User Interface Interventions for Improving Robot Learning from Demonstration

Ornnalin Phaijit
l.phaijit@unsw.edu.au
The University of New South Wales
Sydney, Australia

Claude Sammut
c.sammut@unsw.edu.au
The University of New South Wales
Sydney, Australia

Wafa Johal
wafa.johal@unimelb.edu.au
The University of Melbourne
Melbourne, Australia

ABSTRACT

Teaching robots can be challenging, particularly for novice human users who struggle to understand the robot's learning process. Current research in interactive robot learning lacks effective methods for assessing a user's interpretation of the robot's learning state, which makes it difficult to compare different teaching approaches. To address these issues, we propose and demonstrate a method for assessing the user's interpretation of the robot's learning state in an interactive learning scenario with a robotic manipulator. Additionally, we draw on existing literature to categorise types of interface interventions that can enhance the human-robot teaching process for novice users – both pragmatically and hedonically. In a user study (N=30), we implement two of these interventions and show how they improve robot performance, teaching efficiency and interpretability. These findings provide preliminary insights into the design of effective human-robot teaching interfaces and can be used to assist the development of future teaching approaches.

CCS CONCEPTS

• **Human-centered computing → HCI design and evaluation methods**; *Mixed / augmented reality*; **Empirical studies in HCI**.

KEYWORDS

human-robot interaction, learning from demonstration, interpretability, explainability, mixed reality

ACM Reference Format:
Ornnalin Phaijit, Claude Sammut, and Wafa Johal. 2023. User Interface Interventions for Improving Robot Learning from Demonstration. In *Proceedings of the 11th International Conference on Human-Agent Interaction (HAI '23).* ACM, New York, NY, USA, 10 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Programming a robot can be a challenging and time-consuming task that requires specialized skills and knowledge. It may be beyond the capabilities of many people, and pre-programmed robots cannot accommodate the specific preferences of every user. To make human-robot interaction more intuitive and accessible to a wider audience, it is important to enable laypeople to teach robots tasks just as they would teach another person. A common approach for intuitive human-robot teaching is Learning from Demonstration (LfD), where the teacher demonstrates the task to the robot. However, one major challenge in this approach is the user's inability to understand the robot's learning process, especially when the teacher lacks expertise in programming or robotics. This knowledge gap presents a significant obstacle to effective human-robot teaching, and addressing it is critical to developing more intuitive and user-friendly human-robot interfaces.

With the need to facilitate the communication between human and robot, Explainable Artificial Intelligence (XAI) is increasingly being employed in Human-Robot Interaction (HRI), aka Explainable Robotics [48]. The metrics in XAI are however typically for assessing the system's explainability after its task execution [23]. In other words, the focus has been on explaining to the user why the agent has made a particular decision in *one* scenario. This presents limitations in human-to-robot teaching. How could the teacher understand the *overall* picture of the robot's mental model during teaching?

In a human-human teaching process, a way to gauge the learner's current learning state has been to evaluate the student's test performance. However, it would be unrealistic to be examining the robot at every step due to the greater time and resources required for the robot's physical movements. In between tests, the teacher should also have an idea of the student's learning state. In human-human teaching, this could be based on the student's facial expressions (i.e. confusion or confidence), questions, or remarks. As human teachers have also been learners, they can rely on the Theory of Mind (ToM) [16] to use their past learner experience to build a model of their human students. In the case of a robot learner, the teacher cannot use ToM as the robot does not learn like a human student. This often leads to misinterpretation of the robot's learning status and may therefore induce teaching mistakes [27, 51]. Hence this exacerbates the need for an approach that allows a novice user to better understand the robot.

In order to more effectively evaluate between different teaching approaches, the interpretability of the robot's learning state must be considered. However, there does not yet exist in literature a method for quantitatively assessing the human teacher's interpretation of the robot's learning state.

How can the human-robot interface, or the bi-directional communication between human and robot [60], be used to improve the human-robot teaching experience? How can we quantitatively assess robot interpretability between different approaches? In this paper, we address these issues. Our contribution are:

- A categorisation and definition of four user interface (UI) interventions to the LfD pipeline in HRI used for improving

HAI '23, December 4–7, 2023, Gothenburg, Sweden

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/XXXXXXX.XXXXXXX

the human experience. We provide examples for each of these categories and demonstrate in a user study to show how two of these interventions can be used to improve human-robot teaching.
- A method for quantitatively assessing the human's interpretation of the robot's learning state. We also demonstrate our proposed method in a user study.

Our proposed categories offer valuable guidance for both researchers and practitioners in the field of human-robot teaching, serving as a useful design framework. Additionally, our proposed quantitative interpretability assessment provides a crucial tool for evaluating human-robot teaching approaches, offering a clear path forward for improving the field's methods and outcomes.

2 BACKGROUND AND FORMALISATION OF LEARNING FROM DEMONSTRATION

In 2010, Billing and Hellström [3] formalised the concepts in Learning from Demonstration (LfD). The task space B represents the goal behaviour for the robot that the human is trying to teach. In LfD, the human teacher performs a set of demonstrations $b = \{\beta_1, ..., \beta_n\}$ for the robot learner, where β_i represents a single demonstration. The policy derivation, or learning, of the robot is the selection of π from the controller space Π by using the learning function λ:

$$\pi = \lambda(b) \in \Pi \tag{1}$$

The robot then uses the realization function Λ to map the controller π to a task realisation r from the set of all possible tasks the robot is capable of performing I_h.

$$r = \Lambda(\pi) \in I_h \tag{2}$$

In 2020, Sena and Howard [47] extended this formalization to show how the human teaching behaviour for a robot learner could be quantified (Fig. 1). Their framework introduced a model of the teacher's belief space. In particular, the human teacher's estimated robot's learning state $\tilde{R}$ is obtained from using an interpretation function ω on the observed robot's task execution r and the teacher's bias, Q.

$$\tilde{R} = \omega(r, Q) \tag{3}$$

In this context, Q (the teacher's bias) has been defined to capture factors which are difficult to measure, such as the human mental state or their prior expectation of the learning. In order to select the next set of demonstrations $b^{(n+1)}$ for the robot, the teacher uses their estimate of the robot's current learning state $\tilde{R}$, the teacher's estimation of the task space $\tilde{B}$, the set of demonstrations already given $b^{(n)}$ and the teacher's bias Q.

$$b^{(n+1)} = \Omega(\tilde{R}, \tilde{B}, b^{(n)}, Q) \tag{4}$$

As the purpose of teaching is to allow for the robot to be able to perform the intended task, the human's teaching efficacy can be defined as the portion of the task space that the robot is able to perform correctly:

$$\epsilon = \frac{|R \cap B|}{|B|}, \epsilon \in [0, 1] \tag{5}$$

The teaching efficiency η is the teaching efficacy ϵ normalized by the number of demonstrations required to achieve efficacy ϵ:

$$\eta = \frac{\epsilon}{|b|}, \eta \in [0, 1] \tag{6}$$

Indeed, with higher quality demonstrations, the robot would be able to achieve the same efficacy with a smaller number of demonstrations. Note that here the term "efficiency" is limited in that it only takes into account the number of demonstrations and not the complexity of the demonstrations, the time spent, or the resources required.

In the next section, we describe four types of user interface (UI) interventions to the LfD pipeline which may help to improve the teaching process.

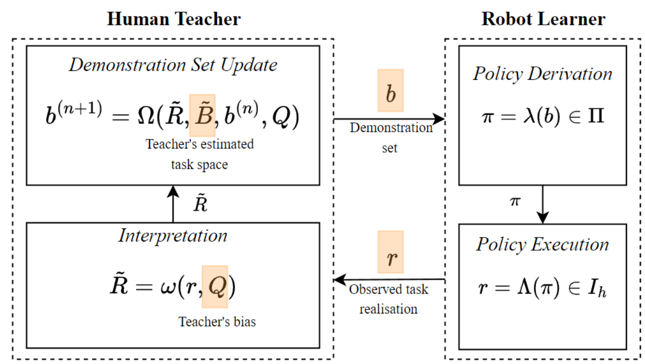

Figure 1: The human-robot demonstration teaching pipeline (adapted from [47]). The highlighted areas are our proposed placements of user interface interventions.

3 USER INTERFACE INTERVENTIONS FOR IMPROVING HUMAN-ROBOT TEACHING

The user experience of robot teaching can be improved in two different ways: via pragmatic or hedonic qualities [7, 17]. Improving robot teaching by focusing on pragmatic qualities would, for instance, be working towards increasing the teaching efficacy or efficiency. On the other hand, improving hedonic qualities could be comprised of reducing the teacher's mental load, increasing user engagement, or increasing user trust. For instance, by being transparent and visualising information concerning the robot or its environment to the user, this can increase user trust [46]. By presenting information to the user instead of relying on the user to memorise every detail, this can reduce the cognitive load for the human teacher – which may increase human trust in the robot [2]. Note that improving pragmatic qualities could also in turn improve hedonic qualities. For example, by improving the teaching effectiveness, this *may* increase perceived agent competency, thereby possibly inducing greater user trust and engagement [40].

In this paper, we focus on improving pragmatic qualities. The components in the human-robot teaching pipeline can be seen in Fig. 1. What user interface (UI) interventions can be made to the pipeline to improve the process? Firstly, the learner's internals, such as its policy derivation and execution, are dependent upon the algorithms of the robot and hence are removed from the candidates for UI interventions. Hence, now there remain three parts: 1) the teacher's internals, 2) the demonstration set given to the robot b, and 3) the observed task realisation r. UI interventions can be applied to these three locations. The teacher's internals however lack precision due to the many variables involved, and hence need to be broken down. The interpretation function ω and the demonstration set update function Ω vary from person to person. Hence we instead focus on the variables given to them: $\tilde{R}, \tilde{B}, b, Q$, and r. The variables b and r have already previously been identified as places of possible interventions. This leaves the teacher's bias Q, the teacher's interpretation of the learner state $\tilde{R}$ and the teacher's estimation of the task space $\tilde{B}$. As $\tilde{R}$ is fully reliant on r and Q – see Eq. (3) – which have both already been identified as possible places of UI interventions, $\tilde{R}$ can be left out. We therefore identify two additional areas that can be improved on: Q and $\tilde{B}$. Hence, we propose the following four categories of UI interventions that can be made to assist in the LfD pipeline:

(1) Intervention on b (demonstrations given to the robot)
(2) Intervention on r (the observed task realization)
(3) Intervention on Q (the teacher's bias)
(4) Intervention on $\tilde{B}$ (the teacher's beliefs regarding the task space)

Note that by improving one of the factors mentioned (e.g. r), this may consequently improve the quality of the demonstrations given to the robot b. However, here these intervention types signify the exact variable that is being directly improved on.

We produced visual examples of these different interventions in mixed reality (MR), using the development platform Unity in Figs. 2, 3, 4, and 6. We used MR as it is a technology that has the ability to overlay digital information over the physical world [34]. These interventions may be done via other means (e.g. the traditional computer screen), however, comparing MR with other methods is out of our work's scope.

3.1 Intervention on the demonstrations given to the robot b

The human demonstration from a novice can be sub-optimal due to several reasons including the difficulty of operating the robot, or the high cognitive load required to memorise given demonstrations. Here we present *some* possible interventions that may help to improve the human demonstrations b to the robot.

3.1.1 Incorrect demonstrations. The authors in [47] suggested that a common teaching failure includes the teacher providing incorrect demonstrations. This may arise due to a number of factors including the teacher not being able to operate the robot correctly, or the teacher making errors due to the high cognitive load required when teaching a robot. The UI intervention can be introduced to allow for the ability to go back in time and delete or modify past demonstrations to decrease the number of incorrect demonstrations for the robot (Fig. 2a). Timelines can be shown in the display to the user (e.g. [33]). Users rewinding to a particular point in time to give corrective demonstrations to the agent learner is already possible [59]. This could be extended to allow for users to modify their *own* past demonstrations – an "artificial timeline" [39]. By letting the

(a)

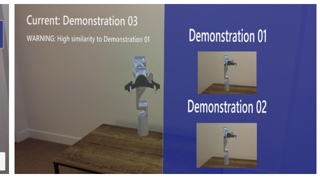

(b)

(c)

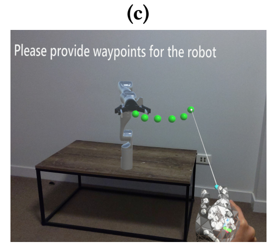

(d)

Figure 2: Improving demonstrations given to the robot. (a) The user could choose to edit or delete the demonstration upon an incorrect demonstration, (b) the user could review past demonstrations and the application warns the user of high similarity to past demonstrations, (c) the user can control the robot by interacting with the virtual robot in MR, (d) the user provides way points (green spheres) for the robot arm by indicating the point in the MR interface

user "time travel", the user can more closely examine their teaching and mistakes can be corrected.

3.1.2 Redundant demonstrations. Redundant demonstrations have also been put forward as another common reason for teaching failures [12, 47]. These are demonstrations which do not sufficiently differ from the prior demonstrations given, and hence lower the teaching efficiency. By being able to review past demonstrations, this can also help the user detect redundant demonstrations. A system for calculating and detecting demonstration similarity could be implemented to notify the user during a new demonstration (Fig. 2b). As proposed in previous literature [47], this can be accomplished by having some measure of similarity s between the current demonstration β_n and previous demonstrations $\beta_1, ..., \beta_{n-1}$ and checking that it does not exceed the ambiguity threshold δ_a:

$$s(\beta_n, \{\beta_1, ..., \beta_{n-1}\}) \leq \delta_a \tag{7}$$

Once the similarity is exceeded, the user can be informed. S/he now has the option to abort or continue with the demonstration. By minimising demonstration ambiguity, the teaching efficiency can be improved.

3.1.3 Action Mapping. A known difficulty in LfD is the retargeting problem [8] – where human actions from demonstrations must be mapped to robot actions. This problem can traditionally be bypassed by guiding the robot through kinesthetic interactions [6, 25, 26]. However, direct physical contact with the robot may not always be ideal or possible. Robots' motions can instead be guided through human teleoperation via a joystick such as in [10, 24, 50], or via interacting with the virtual model of the robot/other virtual objects (Fig. 2c) such as in [15, 18, 28, 41, 58] where the user may click and drag a virtual object to control the robot movements. In the case of motor impairment, user gaze may be mapped to robot actions in brain-computer interfaces [11, 56]. It can be referred to as providing "Augmented Control" [39] to the user.

UI interventions may therefore allow for novices to better translate their actions to those of the robots'.

3.1.4 Spatial mapping. For the novice user, translating their 3D space problem into coordinates for the robot could be difficult and unintuitive. UI interventions however can be employed to visualise to the user the points in 3D space (Fig. 2d). This is already being done in multiple non-learning HRI tasks such as allowing for the user to provide waypoints for the robot via MR [29, 42, 54] or presenting data more intuitively by spatially constructing virtual scenes [43, 44].

3.2 Intervention on the observed task realization r

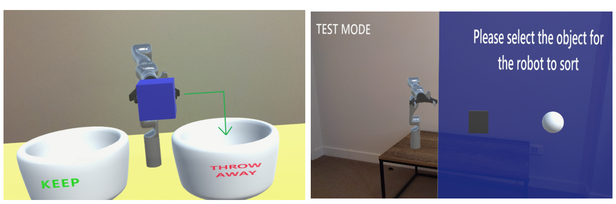

(a) (b)

Figure 3: Improving the observed task realization. (a) The application shows the simulated path of the robot arm without the need for the robot to move physically, (b) To better understand the robot's current knowledge, the user can enquire about specific tasks via the interface. In the figure, the user can select a test object for the robot to sort.

For the teacher to better understand the learner's state, the task executions observed by the teacher are crucial. The following are some possible ways UI interventions could help to improve the observed task realisation:

3.2.1 Increasing the number of task realizations. To better understand the learner's state, the teacher may wish to see the robot's performance after a certain number of demonstrations. However, by requiring the robot to perform the tasks physically, this may impose time and resource constraints. By allowing for the simulated trajectories of the robot to be shown such as in [13, 36, 38], or the visual example in Fig. 3a, this can avoid these constraints, particularly time. This allows for a greater number of task executions for the same time length, and can therefore improve the set of r for the same amount of time and resources.

3.2.2 Enquire about specific task executions. In human-human teaching, the teacher typically asks the student specific questions to have a better idea of the student's state. This is akin to the use of validation sets in traditional machine learning. Validation set results help to shape the machine's subsequent learning, while informing the human of its current learning state [20, 45]. In interactive robot learning however, the teacher may wish to test the robot's ability by assigning a task to it at arbitrary points in the teaching process, rather than only at the end of epochs. Moreover, the teacher may wish to change the test each time to better understand the robot. To enable this, an interface can be employed to facilitate the process of asking the robot to execute specific tasks (see our example in Fig. 3b). The UI, for instance, may assist with this by superimposing a virtual robot onto the physical robot to carry out the task without the need to reset the real robot, or objects such as in [22, 35].

3.3 Intervention on the teacher's bias Q

(a)

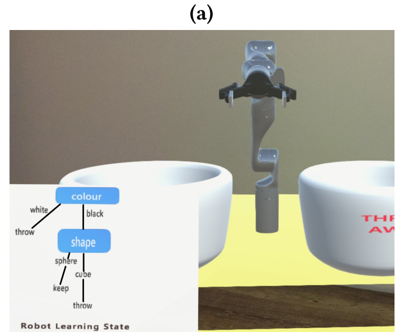

(b)

Figure 4: Intervention on the teacher's bias. (a) The range of the robot arm is indicated on the display as an orb around the robot, (b) The robot's decision tree is presented on the user display,

The variable Q captures the teacher's bias. However, in order to better formulate how to improve teaching, Q can be broken down to be composed of the human's:

- beliefs of the robot's operational capabilities,
- beliefs of the robot's mental model, and

- other biases

Other biases include the human's personal biases, mental state at the time and any other biases. Improvement to these is out of this work's scope and henceforth we will focus on the first two for this subsection.

3.3.1 The robot's operational capabilities. A novice user may not be aware of the range limit of the physical robot's arm, or the positional/rotational constraints of the robot's physical parts. By showing these capabilities visually to the user such as in [1, 32, 37] and Fig. 4a, this could lead to a more efficient teaching process due to less need for adjustments. Furthermore, instead of simply presenting this information to the user in the form of a handbook, visually displaying these thresholds and limits can be more comprehensible, especially to novices.

3.3.2 The robot's mental model. A major challenge in human-robot teaching is the mental barrier between the human teacher and the robot learner. Usually in human-human teaching, the human teacher is able to put her/himself in the student's perspective as s/he has been in that position before as a student. However, this is indeed not the case for human-robot teaching as the teacher has not been a *robot* learner before. The complication further increases when the teacher in question is a novice user. Therefore information regarding the robot's understanding needs to be conveyed in a way that is easy for a human to understand. The UI has been used to show graphical [31] (e.g. Fig 4b) and semantic representations [14, 55] of the robot's understanding and learning process for the task. In previous literature, it has also been used by the robot to point out particular objects in the user's view by, for instance, overlaying a virtual circle [57], an arrow or even a virtual arm for the robot for referring to objects [21]. The robot's intent has also been shown through arrows, waypoints or even the robot's virtual eye gaze [53]. Although some of these methods may not be specifically intended for human-robot teaching, it is evident that there are many different ways for the human to better understand the robot learner's mental model.

3.4 Intervention on the teacher's beliefs of the task space B

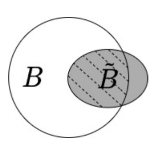

Figure 5: Venn diagram of the actual task space B and the user's understanding of the task space $\tilde{B}$.

Billing and Helström [3] formalised the concept of the task space, or intended behaviour B. Sena and Howard [47] then extended the formalisation to include $\tilde{B}$ to represent the teacher's estimation of the task space. Although there is a distinction between B and $\tilde{B}$ (see Fig. 5), there has not yet been much effort to examine the significance of the distinction. Here we dissect what the areas of misalignment can imply, and propose how a UI intervention may help.

(a) (b)

Figure 6: Improving the teacher's understanding of the task space. (a) The MR interface presents a scenario in which an object is out of the robot's range and asks if it needs to learn what to do in this situation (b) The user is notified via the display that this demonstration may be outside of the task space intended.

Where B is the actual task space and $\tilde{B}$ is the human teacher's beliefs of the task space, $B \setminus \tilde{B}$ represents the actual task space that the user is unaware of. The mistake of underestimating the task is commonly made by experts, and even more so by novices who are less familiar with programming/robotics. Active learning (AL) is when the learning agent interactively asks the teacher specific questions in order to learn better. This method has been introduced to HRI [4, 5, 9, 19]. Employing AL, the teaching interface could help to identify edge cases for the user by allowing for the robot to raise "Demonstration Queries" (term coined by Cakmak [5]). This is when the robot finds a situation that its mental model does not yet cover. For example, should the robot be taught how to handle an object of a particular colour and of a particular shape? A UI intervention could include graphics being overlaid in the display to visualise to the user what these edge cases may look like, such as showing a wall in front of the robot, or displaying an object out of the robot's reach (see Fig. 6a), if it has not been taught what to do in these cases. The user may then be prompted to teach for these cases, or the user may indicate that these conditions are outside of the task space.

In contrast, $\tilde{B} \setminus B$ represents the user-inferred task space which is not within the intended task space. For instance, in a hypothetical fruit handling task, fruits which the user dislikes are first discarded by the robot. Then, the robot peels the remaining fruits. In this scenario, the robot does not need to be taught how to peel an orange if it would have already been discarded. By alerting the teacher of this, this could avoid wasting time and resources training the robot, thereby also enhancing the teaching efficiency η. Realistically, however, the user-inferred task space outside the actual task space would be much smaller than $B \setminus \tilde{B}$.

4 ASSESSING THE TEACHER'S INTERPRETATION OF THE ROBOT LEARNING STATE

The interpretability of the robot's learning state is important in order for the human teacher to teach effectively. There have been attempts at evaluating the explainability in artificial intelligence (AI) [49]; however, these are for explaining to a user why a particular single decision has been made by the agent. Although there have

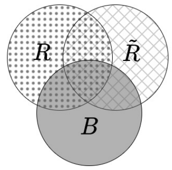

Figure 7: The task B, the robot's actual learning state R and the teacher's estimation of the robot's learning state $\tilde{R}$.

been several methods introduced to help the human teacher keep up with the robot's learning state, to our knowledge there has not been any literature on how the teacher's interpretation of the robot learner's state can be assessed quantitatively.

Here we draw conclusions from the models described from earlier works [3, 47] to build a model of how interpretation can be assessed. We show in Fig. 7 the Venn diagram of B, R and $\tilde{R}$ together. As addressed by the authors in [47], practically, it is difficult to measure R as the space could be unfathomably large. Since the robot will only be in the environment as defined within B, space outside it can be ignored. $R \cap \tilde{R} \cap B$ describes the set of tasks $\in B$ the robot is able to perform correctly both in reality *and* the user's beliefs – which can be termed as true positives. $B \setminus (R \cup \tilde{R})$ describes the set of tasks that the robot is currently unable to perform correctly both in reality and the user's beliefs (true negatives). False positives, or $\tilde{R} \cap B \setminus R$, are the set of tasks whereby the user incorrectly assumes the robot would be able to perform correctly. This translates to a lower teaching efficacy ϵ, possibly resulting in the robot making mistakes during the actual task execution, which could waste resources and impose danger. Lastly, $R \cap B \setminus \tilde{R}$ represents false negatives – the tasks in which the teacher underestimates the robot. Although this sounds ideal, it may instead cause the teacher to provide unnecessary demonstrations, decreasing the teaching efficiency η. This could result in greater time and resources spent than necessary.

Theoretically, the human teacher's interpretation accuracy could then be calculated as follows:

$$i_a = \frac{|R \cap \tilde{R} \cap B| + (B \setminus (R \cup \tilde{R}))|}{|B|}, i_a \in [0, 1] \tag{8}$$

In practice, it may not be possible to measure the task space B depending on how large it is. Hence, more realistically, an approximated task space $\hat{B}$ should be used for tasks with large $|B|$. This may be done by sampling of the task space using Monte Carlo methods, for example. The user then predicts whether the robot can perform each of the tasks in $\hat{B}$ correctly. Upon the robot's sampled task executions, we can then obtain the numbers of true positives, true negatives, false positives, and false negatives.

Due to the definitions provided for this model, the assessment is reduced to simple binary questions: yes or no to whether the robot has learned sufficiently to perform a specific task correctly. To more accurately assess the teacher's interpretation, the assessment should take place regularly throughout the teaching process. To account for the uneven class distribution such as when the robot is unable to perform the majority of the task space, it may be beneficial to use metrics such as the F1 score rather than accuracy.

This method of evaluating the interpretability is limited to whether the task can be executed correctly, but not how close it is to being correct, or if it is optimal. It also does not evaluate the reasoning in the decision-making of the robot. In this work, we have limited the assessment to these measures only as they are quantitative and straightforward.

5 CASE STUDY

In this section, we describe a study we undertook illustrating some of our aforementioned suggestions. We demonstrate how UI interventions can improve robot teaching via improving the teacher's estimation of the robot's mental model and learning process. We also demonstrate how the proposed method of the interpretation accuracy assessment can take place. Note that here we are not comparing MR and non-MR interfaces, but rather the two different UI intervention approaches.

We developed an MR application for the Microsoft HoloLens 2 head-mounted display (HMD) device using Unity. The virtual robot used in the application is the Kinova Gen3 robot arm. We introduced a simple task of teaching the robot to sort objects into one of the two bins according to shape and colour. Objects which were light-coloured and contained edges were to be sorted into one bin, while the others were to be put into the other.

The objects were created virtually on Unity due to the convenience of the creation of different shapes and colours. The shapes were: cube, sphere and spheroid. The colours were: blue, black and white. There were therefore nine unique combinations of shape and colour. There were two of each combination of shape and colour, resulting in 18 objects altogether.

5.1 Participants and Protocol

We recruited 30 participants (M=16, F=14) with no background in AI or robotics to demonstrate a sorting task to the robot via the HoloLens 2 headset. The research was approved by the University's Human Research Ethics Team. The participants' ages ranged between 18 and 32 (M=22.83, SD=3.69). They were not told that there were duplicates for each of the unique objects. The robot learned by using demonstrations to build an ID3 (Iterative Dichotomiser 3) decision tree. We split the participants into three groups: No Feedback, Decision Tree and Demonstration Level. The participants were briefed on the teaching task, how to operate the device, how to read the Decision Tree (if applicable), and the definitions of the Demonstration Levels (if applicable).

The No Feedback (NF) Group: This was the control group. The participants were not shown any extra information from the robot (Fig. 8a).

The Decision Tree (DT) Group: The participants were shown in real-time the robot's decision tree graph throughout the whole teaching process (Fig. 8b). This was to illustrate to the human teacher the current thought process of the learning robot. This information was anchored onto the user's HMD screen such that it stayed in the bottom left corner of the user's view.

The Demonstration Level (DL) Group: Each of the objects was labelled to the participant as one of the three options: "Demonstrated", "Not Demonstrated" and "Partially Demonstrated" (Fig. 8c). This was to indicate whether an object with the same shape and

(a) No Feedback **(b) Decision Tree** **(c) Demonstration Level**

Figure 8: Different modalities of robot teaching in the experiment

colour had previously been sorted by the user. These labels were chosen as it is common in human-human teaching scenarios where the student can clearly say that they are able to do a certain question as they have been shown how to do it before. "Demonstrated" indicates that an object with the exact same shape and colour has been demonstrated before. Relating to a human teaching scenario, this is akin to the student expressing that s/he knows exactly how to solve the problem since the teacher has shown how to do this exact same problem before. "Partially demonstrated" means the robot's decision tree does not contain the path for that exact combination of shape and colour of the object, but has a path for one of the features. This is similar to the student encountering a problem they have not yet seen but it resembles a similar problem they have seen before, and the student has an idea of how they may solve it. "Not demonstrated" means the robot's decision tree does not currently contain a single path to sort this object. Demonstration Levels were used here as we would predict that this could make it easier for the human teacher to pick which objects s/he should teach to the robot to sort, compared to not having any UI interventions.

By providing information about the robot's learning state via the interface, this may allow for the user to select better demonstrations for the robot. As the DT condition shows the entire learning tree for the robot, this may mean that it is the most interpretable. We therefore posited the following hypotheses:

H1: The control (NF) group will have lower robot performance than the intervention groups (DT and DL)
H2: The control (NF) group will have lower teaching efficiency than the intervention groups
H3: The control (NF) group will have lower interpretability than the intervention groups
H4: The DT group will have the highest robot performance
H5: The DT group will have the highest teaching efficiency
H6: The DT group will have highest interpretability

5.2 Measures

To assess the robot performance, both the F1 score and accuracy were recorded after nine demonstrations had been completed. Although the F1 score and accuracy values displayed similar trends, we decided to use the F1 score as it is more appropriate for binary classification tasks.

For the evaluation of the teaching efficiency, we recorded the number of teaching demonstrations required for the robot to consistently sort all objects correctly.

The user's interpretation of the robot learning state was assessed twice: after five demonstrations, and after ten demonstrations. For the assessment, the participants were shown the nine unique objects and asked whether at that point the robot would be able to sort each one correctly. The selectable answers were "yes" and "no". The F1 score was then calculated for each participant.

5.3 Results

To analyse the number of teaching demonstrations required, the interpretation accuracy, and agent performance for the three groups, we used the Kruskal-Wallis H-test for the preliminary statistical test and the Dunn's Test for the post-hoc test.

5.3.1 Performance. There were significant statistical differences between the F1 Scores of the three groups ($H(2) = 14.36, p < .001$). Post-hoc pairwise comparisons using Dunn's test indicated that F1 scores for the NF group (M=0.69 SD=0.29) were significantly lower than those of the intervention groups ($p < .001$ and p=.013 for DT and DL groups respectively) – see Fig. 9a. No statistically significant differences were observed between the F1 scores of the DT (M=0.94 SD=0.07) and DL (M=0.97 SD=0.05) groups ($p = .35$). Overall, our results support H1 but fail to support H4.

5.3.2 Number of Demonstrations Required. There were significant statistical differences between the numbers of demonstrations required of the three groups ($H(2) = 17.34, p < .001$). As can be observed in Fig. 9b, the numbers of demonstrations required to complete the teaching for the NF group (M=14.30 SD=2.83) were significantly greater than those of the intervention groups ($p < .001$ and $p = .0017$ for DT and DL groups respectively). Our results therefore support H2. The number of demonstrations required for the DT (M=8.70, SD=1.64) and DL (M=9.00, SD=1.15) groups did not differ significantly ($p = .63$) and hence this fails to support H5.

5.3.3 Interpretability of the Robot. There were significant statistical differences between the robot interpretability of the three groups ($H(2) = 18.35, p < .001$). As can be observed in Fig. 9c, the interpretability scores for the NF group (M=0.77 SD=0.13) were significantly lower than those of the intervention groups ($p < .001$ and p=.018 for DT and DL groups respectively). The interpretability scores for the DT (M=0.98, SD=0.02) and DL (M=0.94, SD=0.03) groups did not differ significantly ($p = .07$). Overall, our results support H3 but fail to support H6.

Overall, it is clear that implementing additional information regarding the robot can increase interpretability during teaching,

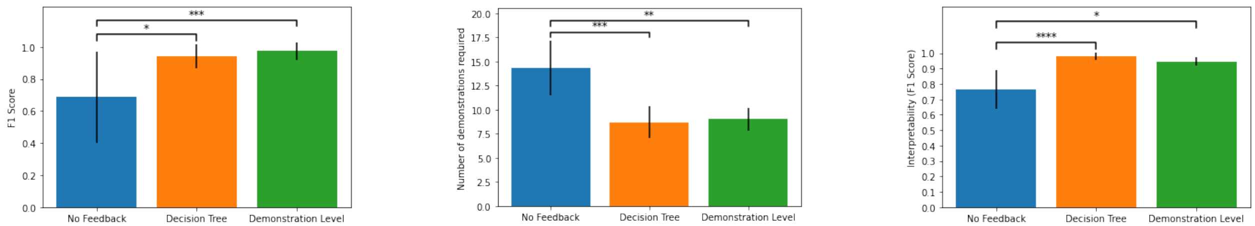

(a) Robot performance (F1 Score) **(b) Number of demonstrations required** **(c) Robot interpretability (F1 Score)**

Figure 9: The robot's performance, the teaching efficiency and the interpretability measured across the three modes

generate higher agent performance and increase teaching efficiency. It was as expected that showing labels of whether the object had been demonstrated prior would improve the teaching efficiency as it put less of the burden on the teacher to memorise what examples they had taught to the robot. In practice, several thousands of examples may need to be given to the robot and hence the teacher should not always be expected to remember. Overlaying information in the robot's environment, seems to be a good way to avoid the need for the teacher to remember or look up the status of each object in a list.

Another interesting observation to note was that the participants did not completely understand how decision trees work. For instance, before the briefing on decision trees, many did not understand why a decision tree would not show *all* attributes. This shows that in order to make teaching robots accessible to novices, these graphical representations may in fact still be too complex. This emphasises the need for a more novice-friendly robot data visualisation.

6 CONCLUSION

In this work, we presented possible UI interventions that could be introduced to improve the human-robot teaching process. We proposed how the teacher's interpretation of the robot learner can be assessed quantitatively. We developed a teaching interface with some of our proposed interventions and showed how using these interventions to display the robot's decision tree or the demonstration levels of tasks can improve the interpretability, robot performance and teaching efficiency.

Although we demonstrated how some UI interventions could improve interpretability, teaching efficiency and robot performance, we did not test other methods such as whether showing the robot's reach range could improve the teaching process. These proposed methods will require testing to obtain empirical evidence. Moreover, for each of the methods introduced for improving the teaching process, there could be several different approaches. For example, when showing the robot's mental model, its uncertainty about a task could be shown, or the objects within its world model could be labelled visually to the user. Even then, if we are to show the robot's uncertainties, there could be several different ways of doing so. How, for example, would we be showing uncertainties to the layperson – through percentages, binary values, or multiple discrete levels? Would we be using numbers, or perhaps a smiling or frowning face icon?

The learning model of the robot in our case study was a simple decision tree. However, with a more complex decision tree, it may in fact not improve the human-robot teaching due to the lack of readability. It would also be beneficial to explore how this could translate to representations of other interpretable models such as rules [30] or sparse linear models [52]. Based on our descriptions of UI interventions, more empirical research should be conducted to build a visual grammar and guidelines for interfaces designed to assist humans in teaching robots.

ACKNOWLEDGMENTS

This research is partially supported by the Australian Research Council Discovery Early Career Research Award (Grant No.DE210100858).

REFERENCES

[1] Syed Mohsin Abbas, Syed Hassan, and Jongwon Yun. 2012. Augmented reality based teaching pendant for industrial robot. In *2012 12th International Conference on Control, Automation and Systems*. 2210–2213.
[2] Muneeb Imtiaz Ahmad, Jasmin Bernotat, Katrin Lohan, and Friederike Eyssel. 2019. Trust and cognitive load during human-robot interaction. *arXiv preprint arXiv:1909.05160* (2019).
[3] Erik A Billing and Thomas Hellström. 2010. A formalism for learning from demonstration. *Paladyn, Journal of Behavioral Robotics* 1, 1 (2010), 1–13.
[4] Maya Cakmak, Crystal Chao, and Andrea L Thomaz. 2010. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development* 2, 2 (2010), 108–118.
[5] Maya Cakmak and Andrea L. Thomaz. 2012. Designing robot learners that ask good questions. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 17–24. https://doi.org/10.1145/2157689.2157693
[6] Sylvain Calinon and Aude Billard. 2007. Active teaching in robot programming by demonstration. In *RO-MAN 2007-The 16th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 702–707.
[7] Mehmet Ege Cansev, Honghu Xue, Nils Rottmann, Adna Bliek, Luke E. Miller, Elmar Rueckert, and Philipp Beckerle. 2021. Interactive Human–Robot Skill Transfer: A Review of Learning Methods and User Experience. *Advanced Intelligent Systems* 3, 7 (2021), 2000247. https://doi.org/10.1002/aisy.202000247
[8] Joyce Y Chai, Maya Cakmak, Candace Sidner, and J Lupp. 2017. Teaching robots new tasks through natural interaction. In *Interactive Task Learning: Agents, Robots, and Humans Acquiring New Tasks through Natural Interactions, Strüngmann Forum Reports, J. Lupp, series editor*, Vol. 26.
[9] Crystal Chao, Maya Cakmak, and Andrea L Thomaz. 2010. Transparent active learning for robots. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 317–324.
[10] Jason Chen and Alex Zelinsky. 2003. Programing by demonstration: Coping with suboptimal teaching actions. *The International Journal of Robotics Research* 22, 5 (2003), 299–319.
[11] Xiaogang Chen, Xiaoshan Huang, Yijun Wang, and Xiaorong Gao. 2020. Combination of Augmented Reality Based Brain- Computer Interface and Computer Vision for High-Level Control of a Robotic Arm. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 28, 12 (Dec. 2020), 3140–3147.
[12] Sonia Chernova and Manuela Veloso. 2009. Interactive policy learning through confidence-based autonomy. *Journal of Artificial Intelligence Research* 34 (2009), 1–25.

[13] J. W. S. Chong, S. K. Ong, A. Y. C. Nee, and K. Youcef-Youmi. 2009. Robot programming using augmented reality: An interactive method for planning collision-free paths. *Robotics and Computer-Integrated Manufacturing* 25, 3 (June 2009), 689–701. https://doi.org/10.1016/j.rcim.2008.05.002

[14] Maximilian Diehl, Alexander Plopski, Hirokazu Kato, and Karinne Ramirez-Amaro. 2020. Augmented Reality interface to verify Robot Learning. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 378–383. https://doi.org/10.1109/RO-MAN47096.2020.9223502

[15] Jared Alan Frank, Sai Prasanth Krishnamoorthy, and Vikram Kapila. 2017. Toward Mobile Mixed-Reality Interaction With Multi-Robot Systems. *IEEE Robotics and Automation Letters* 2, 4 (Oct. 2017), 1901–1908. https://doi.org/10.1109/LRA.2017.2714128 Conference Name: IEEE Robotics and Automation Letters.

[16] Chris Frith and Uta Frith. 2005. Theory of mind. *Current biology* 15, 17 (2005), R644–R645.

[17] Nora Fronemann. 2022. Should my robot know what's best for me? Human–robot interaction between user experience and ethical design. (2022), 17.

[18] Fabrizio Ghiringhelli, Jérôme Guzzi, Gianni A. Di Caro, Vincenzo Caglioti, Luca M. Gambardella, and Alessandro Giusti. 2014. Interactive Augmented Reality for understanding and analyzing multi-robot systems. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 1195–1201. https://doi.org/10.1109/IROS.2014.6942709

[19] 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* (2022).

[20] Gavin Hackeling. 2017. *Mastering Machine Learning with scikit-learn*. Packt Publishing Ltd.

[21] Jared Hamilton, Thao Phung, Nhan Tran, and Tom Williams. 2021. What's The Point? Tradeoffs Between Effectiveness and Social Perception When Using Mixed Reality to Enhance Gesturally Limited Robots. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21)*. Association for Computing Machinery, New York, NY, USA, 177–186. https://doi.org/10.1145/3434073.3444676

[22] Juan David Hernández, Shlok Sobti, Anthony Sciola, Mark Moll, and Lydia E. Kavraki. 2020. Increasing Robot Autonomy via Motion Planning and an Augmented Reality Interface. *IEEE Robotics and Automation Letters* 5, 2 (April 2020), 1017–1023. https://doi.org/10.1109/LRA.2020.2967280

[23] Robert R Hoffman, Shane T Mueller, Gary Klein, and Jordan Litman. 2018. Metrics for explainable AI: Challenges and prospects. *arXiv preprint arXiv:1812.04608* (2018).

[24] Tetsunari Inamura. 1999. Acquisition of probabilistic decision model based on interactive teaching method. In *Proc. 9th Int. Conf. on Advanced Robotics, Tokyo, 1999*. 523–528.

[25] Tetsunari Inamura, Naoki Kojo, and Masayuki Inaba. 2006. Situation recognition and behavior induction based on geometric symbol representation of multimodal sensorimotor patterns. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5147–5152.

[26] Masato Ito, Kuniaki Noda, Yukiko Hoshino, and Jun Tani. 2006. Dynamic and interactive generation of object handling behaviors by a small humanoid robot using a dynamic neural network model. *Neural Networks* 19, 3 (2006), 323–337.

[27] Taylor A. Kessler Faulkner and Andrea Thomaz. 2021. Interactive Reinforcement Learning from Imperfect Teachers. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, Boulder CO USA, 577–579. https://doi.org/10.1145/3434074.3446361

[28] Nishanth Kumar, Eric Rosen, and Stefanie Tellex. 2019. Knowledge Acquisition for Robots Through Mixed Reality Head-Mounted Displays. (2019). https://doi.org/10.13140/RG.2.2.23780.83843

[29] Jens Lambrecht and Jörg Krüger. 2014. Spatial programming for industrial robots: efficient, effective and user-optimised through natural communication and augmented reality. In *Advanced Materials Research*, Vol. 1018. Trans Tech Publ, 39–46.

[30] Benjamin Letham, Cynthia Rudin, Tyler H McCormick, and David Madigan. 2015. Interpretable classifiers using rules and bayesian analysis: Building a better stroke prediction model. *The Annals of Applied Statistics* 9, 3 (2015), 1350–1371.

[31] Hangxin Liu, Yaofang Zhang, Wenwen Si, Xu Xie, Yixin Zhu, and Song-Chun Zhu. 2018. Interactive Robot Knowledge Patching Using Augmented Reality. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. 1947–1954.

[32] Matthew B. Luebbers, Connor Brooks, Carl L. Mueller, Daniel Szafir, and Bradley Hayes. 2021. ARC-LfD: Using Augmented Reality for Interactive Long-Term Robot Skill Maintenance via Constrained Learning from Demonstration. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, Xi'an, China, 3794–3800. https://doi.org/10.1109/ICRA48506.2021.9561844

[33] Stéphane Magnenat, Morderchai Ben-Ari, Severin Klinger, and Robert W. Sumner. 2015. Enhancing Robot Programming with Visual Feedback and Augmented Reality. In *Proceedings of the 2015 ACM Conference on Innovation and Technology in Computer Science Education*. ACM, Vilnius Lithuania, 153–158. https://doi.org/10.1145/2729094.2742585

[34] Paul Milgram and Fumio Kishino. 1994. A taxonomy of mixed reality visual displays. *IEICE TRANSACTIONS on Information and Systems* 77, 12 (1994), 1321–1329.

[35] Paul Milgram, Anu Rastogi, and Julius J Grodski. 1995. Telerobotic control using augmented reality. In *Proceedings 4th IEEE International Workshop on Robot and Human Communication*. IEEE, 21–29.

[36] James F. Mullen, Josh Mosier, Sounak Chakrabarti, Anqi Chen, Tyler White, and Dylan P. Losey. 2021. Communicating Inferred Goals With Passive Augmented Reality and Active Haptic Feedback. *IEEE Robotics and Automation Letters* (Oct. 2021).

[37] S. K. Ong, A. Y. C. Nee, A. W. W. Yew, and N. K. Thanigaivel. 2020. AR-assisted robot welding programming. *Advances in Manufacturing* 8, 1 (March 2020), 40–48.

[38] Yun Suen Pai, Hwa Jen Yap, and Ramesh Singh. 2015. Augmented reality–based programming, planning and simulation of a robotic work cell. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture* 229, 6 (2015), 1029–1045.

[39] Ornnalin Phaijit, Mohammad Obaid, Claude Sammut, and Wafa Johal. 2022. A Taxonomy of Functional Augmented Reality for Human-Robot Interaction. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 294–303.

[40] Ornnalin Phaijit, Claude Sammut, and Wafa Johal. 2022. Let's Compete! The Influence of Human-Agent Competition and Collaboration on Agent Learning and Human Perception. In *Proceedings of the 10th International Conference on Human-Agent Interaction*. 86–94.

[41] David Puljiz, Erik Stöhr, Katharina S. Riesterer, Björn Hein, and Torsten Kröger. 2019. General Hand Guidance Framework using Microsoft HoloLens. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 5185–5190. https://doi.org/10.1109/IROS40897.2019.8967649

[42] Camilo Perez Quintero, Sarah Li, Matthew KXJ Pan, Wesley P. Chan, H.F. Machiel Van der Loos, and Elizabeth Croft. 2018. Robot Programming Through Augmented Trajectories in Augmented Reality. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1838–1844. https://doi.org/10.1109/IROS.2018.8593700

[43] Christopher Reardon, Jason Gregory, Carlos Nieto-Granda, and John G. Rogers. 2020. Enabling Situational Awareness via Augmented Reality of Autonomous Robot-Based Environmental Change Detection. In *International Conference on Human-Computer Interaction*, Jessie Y. C. Chen and Gino Fragomeni (Eds.). Springer, 611–628.

[44] Christopher Reardon, Kevin Lee, John G. Rogers, and Jonathan Fink. 2019. Communicating via Augmented Reality for Human-Robot Teaming in Field Environments. In *2019 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. 94–101.

[45] Brian D Ripley. 2007. *Pattern recognition and neural networks*. Cambridge university press.

[46] Eric Rosen, David Whitney, Michael Fishman, Daniel Ullman, and Stefanie Tellex. 2020. Mixed Reality as a Bidirectional Communication Interface for Human-Robot Interaction. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 11431–11438. https://doi.org/10.1109/IROS45743.2020.9340822

[47] Aran Sena and Matthew Howard. 2020. Quantifying teaching behavior in robot learning from demonstration. *The International Journal of Robotics Research* 39, 1 (2020), 54–72.

[48] Rossitza Setchi, Maryam Banitalebi Dehkordi, and Juwairiya Siraj Khan. 2020. Explainable Robotics in Human-Robot Interactions. *Procedia Computer Science* 176 (Jan. 2020), 3057–3066. https://doi.org/10.1016/j.procs.2020.09.198

[49] Kacper Sokol and Peter Flach. 2020. Explainability fact sheets: a framework for systematic assessment of explainable approaches. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 56–67.

[50] John D Sweeney and Rod Grupen. 2007. A model of shared grasp affordances from demonstration. In *2007 7th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 27–35.

[51] Andrea L. Thomaz and Cynthia Breazeal. 2008. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence* 172, 6 (April 2008), 716–737.

[52] Berk Ustun and Cynthia Rudin. 2016. Supersparse linear integer models for optimized medical scoring systems. *Machine Learning* 102, 3 (2016), 349–391.

[53] Michael Walker, Hooman Hedayati, Jennifer Lee, and Daniel Szafir. 2018. Communicating Robot Motion Intent with Augmented Reality. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. 316–324.

[54] Michael E. Walker, Hooman Hedayati, and Daniel Szafir. 2019. Robot Teleoperation with Augmented Reality Virtual Surrogates. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 202–210.

[55] Chao Wang and Anna Belardinelli. 2022. Investigating explainable human-robot interaction with augmented reality. (2022).

[56] Yanxin Wang, Hong Zeng, Aiguo Song, Baoguo Xu, Huijun Li, Lifeng Zhu, Pengcheng Wen, and Jia Liu. 2017. Robotic arm control using hybrid brain-machine interface and augmented reality feedback. In *2017 8th International IEEE/EMBS Conference on Neural Engineering (NER)*. 411–414. https://doi.org/10.1109/NER.2017.8008377 ISSN: 1948-3554.

[57] Tom Williams, Matthew Bussing, Sebastian Cabrol, and Ian Lau. 2019. Toward allocentric mixed-reality deictic gesture. *VAM-HRI* (2019).

[58] Chung Xue, Yuansong Qiao, and Niall Murray. 2020. Enabling Human-Robot-Interaction for Remote Robotic Operation via Augmented Reality. In *2020 IEEE 21st International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*. 194–196. https://doi.org/10.1109/WoWMoM49955.2020.00046

[59] Peng Zang, Runhe Tian, Andrea L Thomaz, and Charles L Isbell. 2010. Batch versus interactive learning by demonstration. In *2010 IEEE 9th International Conference on Development and Learning*. IEEE, 219–224.

[60] Peng Zhang. 2010. Human-machine interfaces. *Advanced Industrial Control Technology* 1 (2010), 527–555.