the human instructors, and their willingness to continue teaching the robot, we conducted a controlled lab experiment with the following research question: *While teaching the robot a new skill, how do the robot's preexisting proficiency at a different skill and the rate at which it learns the new skill affect the human teacher's self-efficacy, perception of the robot, and willingness to teach?*

Before conducting the study, it is reasonable to assume that a fast-learning robot would be favored more than a slow-learning one. However, since we expect the learning behavior of future commercial robots to be uncontrollable by their end users, it is also important to study how a potentially slow-learning robot would be perceived. Additionally, how the robot's learning rate interacts with other factors that might influence this perception, e.g., preexisting proficiency, is not clearly defined.

3.1 Study Design

To answer this question, we designed a within-subject study with two independent variables: (1) preexisting proficiency and (2) learning rate. *Preexisting proficiency* represents the ability level the robot could demonstrate at a secondary skill different from the one participants teach the robot. This variable has two levels (low/high). The secondary preexisting skill (drawing a square in midair) is different from the new skill the robot will be taught (placement of colored cubes on a flat surface). We chose a midair drawing task to demonstrate preexisting skill as it shares some aspects of the cube placement task, while remaining different in nature. We reason that placing cubes on a flat surface and drawing in midair both involve 2D spatial reasoning, which is needed to correctly place the cubes and the vertices of a square, respectively. This allows some commonality between the tasks while not completely overlapping.

The second variable (*"learning rate"*) represents the rate at which the robot acquires the new target skill (placing a set of colored cubes in the correct order and position). Learning Rate has two (slow/fast) levels. This simulates the learning progression rate. We chose cube placement as the task to teach the robot for two reasons. First, pick-and-place is a common task in HRI user studies [30], and second, this can be an analog to some generic household tasks that a companion robot could be asked to help with, e.g., putting away the dishes. We combined all levels of independent variables to systematically investigate their effects, which resulted in a $2x2$ within-subject study design and the following four experimental conditions:

(1) LS: Low preexisting proficiency/ Slow learning rate
(2) LF: Low preexisting proficiency/ Fast learning rate
(3) HS: High preexisting proficiency/ Slow learning rate
(4) HF: High preexisting proficiency/ Fast learning rate

Participants' task was to teach the robot a new skill, which we chose to be arranging colored (Red, Green, and Blue) cubes in a predefined order and placement, similar to object manipulation tasks in literature [34]. To simulate the learning process, we pre-programmed a set of positions for the cubes that correspond to different types and severity of errors. As participants "teach" the robot, it progresses through the pre-designed placements in ascending

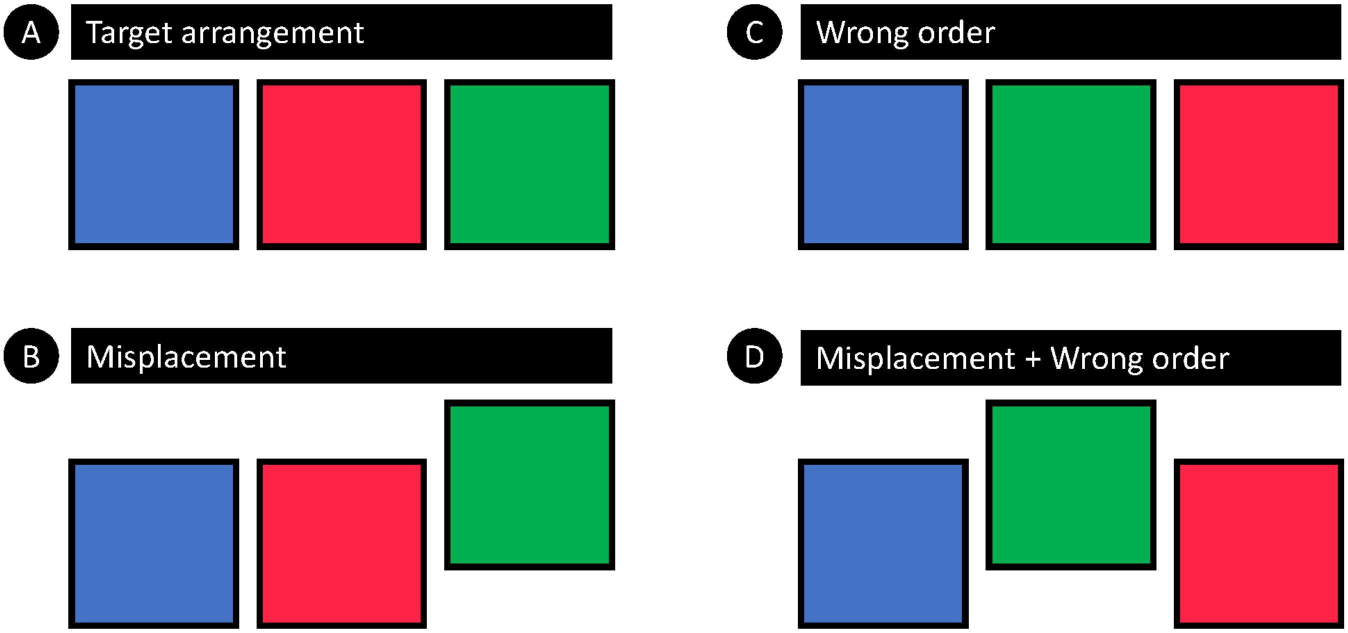

Figure 2: (A): Target order and placement of colored cubes taught to the robot. Three types of errors that the robot makes during the learning process: (B) misplacement, (C) wrong order, and (D) misplacement + wrong order.

levels of correctness until it reaches full proficiency, at which point it perfectly mimics the placement as the participant instructs. The robot could make two types of errors, deviating from the pattern it is taught: arranging the cubes in the wrong order (e.g., Red-Green-Blue when it was instructed with Blue-Red-Green) and placing the cubes incorrectly in 2D space (i.e., such that the position of each cube relative to the others is inaccurate). In addition, the robot can combine both error types (Figure 2). The rotation of the cubes did not matter. This means that cubes would still have been considered in the correct arrangement even with some side misalignment.

The robot's task proficiency could be implicitly represented as a percentage, ranging from 0% (no proficiency) to 100% (full mastery), increased after each valid demonstration depending on its learning speed: slow learners increased their proficiency by 10% per correct human demonstration, requiring 10 iterations to reach full mastery. Fast learners increased their proficiency by 33.34%, requiring three iterations to mastery. The robot attempted the task after each demonstration, and its performance reflected its proficiency level: below 33.34%, it made combined mistakes (misplacement + wrong order); between 33.34% and 66.68%, it made only wrong-order mistakes; and between 66.68% and 100%, it made only misplacement mistakes, with the degree of misplacement becoming smaller as proficiency increased. At 100% proficiency, the robot made no mistakes. These proficiency step sizes were designed to clearly distinguish slow- and fast-learning while keeping the study session at a reasonable length. The possible types of mistakes the robot could make while learning are shown in Figure 2.

3.2 Apparatus

We built the study environment in Unity 2022.3.21f1. To represent the robot, we used the premade default 3D model in the Unity 3rd-person starter asset package (Figure 4). The virtual environment resembled a room with approximate dimensions W x H x D = 14.1m x 6m x 24m. The interaction with the task and cubes took place on a table-like platform with measurements W x H x D = 1.4m x 0.85m x 1.2m. The cubes had a side length of $0.1m$. The VR environment and interaction surface can be seen in Figures 1 and 4. Participants experienced the study environment through the Meta Quest 2 head-mounted display and interacted with the environment using the hand-held controller.

3.3 Measures

To compare all levels of preexisting proficiency and learning rates, we measured the following dependant variables:

- **Teaching time, sec**: we measured the amount of time invested into teaching a robot.
- **Number of teaching attempts**: we counted the number of attempts to teach a robot to a proficient level.
- **Achieved proficiency**: we noted the achieved proficiency of a robot on the scale from 0 to 1.0 with a step of 0.1 when participants decided to stop the teaching process.
- **Initial observing time, sec**: we measured the amount of time participants used to observe a robot performing the shape-drawing task before teaching began. While unrelated to teaching, we put this measure in place to gauge whether or not each participant paid attention to the robot's preexisting skill, the effect of which is core to our research question.

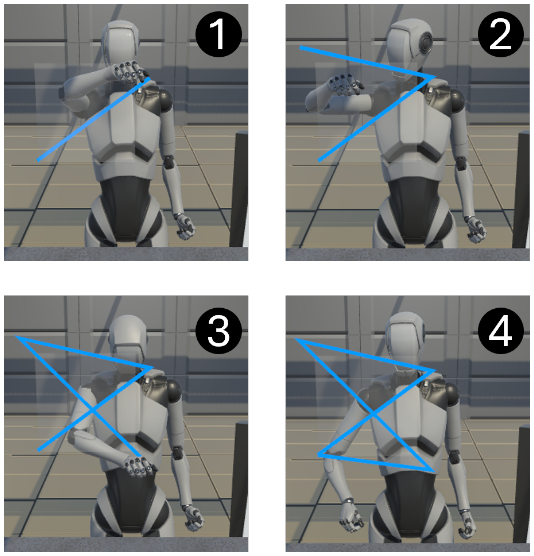

(a)

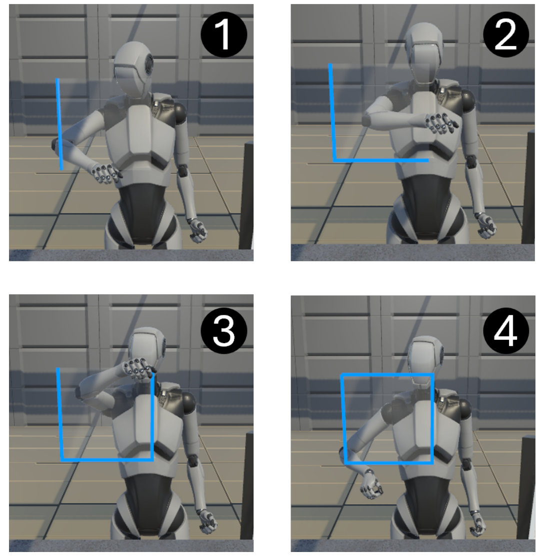

(b)

Figure 3: Robots in our study demonstrated their preexisting proficiency at a secondary task by drawing a square in midair. a) robot with low preexisting proficiency fails to properly draw a square. b) robot with high preexisting proficiency correctly draws a square.

- **Robot perception**: we assessed how participants perceived a robot in terms of intelligence, safety, likeability, anthropomorphism, and animacy, based on the Godspeed series [4].
- **Teaching self-efficacy**: after each teaching session, participants indicated their self-efficacy concerning their ability to teach and interact with the robot using items 2, 8, and 10 extracted from the Self-efficacy in HRI scale [35].
- **Teaching experience**: after each teaching session, participants indicated their experience teaching a robot using a UEQ-S [40].
- **Willingness to teach**: after each teaching session, participants assessed their desire to continue the teaching session or engage with a robot.

3.4 Participants

We recruited 24 participants (self-identified: 11 M, 13 F) aged between 22 and 40 ($M = 29.54, SD = 4.14$) through a convenience sampling and university mailing lists. Their experience levels with teaching included teaching children outside of school (e.g. summer camp) (3), workplace training (1), school teacher/university tutor (11), and none (9). Participants rated their familiarity with computer science, robotics, and AI, on a scale of 1 (novice) to 5 (expert) ($M = 2.54, SD = 1.22$). They also responded to the Ten-item Personality Inventory (TIPI) [19], Extraversion ($M = 4.44, SD = 1.28$), Agreeableness ($M = 4.83, SD = 1.00$), Conscientiousness ($M = 4.88, SD = 1.27$), Emotional Stability ($M = 4.96, SD = 1.00$), and Openness to Experience ($M = 5.54, SD = 0.85$). One participant reported some feelings of motion sickness after finishing the study. All participants had normal or corrected-to-normal vision.

3.5 Procedure

Participants were briefed about the study environment and task of teaching the robot to place colored cubes. After providing their informed consent to take part in the study, participants filled out the onboarding questionnaire which captured their demographic information, computer science/robotics familiarity, teaching competencies, and personality information. We then provided a guided tutorial to familiarize participants with the VR environment and the basics of interacting with the robot. Participants are then told that they will be teaching a task to four different robots that visually look identical but represent different learning algorithms. For each of the four conditions, the interaction begins with the robot waving at the participants to request their attention, followed by preexisting proficiency demonstration where the robot displays its preexisting proficiency in an unrelated skill by attempting to draw a square in midair. Participants are told that the robot was taught this skill previously, reaching its current exhibited skill level. This happens once. Participants were not explicitly told anything else about the robot's capabilities. The rest of the interaction followed these steps:

(1) **Teaching session**: Each teaching session starts with an identical setup, with the cubes initially placed in the same positions (the area marked by "unsorted objects" in Figure 1.B) ready to be picked up by participants. Participants are asked to teach the robot to place three cubes in a specific order. They are instructed to pick up the cubes using the handheld controller and to place them in a "target" quadrant on top of a work bench in the VR environment (Figure

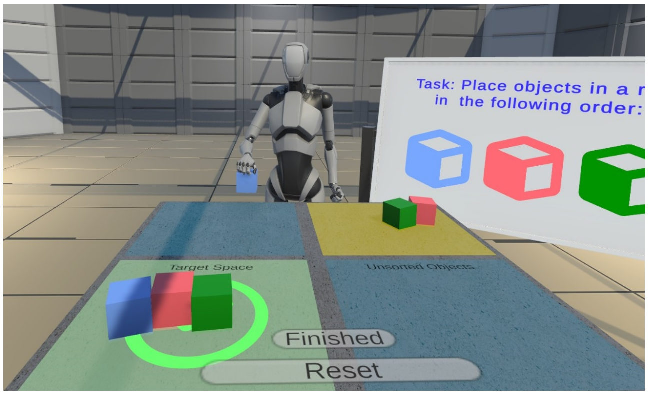

Figure 4: Study setup in VR: Participants stood in front of a virtual desk in front of a robot that demonstrated to them its preexisting skills and its intermediary progress of learning a new skill.

4). This step ends only when the participant uses the hand-held controller to tap the *finished* button, signifying that the correct arrangement was fully demonstrated.

(2) **Learning progress demonstration**: the robot demonstrates its current learning progress by attempting to repeat the cube placement task. This is when participants can see the progress made by the robot in learning.

(3) **Continue to teach?**: a prompt appears in front of the participant asking them whether they wish to continue teaching the robot, or stop the interaction. This appears only if participant has completed four or more teaching sessions.

For each of the conditions, we designed the *continue to teach* prompt to only appear from the fourth teaching session onwards. This was done to ensure that participants can see the robot make some amount of progress, and to dispel any wrong impression that the robot is meant to perform one-shot learning. The order of experimental conditions was counterbalanced using a Balanced Latin square. After each condition, participants complete questionnaires assessing their perception of the robot and their self-efficacy. After all four conditions, participants are interviewed and asked open-ended questions about their experience during the study as well as their thoughts on the teaching process, learning, and their own self-efficacy. All study sessions lasted approximately 60 minutes.

3.6 Data analysis

Quantitative data: Given the non-parametric nature of the collected data, we applied the aligned rank transform for non-parametric factorial analyses [49]. Therefore, we applied an Aligned Rank Transform (ART) ANOVA for all statistical analyses presented below. For pairwise comparisons, we used a Bonferroni correction.

Qualitative data: Two authors performed inductive thematic analysis [42] on interview transcripts and open questionnaire items. Themes were derived by each researcher separately, then consolidated together in a second iteration.

4 Results

Our results indicate that participants spent more time observing robots with lower than higher preexisting skill. Moreover, robots achieved higher proficiency with high preexisting skill than low. Participants perceived robots with higher preexisting skill to be more anthropomorphic. Faster learners were also perceived as more anthropomorphic than their slower counterparts.

Additionally, robots exhibiting faster learning were perceived to be more animated, likeable, and intelligent. Participants reported higher self-efficacy, better experience, and more willingness to teach when teaching fast learners.

4.1 Teaching time

We found that participants spent comparable amount of teaching robots with low ($Md = 114.5sec, IQR = 61$) existing skill than higher ($Md = 113.5sec, IQR = 55.5$). This finding was supported by the non-statistically significant main effect for the existing skill ($F(1, 23) = 0.051, p = 0.82, \eta^2 = 0.0022$). As for the learning rate, participants spent more time teaching slow learners ($Md = 139sec, IQR = 54.5$) than fast learners ($Md = 94sec, IQR = 55$). This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 18.4, p < 0.001, \eta^2 = 0.44$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 0.66, p = 0.42, \eta^2 = 0.03$).

4.2 Number of attempts

We found that participants used a comparable number of attempts between teaching a robot with low ($Md = 5.5, IQR = 4$) and high ($Md = 6, IQR = 3$) existing skill. This finding was supported by the non-statistically significant main effect for the existing skill ($F(1, 23) = 0.33, p = 0.56, \eta^2 = 0.014$). As for the learning rate, participants made more attempts to teach slow learners ($Md = 7, IQR = 3.25$) than fast learners ($Md = 5, IQR = 1.25$). This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 28.9, p < 0.001, \eta^2 = 0.56$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 1.72, p = 0.2, \eta^2 = 0.069$).

4.3 Achieved proficiency

We found that when participants decided to stop teaching, robots had achieved lower simulated proficiency with low ($M = 86\%, SD = 21$) than high($M = 88\%, SD = 17$) existing skill. This finding was supported by the statistically significant main effect for the existing skill ($F(1, 23) = 19.1, p < 0.001, \eta^2 = 0.45$). As for the learning rate, slow learners achieved lower proficiency ($M = 74\%, SD = 19$) than fast ($M = 100\%, SD = 0$). This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 56.4, p < 0.001, \eta^2 = 0.71$). Finally, we observed a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 19.19, p < 0.001, \eta^2 = 0.45$). However, none of the pairwise comparisons were statistically significant ($p > 0.05$) due to the p-value correction.

4.4 Initial observing time

We found that participants spent more time initially observing robots with low ($Md = 25sec, IQR = 3.25$) existing skill than high

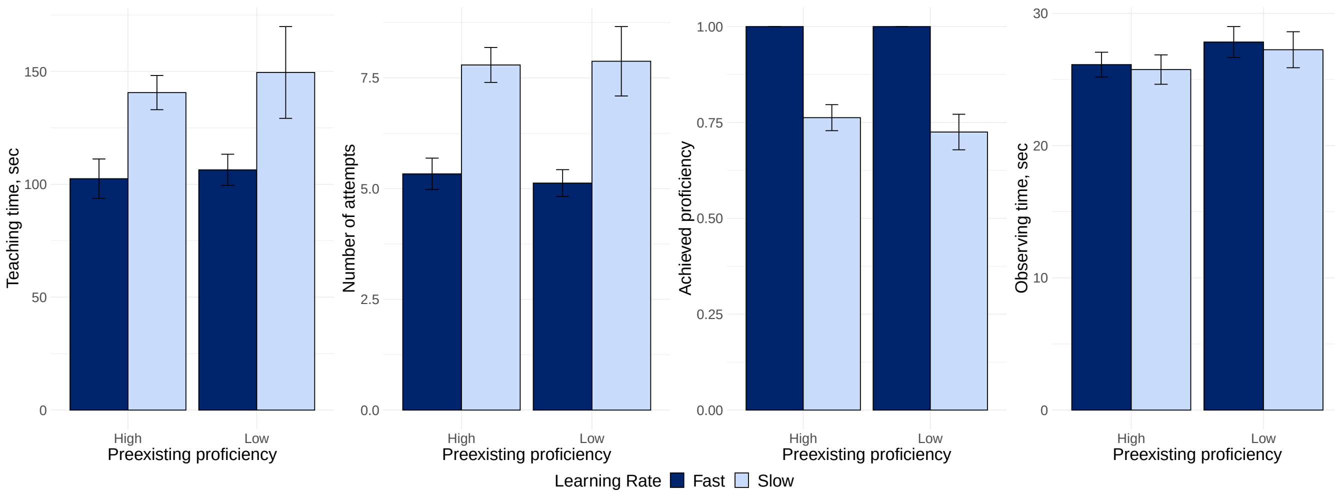

Figure 5: Overview of results: means and standard errors for the teaching time, number of attempts, achieved proficiency, and initial observing time. The results are grouped by two independent variables: (1) preexisting proficiency and (2) learning rate.

($Md = 24sec, IQR = 2$). This finding was supported by the statistically significant main effect for the existing skill ($F(1, 23) = 9.53, p = 0.005, \eta^2 = 0.3$). As for the learning rate, the time was comparable for slow ($Md = 25, IQR = 2$) and fast ($Md = 25, IQR = 3$) learners. This finding was supported by the non-statistically significant main effect for the learning rate ($F(1, 23) = 0.96, p = 0.33, \eta^2 = 0.04$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 0.005, p = 0.94, \eta^2 = 0.0002$).

4.5 Robot perception

Our results indicate that participants perceived fast-learning robots as more animated, likeable, and intelligent. Moreover, robots with low existing skills were perceived as less anthropomorphic than skilled robots. There was no influence on perceived safety from the existing skill or the learning rate.

Across the four conditions, we observed high internal reliability, measured using Cronbach's α, for the Godspeed subscales Anthropomorphism (0.874 to 0.939), Animacy (0.838 to 0.885), Likeability (0.914 to 0.970), Perceived Intelligence (0.901 to 0.957). Perceived Safety exhibited lower reliability (0.586 to 0.858). Overall internal reliability is summarized in table 1.

Table 1: Cronbach's α and 95% Confidence Intervals for Godspeed Questionnaire [4] Subscales

Subscale	Cronbach's α [95% CI]
Anthropomorphism	0.907 [0.874, 0.933]
Animacy	0.897 [0.860, 0.925]
Likeability	0.935 [0.912, 0.953]
Perceived Intelligence	0.936 [0.914, 0.953]
Perceived Safety	0.702 [0.577, 0.795]

4.5.1 Anthropomorphism. Participants perceived that robots with low ($Md = 2, IQR = 1.45$) existing skills were less anthropomorphic than those with high ($Md = 2.2, IQR = 1.45$). This finding was supported by the statistically significant main effect for the existing skill ($F(1, 23) = 4.8, p = 0.04, \eta^2 = 0.17$). Similarly, fast-learning robots ($Md = 2.3, IQR = 1.4$) were perceived as more anthropomorphic than slow-learning ones ($Md = 1.8, IQR = 1.05$). This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 13.2, p = 0.001, \eta^2 = 0.36$). Finally, we observed a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 7.62, p = 0.01, \eta^2 = 0.24$). However, none of the pairwise comparisons were statistically significant ($p > 0.05$) due to the p-value correction.

4.5.2 Animacy. Participants did not report differences in animacy for robots with low ($Md = 2.2, IQR = 1.65$) existing skills than those with high ($Md = 2.3, IQR = 1.05$). This finding was supported by the statistically non-significant main effect for the preexisting skill ($F(1, 23) = 2.7, p = 0.11, \eta^2 = 0.1$). However, fast-learning robots ($Md = 2.4, IQR = 1$) were perceived as more animated than slow-learning ones ($Md = 1.8, IQR = 1.25$). This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 26.7, p < 0.001, \eta^2 = 0.53$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 5.7, p = 0.03, \eta^2 = 0.19$).

4.5.3 Likeability. Participants did not have preferences in likeability between robots with low ($Md = 3.1, IQR = 1.05$) preexisting skills and those with high ($Md = 3.1, IQR = 1.05$) preexisting skills. This finding was supported by the statistically non-significant main effect for the existing skill ($F(1, 23) = 0.41, p = 0.52, \eta^2 = 0.017$). However, fast-learning robots ($Md = 3.4, IQR = 0.8$) were more likeable than slow-learning ones ($Md = 2.8, IQR = 1$). This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 17.4, p < 0.001, \eta^2 = 0.43$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 3.5, p = 0.07, \eta^2 = 0.13$).

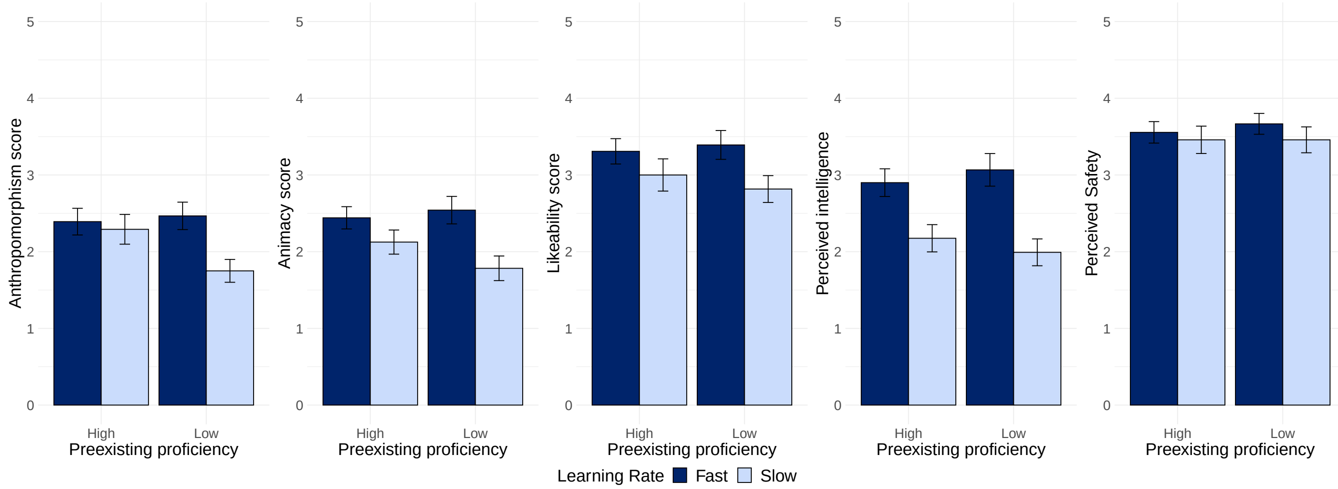

Figure 6: Overview of results: means and standard errors for subscales of robot perception. The results are grouped by two independent variables: (1) preexisting proficiency and (2) learning rate.

4.5.4 Perceived Intelligence. Similarly, participants did not perceive robots with low ($Md = 2.4, IQR = 1.65$) existing skills as the ones with lower intelligence than with high ($Md = 2.4, IQR = 1.45$). This finding was supported by the non-statistically significant main effect for the existing skill ($F(1, 23) = 0.21, p = 0.64, \eta^2 = 0.009$). However, fast-learning robots ($Md = 3.2, IQR = 1.2$) were perceived as more intelligent than slow-learning ones ($Md = 2, IQR = 1$). This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 38.2, p < 0.001, \eta^2 = 0.62$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 1.69, p = 0.2, \eta^2 = 0.07$).

4.5.5 Perceived Safety. Participants did not observe any differences in safety neither between robots with low ($Md = 3.6, IQR = 1.1$) and high ($Md = 3.3, IQR = 1.1$) existing skills nor slow ($Md = 3.3, IQR = 1.3$) and fast ($Md = 3.5, IQR = 1$) learners. Both of these results are supported by the non-statistically significant main effects for the existing skill ($F(1, 23) = 0.4, p = 0.53, \eta^2 = 0.017$) and learning rate ($F(1, 23) = 1.16, p = 0.29, \eta^2 = 0.05$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 0.37, p = 0.54, \eta^2 = 0.016$).

4.6 Teaching self-efficacy

Participants did not observe any differences in their own assessment of teaching efficacy between robots with low ($Md = 4, IQR = 3$) and high ($Md = 4.375, IQR = 3$) existing skills. This finding was supported by the non-statistically significant main effect for the existing skill ($F(1, 23) = 0.44, p = 0.51, \eta^2 = 0.018$). However, participants perceived themselves as more efficient teachers with fast- ($Md = 4.875, IQR = 1.375$) than with slow-learning ($Md = 2.375, IQR = 2.125$) robots. This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 47.7, p < 0.001, \eta^2 = 0.67$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 0.91, p = 0.34, \eta^2 = 0.038$).

4.7 Teaching experience

4.7.1 Pragmatic experience. Participants did not observe any differences in the pragmatic aspects of their teaching robots with low ($Md = 4, IQR = 2.75$) and high ($Md = 4, IQR = 2.5$) existing skills. This finding was supported by the non-statistically significant main effect for the existing skill ($F(1, 23) = 0.17, p = 0.68, \eta^2 = 0.007$). However, participants had a better pragmatic experience with fast- ($Md = 4.75, IQR = 1.25$) than with slow-learning ($Md = 2.875, IQR = 1.5$) robots. This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 52.6, p < 0.001, \eta^2 = 0.67$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 0.91, p = 0.34, \eta^2 = 0.038$).

4.7.2 Hedonic experience. Similarly, participants did not observe any differences in the hedonic aspects of their teaching robots with low ($Md = 3.75, IQR = 2.06$) and high ($Md = 3.375, IQR = 2$) existing skills. This finding was supported by the non-statistically significant main effect for the existing skill ($F(1, 23) = 1.2, p = 0.27, \eta^2 = 0.05$). However, participants had a better hedonic experience with fast- ($Md = 4.125, IQR = 1.625$) than with slow-learning ($Md = 2.75, IQR = 1.56$) robots. This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 25.5, p < 0.001, \eta^2 = 0.52$). Finally, we observed a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 5.2p = 0.034, \eta^2 = 0.18$). However, none of the pairwise comparisons were statistically significant ($p > 0.05$) due to the p-value correction.

4.8 Willingness to continue teaching

Lastly, robots' existing skill did not make a difference on participants' self-reported willingness to continue teaching them, if for low ($Md = 4, IQR = 3$) or high ($Md = 4, IQR = 3$) existing skills. This finding was supported by the non-statistically significant main effect for the existing skill ($F(1, 23) = 1.14, p = 0.29, \eta^2 = 0.05$).

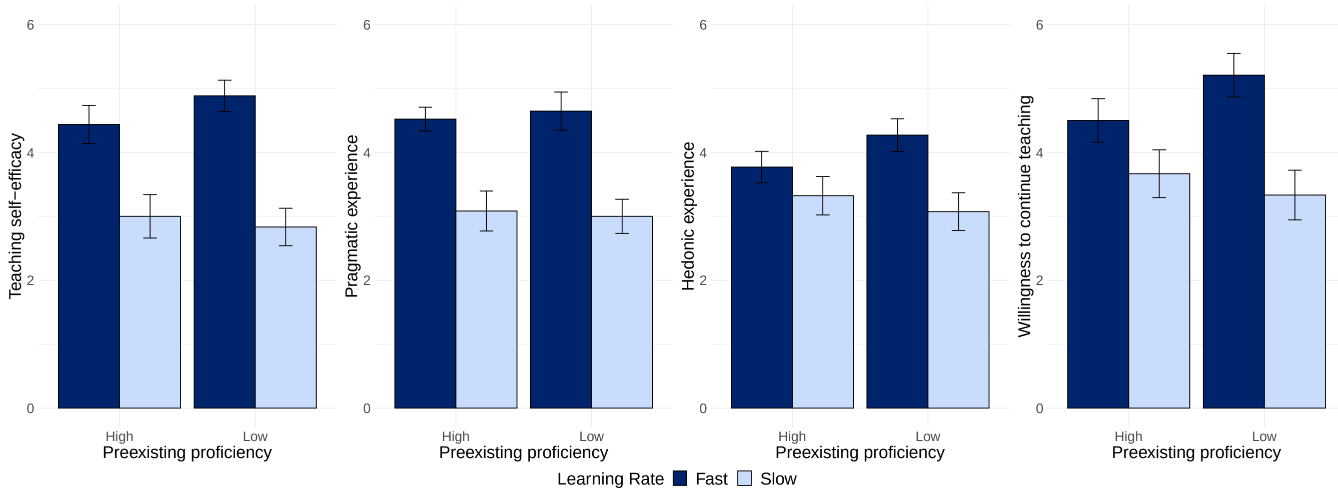

Figure 7: Overview of results: means and standard errors for teaching experience and willingness to teach. The results are grouped by two independent variables: (1) preexisting proficiency and (2) learning rate.

However, participants' willingness to teach was higher for fast- ($Md = 5, IQR = 2$) than slow-learning ($Md = 3, IQR = 3$) robots. This finding was supported by the statistically significant main effect for the learning rate ($F(1, 23) = 24.9, p < 0.001, \eta^2 = 0.52$). Finally, we did not observe a statistically significant interaction effect for existing skill * learning rate ($F(1, 23) = 3.88, p = 0.06, \eta^2 = 0.14$).

4.9 Qualitative results

We observed several themes in participants responses, related to perception of the robots, their performance, and the teaching process. Below are the major themes observed and participant quotes.

4.9.1 Influence of preexisting proficiency. Robots with higher preexisting proficiency caused participants to expect faster learning speed, and vice versa. "*I expected [the robot] to be faster because I saw that the rectangle was perfect and I thought, okay, then you also do the task easier [...]*"***[P1]*** "*[...]if the first task is already completely wrong, then I might lose motivation to continue.* "***[P3]*** "*I thought if [the robots] are good at drawing, they should be smarter[...]*"***[P10]*** "*When I saw the robot drawing a rectangle that was not a rectangle at all, I already had the preconception that, oh, he might be a little bit dumber or something.*"***[P17]***

Participants reported being (positively) surprised when a robot that exhibited low preexisting proficiency was learning fast, finding this behavior a sign of complexity. "*[...] the last [LF] one was did completely opposite of what it's supposed to do [on the drawing task] but it was more complex and it was pretty smooth.*"***[P12]*** "*[...] that's why I was surprised that it did learn the order even if it drew the hourglass shape at first.*"***[P19]*** "*[...] for the last [LF] one, for instance, the performance of the previous task was poor again, and then the robot learned really fast from what I've shown to the robot, and then I was happy.*"***[P24]***

Additionally, some participants noted that during teaching they perceived the preexisting skill level to be at an unrelated, therefore irrelevant, task. " *[...] the first [LF] one did it wrong, but learned pretty good. And the second [HS] one did the drawing right, but didn't learn as well. So I was like it doesn't matter how they behave in drawing, because it's a different task.*"***[P2]*** "*Well, first I thought [...] that it said something about how well they would perform in the next task, but I think it didn't. I think there was not connection [...] (between) the rectangle task and the cube task [...]*"***[P14]***

4.9.2 Influence of learning speed. Multiple participants reported higher motivation when the robot exhibited quick progress, and vice versa. "*[...]but if I see some progress and I see it's okay [...] just one small color mistake, and I would teach again.*"***[P3]*** Slow learning was described as frustrating and led some participants to cut the training short before the robot achieved the highest proficiency. "*The robots that learned quickly were very motivating and the other ones were very frustrating.*"***[P11]*** "*[...] some of them did really bad and I thought, okay, let's give them one other try, but it was rather frustrating.*"***[P14]*** "*I felt like its hard to teach this robot, so I quit.*"***[P21]*** Conversely, even small progress or change in output seems to have increased participants' motivation to teach the robot. "*[...] if they made a little progress, I thought, okay, you can do something, I can do something if I continue.*"***[P14]*** "*It was important for me to see differences in each iteration, so that I could see that the robot learned something or he improved from iteration to iteration. That's also how I decided when to stop continuing to teach.*"***[P17]***

4.9.3 Other remarks. Participants reported adapting their teaching approach for different robots depending on learning progress, as well as to explore the robot's responses to different teaching approaches: "*I tried different techniques with different robots. So with one robot, I tried to just do [the teaching] quickly, [...] just like a human being would be because it's like an easy task. [...] I also bring the [blocks] in a straight line because I wanted to know if [the robot] just throw it on there, or if they just really copy the positions as well.*"***[P1]***

Some participants seemed to have perceived one of the robots differently in aspects that were common to all robots. For example, one participant perceived the fast-learning, high preexisting proficiency (HF) robot to be more adaptive because it waved at them.

Notably, this waving animation was common to all four conditions. "*[...] I feel like the last [FS] one was a more adaptive one because I'm not sure if the others did it as well, but [the robot] waved when it started."**[P2]*** Similarly, one participant reported perceiving one robot's movement to be more fluid, even when the animations were all identical. "*I felt like there was one robot to move significantly less rigidly than the others."**[P19]*** Another participant reported feeling "sympathy" for the slow learners. "*I sympathize with the slow learners, I don't know."**[P18]***

Interestingly, one participant perceived the low preexisting proficiency, fast-learner (LF) robot to be more advanced than the rest. "*[...] I thought that it [LF] may learn; it has the capacity to learn more [...]"**[P12]*** In a similar vein, some participants expressed the feeling that the slow-learning robot "was not trying", as if ignoring their teaching. "*I was more patient with the first [LF] one. I gave more attempts.[...] like you realized he's really trying. Because the second [HS] one was not trying, or the third one [LS]."**[P2]***

5 Discussion and Future Work

It can be generally seen from the results that participants heavily favored fast learners, rating them as more likeable, intelligent, and animated. The hedonic experience ratings support this interpretation. It also makes sense that participants would be more willing to teach fast learners; faster learners provide earlier gratification and appear more responsive, which leads teachers to feel more effective and confident in their teaching. This is reflected in the higher teacher self-efficacy ratings when teaching fast learners. As we mentioned in section 3, participants favoring a fast-learning robot is to be expected. However, seeing as end users are not expected to have any control over how fast their commercial robots can learn, our results can help make sense of how the learning rate can affect end users. Additionally, after conducting our study, we now have an understanding of how the learning rate interacts with other properties of the robot to shape the perception of end users.

5.1 Do higher preexisting proficiency or faster learning rates make robots human-like?

Participant ratings associated a higher perception of anthropomorphism with robots having higher preexisting proficiency or faster learning rates. We interpret this to mean that participants found those qualities more human-like. This seems to contradict with previous work showing that participants may perceive a robot that makes mistakes as more *human* than one with perfect performance [33]. Further, participant comments reflected a tendency to anthropomorphize the robots, regardless of their capabilities. This is reflected in comments describing the robots and whether or not they were "trying" to learn, endowing the robots with willpower they do not have. This is consistent with existing work on tendencies of humans to anthropomorphize robots, machines, and other kinds of non-sentient agents [4, 12, 17, 28]. Another sign of anthropomorphizing can be seen in some participant comments assigning a gendered "he" to the genderless robot. However, since English is not the native language of some participants, this gendering might be an artifact of some participants' gendered native language. We acknowledge that the robot's humanoid form factor in our study may have contributed to particiants' tendencies to anthropomorphize the robot. The utility of robots looking and behaving like humans is situational [18], which presents *form factor* as another variable to be manipulated so that results can be studied. We would be interested in seeing if the same study conducted with non-humanoid robots in the student position would lead to different anthropomorphizing tendencies. Future work can also examine if the robot's form factor affects the other constructs we measured in our study.

5.2 Expectations and self-fulfilling prophecies

Participant comments repeatedly indicated that preexisting proficiency did not affect their motivation to continue teaching the robots. However, we could see in the session logs that proficient robots were trained to a higher skill level, which might indicate a higher desire to invest time and effort into teaching the robots. We interpret this to point to an unconscious bias formed by expectations; participants expected proficient robots to achieve higher skill and, therefore, continued teaching them until they reached that higher skill level. Conversely, robots that showed modest performance from the outset caused participants to not expect as much from them, leading to lower achieved proficiency. This describes a "self-fulfilling prophecy" that shaped the outcome of the teaching process. This is consistent with some findings indicating this possibility in pedagogy research on human teacher-human student interaction [24, 37].

While participants in our study found slow learners to be frustrating and perceived them as less intelligent, we do not consider algorithms that learn slowly to be inferior. In a more complex context where the task is more intricate, there can be benefits to learning slowly, e.g., avoiding hasty generalizations, which would be analogous to overfitting in the context of machine learning. The negative effects of slower learning can be mitigated by other means, such as the robot asking questions about the task [5, 20]. This can serve the dual benefit of giving the robot more information about the task beyond demonstration, while giving the human teacher the necessary feedback about the robot's learning progress, which may not be perceptible otherwise. It is also noteworthy that the frustration participants felt while teaching slow-learners is partially due to expectations. In the context of a commercial robot that a user can buy and teach a new task, if the robot's learning parameters and process are opaque to the user, it might also be perceived as slow or "stupid". It is then the role of interaction designers to set the right expectations for the user for such a system to be usable. We have also seen evidence that participants would adapt their teaching style and strategy when the robots were not responsive, which is consistent with findings from previous research [44]. Additionally, having lower expectations is not necessarily an aspect to be avoided, as lower expectations can reduce the negative impact of errors on trust [29]. This is another reason for future studies to investigate the effects *form factor* can have as an independent variable on such expectations.

5.3 What have we "learned"?

Our results show that the surest way to design a student robot that human teachers would prefer is to make it a fast learner. In our study, the robot's learning behavior was pre-programmed and rigid for the sake of consistency in simulation. Outside of the lab, when learning-capable assistant robots reach the commercial market, their performance and learning speed will be determined by the underlying learning algorithm, and mostly out of the hands of their owners or teachers. While preference for a faster learning robot may be obvious in hindsight, we argue that other factors affecting human perception of robot learning need to be studied, as well as the interaction between these factors and learning rate. Participant comments show that initial capabilities shape the impression made by the robot, as seen in 4.9.1 when a participant found a seemingly less capable robot (lower initial proficiency) to be more noteworthy and fulfilling to teach than its more capable counterpart. This means that designers of such robots need to consider communication methods and strategies that properly set the expectations of human teachers to avoid negative impressions or frustration. Human-robot teaching patterns following fully human classroom dynamics, coupled with the anthropomorphizing of robots, point to pedagogy research and literature [25, 43] as potential sources of insights for how to get the most out of teaching robots. These insights could potentially inform how to design interactions when humans are in the teaching role, such that the experience is more pleasant and less frustrating. In the context of a user customizing their own companion robot by teaching it a new task, it is to be expected that individual preferences play a role in the complexity of the taught skills, as well as the expectations of the robot.

6 Limitations

The user study we conducted in VR involved a simulated robot with pre-programmed behaviors that give participants the illusion of learning. This introduces some limitations to our approach. For example, VR interaction may not completely mimic interaction with a physical robot in a real situation. Future studies could implement a similar learning situation with a physical humanoid robot to verify the similarity of the findings. In a physical real-world setting, the simulated learning of our robots may not be sufficiently representative of more complex future learning algorithms and AI capabilities. In terms of the participant sample, a more or less "tech-savvy" sample could also produce different results. Future work could benefit from a larger and more diverse participant sample to further generalize our findings. We also acknowledge that the repetitiveness of the task, placing the cubes in the same orientation multiple times, could have influenced the results by introducing some aspects of boredom or fatigue. For consistency, our study involved teaching one task using one teaching method (human demonstrations) and one teaching process, which is fully demonstrating the process to the robot, followed by a demonstration from the robot of what was learned. We acknowledge that this may have restricted participants, thereby affecting their sense of self-efficacy. Future studies can expand on our setting by also accommodating more teaching methods. In our study, the robot was pre-programmed to succeed as long as the participant continues teaching it. Future work can also investigate more and different learning patterns, e.g., a robot with a more complex and dynamic learning process, or a robot that can make more types of mistakes.

7 Conclusion

In this paper, we investigated how humans teaching a student robot a new skill are influenced by the robot's preexisting skills and the learning rate of the new skill. From a controlled lab experiment ($N = 24$) in which robots had to learn object placement from human demonstration, we discovered that learning rate plays an important role in the perception of robots. Fast learners were perceived to be more intelligent, anthropomorphic, and likeable. Additionally, participants spent more time and used more attempts to teach slow learners than fast learning robots, even though participants reported feelings of frustration while teaching slow robots. Moreover, robots achieved lower proficiency with lower preexisting skills than higher, and slow learners achieved lower proficiency than fast learners. Our findings can help understand how future robots capable of learning from human demonstrations will be perceived and expected to behave by users.

References

[1] Baris Akgun, Kalesha Bullard, Vivian Chu, and Andrea Thomaz. 2014. An HRI approach to Learning from Demonstration. In *2014 AAAI Fall Symposium Series*.

[2] Pourya Aliasghari, Moojan Ghafurian, Chrystopher L. Nehaniv, and Kerstin Dautenhahn. 2023. How Do We Perceive Our Trainee Robots? Exploring the Impact of Robot Errors and Appearance When Performing Domestic Physical Tasks on Teachers' Trust and Evaluations. *ACM Transactions on Human-Robot Interaction* 12, 3 (May 2023), 1–41. doi:10.1145/3582516

[3] Albert Bandura. 1977. Self-efficacy: toward a unifying theory of behavioral change. *Psychological review* 84, 2 (1977), 191.

[4] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2008. Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. *International Journal of Social Robotics* 1, 1 (Nov. 2008), 71–81. doi:10.1007/s12369-008-0001-3

[5] Maya Cakmak and Andrea L. Thomaz. 2012. Designing robot learners that ask good questions. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction (HRI'12)*. ACM. doi:10.1145/2157689.2157693

[6] Maya Cakmak and Andrea L. Thomaz. 2014. Eliciting good teaching from humans for machine learners. *Artificial Intelligence* 217 (Dec. 2014), 198–215. doi:10.1016/j.artint.2014.08.005

[7] Sylvain Calinon and Aude G Billard. 2007. What is the teacher's role in robot programming by demonstration?: Toward benchmarks for improved learning. *Interaction Studies* 8, 3 (2007), 441–464. doi:10.1075/is.8.3.08cal

[8] Shruti Chandra, Raul Paradeda, Hang Yin, Pierre Dillenbourg, Rui Prada, and Ana Paiva. 2018. Do Children Perceive Whether a Robotic Peer is Learning or Not?. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction* (Chicago, IL, USA) *(HRI '18)*. Association for Computing Machinery, New York, NY, USA, 41–49. doi:10.1145/3171221.3171274

[9] Sonia Chernova and Andrea L Thomaz. 2014. *Robot learning from human teachers*. Morgan & Claypool Publishers.

[10] Vivienne Bihe Chi and Bertram F. Malle. 2023. Calibrated Human-Robot Teaching: What People Do When Teaching Norms to Robots*. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 1308–1314. doi:10.1109/RO-MAN57019.2023.10309635

[11] Joachim De Greeff and Tony Belpaeme. 2015. Why robots should be social: Enhancing machine learning through social human-robot interaction. *PLoS one* 10, 9 (2015), e0138061.

[12] Roy de Kleijn, Lisa van Es, George Kachergis, and Bernhard Hommel. 2019. Anthropomorphization of artificial agents leads to fair and strategic, but not altruistic behavior. *International Journal of Human-Computer Studies* 122 (Feb. 2019), 168–173. doi:10.1016/j.ijhcs.2018.09.008

[13] Joseph DelPreto, Jeffrey I. Lipton, Lindsay Sanneman, Aidan J. Fay, Christopher Fourie, Changhyun Choi, and Daniela Rus. 2020. 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. doi:10.1109/icra40945.2020.9196754

[14] Mohamed El-Shamouty, Xinyang Wu, Shanqi Yang, Marcel Albus, and Marco F Huber. 2020. Towards safe human-robot collaboration using deep reinforcement learning. In *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 4899–4905.

[15] Bernd Finkemeyer. 2017. Towards safe human-robot collaboration. In *2017 22nd International Conference on Methods and Models in Automation and Robotics (MMAR)*. IEEE, 883–888.

[16] Leonard Friedman. 1977. Robot perception learning. *ACM SIGART Bulletin* 63 (1977), 44–49.

[17] Susan R. Fussell, Sara Kiesler, Leslie D. Setlock, and Victoria Yew. 2008. How people anthropomorphize robots. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction (HRI '08)*. ACM. doi:10.1145/1349822.1349842

[18] Jean-Christophe Giger, Nuno Piçarra, Patrícia Alves-Oliveira, Raquel Oliveira, and Patrícia Arriaga. 2019. Humanization of robots: Is it really such a good idea? *Human Behavior and Emerging Technologies* 1, 2 (April 2019), 111–123. doi:10.1002/hbe2.147

[19] Samuel D Gosling, Peter J Rentfrow, and William B Swann. 2003. A very brief measure of the Big-Five personality domains. *Journal of Research in Personality* 37, 6 (Dec. 2003), 504–528. doi:10.1016/s0092-6566(03)00046-1

[20] Soheil Habibian, Ananth Jonnavittula, and Dylan P. Losey. 2022. Here's What I've Learned: Asking Questions that Reveal Reward Learning. *ACM Transactions on Human-Robot Interaction* 11, 4 (Sept. 2022), 1–28. doi:10.1145/3526107

[21] Erin Hedlund, Michael Johnson, and Matthew Gombolay. 2021. 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 (HRI '21)*. ACM. doi:10.1145/3434073.3444664

[22] Erin Hedlund-Botti and Matthew C. Gombolay. 2023. Investigating Learning from Demonstration in Imperfect and Real World Scenarios. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23, Vol. 189)*. ACM, 769–771. doi:10.1145/3568294.3579980

[23] Seyed Ramezan Hosseini, Alireza Taheri, Minoo Alemi, and Ali Meghdari. 2021. One-shot Learning from Demonstration Approach Toward a Reciprocal Sign Language-based HRI. *International Journal of Social Robotics* 16, 4 (Aug. 2021), 645–657. doi:10.1007/s12369-021-00818-1

[24] Lee Jussim and Kent D. Harber. 2005. Teacher Expectations and Self-Fulfilling Prophecies: Knowns and Unknowns, Resolved and Unresolved Controversies. *Personality and Social Psychology Review* 9, 2 (May 2005), 131–155. doi:10.1207/s15327957pspr0902_3

[25] Tasneem Kaochar, Raquel Torres Peralta, Clayton T. Morrison, Ian R. Fasel, Thomas J. Walsh, and Paul R. Cohen. 2011. *Towards Understanding How Humans Teach Robots*. Springer Berlin Heidelberg, 347–352. doi:10.1007/978-3-642-22362-4_31

[26] W. Bradley Knox and Peter Stone. 2009. Interactively shaping agents via human reinforcement: the TAMER framework. In *Proceedings of the Fifth International Conference on Knowledge Capture* (Redondo Beach, California, USA) *(K-CAP '09)*. Association for Computing Machinery, New York, NY, USA, 9–16. doi:10.1145/1597735.1597738

[27] N. Kubota, Y. Nojima, F. Kojima, and T. Fukuda. 2001. Dual learning for perception and behavior of mobile robots. In *Proceedings Joint 9th IFSA World Congress and 20th NAFIPS International Conference (Cat. No. 01TH8569)*, Vol. 3. 1401–1406 vol.3. doi:10.1109/NAFIPS.2001.943754

[28] Dieta Kuchenbrandt, Friederike Eyssel, Simon Bobinger, and Maria Neufeld. 2013. When a Robot's Group Membership Matters: Anthropomorphization of Robots as a Function of Social Categorization. *International Journal of Social Robotics* 5, 3 (May 2013), 409–417. doi:10.1007/s12369-013-0197-8

[29] Min Kyung Lee, Sara Kiesler, Jodi Forlizzi, Siddhartha Srinivasa, and Paul Rybski. 2010. Gracefully mitigating breakdowns in robotic services. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE. doi:10.1109/hri.2010.5453195

[30] Gregory Lemasurier, Gal Bejerano, Victoria Albanese, Jenna Parrillo, Holly A. Yanco, Nicholas Amerson, Rebecca Hetrick, and Elizabeth Phillips. 2021. Methods for Expressing Robot Intent for Human–Robot Collaboration in Shared Workspaces. *ACM Transactions on Human-Robot Interaction* 10, 4 (Sept. 2021), 1–27. doi:10.1145/3472223

[31] Fabio Dalla Libera, Takashi Minato, Ian Fasel, Hiroshi Ishiguro, Emanuele Menegatti, and Enrico Pagello. 2007. Teaching by touching: An intuitive method for development of humanoid robot motions. In *2007 7th IEEE-RAS International Conference on Humanoid Robots*. 352–359. doi:10.1109/ICHR.2007.4813893

[32] Maximilian Metzner, Daniel Utsch, Matthias Walter, Christian Hofstetter, Christina Ramer, Andreas Blank, and Jörg Franke. 2020. A system for human-in-the-loop simulation of industrial collaborative robot applications. In *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*. IEEE, 1520–1525.

[33] Nicole Mirnig, Gerald Stollnberger, Markus Miksch, Susanne Stadler, Manuel Giuliani, and Manfred Tscheligi. 2017. To Err Is Robot: How Humans Assess and Act toward an Erroneous Social Robot. *Frontiers in Robotics and AI* 4 (May 2017). doi:10.3389/frobt.2017.00021

[34] Nina Moorman, Erin Hedlund-Botti, Mariah Schrum, Manisha Natarajan, and Matthew C. Gombolay. 2023. Impacts of Robot Learning on User Attitude and Behavior. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (New York, NY, USA, 2023-03-13) *(HRI '23)*. Association for Computing Machinery, 534–543. doi:10.1145/3568162.3576996

[35] Astrid Rosenthal-Von Der Pütten and Nikolai Bock. 2018. Development and Validation of the Self-Efficacy in Human-Robot-Interaction Scale (SE-HRI). *J. Hum.-Robot Interact.* 7, 3, Article 21 (dec 2018), 30 pages. doi:10.1145/3139352

[36] Sandra Robla-Gómez, Victor M Becerra, José Ramón Llata, Esther Gonzalez-Sarabia, Carlos Torre-Ferrero, and Juan Perez-Oria. 2017. Working together: A review on safe human-robot collaboration in industrial environments. *Ieee Access* 5 (2017), 26754–26773.

[37] Robert Rosenthal and Lenore Jacobson. 1968. Pygmalion in the classroom. *The Urban Review* 3, 1 (Sept. 1968), 16–20. doi:10.1007/bf02322211

[38] Akanksha Saran, Elaine Schaertl Short, Andrea Thomaz, and Scott Niekum. 2019. Enhancing robot learning with human social cues. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 745–747.

[39] Joe Saunders, Chrystopher L. Nehaniv, and Kerstin Dautenhahn. 2006. Teaching robots by moulding behavior and scaffolding the environment. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction (HRI06, Vol. 7)*. ACM, 118–125. doi:10.1145/1121241.1121263

[40] Martin Schrepp, Andreas Hinderks, et al. 2017. Design and evaluation of a short version of the user experience questionnaire (UEQ-S). (2017).

[41] Mariah L. Schrum, Erin Hedlund-Botti, Nina Moorman, and Matthew C. Gombolay. 2022. MIND MELD: Personalized Meta-Learning for Robot-Centric Imitation Learning. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE. doi:10.1109/hri53351.2022.9889616

[42] David R Thomas. 2006. A general inductive approach for analyzing qualitative evaluation data. *American journal of evaluation* 27, 2 (2006), 237–246.

[43] Andrea L. Thomaz and Cynthia Breazeal. 2008. Experiments in socially guided exploration: lessons learned in building robots that learn with and without human teachers. *Connection Science* 20, 2–3 (Sept. 2008), 91–110. doi:10.1080/09540090802091917

[44] Andrea L. Thomaz and Cynthia Breazeal. 2008. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence* 172, 6 (2008), 716–737. doi:10.1016/j.artint.2007.09.009

[45] Andrea L. Thomaz and Maya Cakmak. 2009. Learning about objects with human teachers. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction* (La Jolla, California, USA) *(HRI '09)*. Association for Computing Machinery, New York, NY, USA, 15–22. doi:10.1145/1514095.1514101

[46] Anna-Maria Velentza, Nikolaos Fachantidis, and Ioannis Lefkos. 2022. Human-robot interaction methodology: Robot teaching activity. *MethodsX* 9 (2022), 101866. doi:10.1016/j.mex.2022.101866

[47] George-Christopher Vosniakos, Lucas Ouillon, and Elias Matsas. 2019. Exploration of two safety strategies in human-robot collaborative manufacturing using Virtual Reality. *Procedia manufacturing* 38 (2019), 524–531.

[48] Auriel Washburn, Akanimoh Adeleye, Thomas An, and Laurel D. Riek. 2020. Robot Errors in Proximate HRI: How Functionality Framing Affects Perceived Reliability and Trust. *ACM Transactions on Human-Robot Interaction* 9, 3 (May 2020), 1–21. doi:10.1145/3380783

[49] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. *The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only Anova Procedures*. Association for Computing Machinery, New York, NY, USA, 143–146.

[50] Setareh Zafari, Isabel Schwaninger, Matthias Hirschmanner, Christina Schmidbauer, Astrid Weiss, and Sabine T. Koeszegi. 2019. "You Are Doing so Great!" – The Effect of a Robot's Interaction Style on Self-Efficacy in HRI. In *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE. doi:10.1109/ro-man46459.2019.8956437