Improving Robot Learning from Demonstration via Competitive Interactions and User Interface Interventions

Ornnalin Phaijit[1], Claude Sammut[1], Wafa Johal[2]
[1]The University of New South Wales, Australia, [2]The University of Melbourne, Australia
l.phaijit@unsw.edu.au, c.sammut@unsw.edu.au, wafa.johal@unimelb.edu.au

Abstract

For novice users, teaching robots by demonstration can be a challenging and often frustrating process due to communication barriers and the repetitive nature of the task. Competitive interactions and gamification have been successfully applied in fields such as education to enhance engagement, yet their potential in human-robot teaching remains largely unexplored. In this paper, we investigate how these methods can improve the teaching experience for novice users by conducting an experiment in which participants taught a robot to perform two tasks: 1) grasping objects and 2) sorting them. Twenty participants took part in a lab study in which they experienced teaching a robot in a gamified competitive way or not. Our results show that competitive human-robot interactions significantly increased teaching efficiency, enhanced the perceived competence of the robot and increased user enjoyment. Additionally, participants were shown both decision trees and tables representing the robot's learning process, and we collected comparative feedback on these visualisations. Qualitative responses highlighted the importance of familiarity with the visualisations and the need for clear guidance on interpreting them. These findings suggest that integrating competitive elements and intuitive visual representations can enhance the human-robot teaching experience for novice users, making it more engaging and effective.

1 Introduction

Robots are increasingly being adopted by users with little to no experience in programming or robotics. Since it is impractical to pre-program robots for every possible task or user preference before deployment, it is crucial that these novice users are able to teach robots to perform tasks in ways that align with their specific needs and preferences. To address this challenge, interactive robot learning has emerged as a promising approach, enabling robots to adapt and learn from human input through direct interactions.

In recent years, research on interactive robot learning has largely focused on improving the robot's performance, often overlooking the human teacher's role in the process. Since the effectiveness of robot learning heavily depends on the human teacher, it is crucial to examine the teacher's experience in order to improve human-in-the-loop robot learning systems [Amershi et al., 2014]. User perception plays a key role, not only in robot performance but also in fostering user trust, which influences engagement and overall satisfaction. A robot perceived as competent increases trust, which can lead to greater user acceptance and a more enjoyable training process [Bryant et al., 2020] [Hancock et al., 2011]. Trust is critical in human-robot interaction, as users are more likely to rely on systems they trust and abandon those they do not [Lewandowsky et al., 2000] [Muir and Moray, 1996]. When users are aware of a system's potential faults, their trust remains intact, provided the system is predictable and transparent in its learning process [Lewis et al., 2018].

Furthermore, the repetitive and often monotonous nature of training by demonstration can lead to human disengagement, negatively impacting task performance and teaching quality [Pekrun et al., 2010]. The novelty of interacting with robots tends to fade quickly, leading many users to abandon their systems due to a lack of sustained engagement [Robinson et al., 2019], [Gockley et al., 2005][De Graaf et al., 2017]. This underscores the need for more engaging and efficient training methods.

Another key challenge in interactive robot learning is that a teacher is typically required to have a good understanding of the learning robot's state in order to teach it tasks effectively. As the human and the robot do not share the same mental model or language, especially when the human teacher is a robotics novice, this can make the teaching task inefficient, worsening the teacher's experience [Bilal et al., 2024]. In order to address this, several approaches to convey the robot learning state have been proposed such as displaying the robot's decision-making using And-Or graphs [Liu et al., 2021] and decision trees [Phaijit et al., 2023], visualising estimated trajectories [Sena and Howard, 2020], or allowing the robot to ask questions [Cakmak et al., 2010; Chao et al., 2010; Habibian et al., 2022; Cakmak and Thomaz, 2012]. Although these methods have been shown to be effective, we currently lack more insightful qualitative data from the novice users. For instance, in the study by [Phaijit et al., 2023], decision trees and demonstration levels were shown to lead to similar teaching efficiencies. However, users were not asked regarding their personal preferences such as which one they preferred or found the most helpful.

In this paper, we present an experiment comparing competitive and non-competitive human-robot interactions during robot Learning from Demonstration (LfD) in two common household tasks: grasping and sorting. We explore how competitive interactions can enhance perceived robot competency, user enjoyment, and teaching efficiency by reducing the number of demonstrations required. Additionally, we gather qualitative data on users' opinions of two different methods for visualising the robot's learning process: decision trees and tabular data.

2 Related Work

2.1 Human-Robot Competitive Interaction

Human-robot interactions in learning environments can be collaborative or competitive. While the effectiveness of human advisory teaching is well-established, adversarial teaching, or human-robot competition in teaching, remains relatively new. Duan et al. [Duan et al., 2020] found that adversarial input improved robot grasping performance compared to learning without human intervention. However, participants often struggled to apply perturbations effectively, despite being instructed to do so. To address this, Yoon et al. [Yoon and Nikolaidis, 2020] introduced a framework that distinguished between collaborative and adversarial interactions, ensuring the robot learned only from adversarial interventions, resulting in improved performance. Hamaya et al. [Hamaya et al., 2021] further demonstrated that combining human advisory and adversarial input in a robotic peg-in-hole task enhanced performance compared to random or no interactions. The existing research focuses on reinforcement learning, requiring humans to determine when and how to intervene. In contrast, Learning from Demonstration (LfD) may simplify this process by allowing the human to perform the task directly, without needing to deliberately advise or hinder the agent.

Although competitive interactions and gamifications are vastly employed in human educational tasks due to their strong evidence of increasing user engagement [Caponetto et al., 2014], there is however a lack of research in how these methods could be used to assist in the human's robot teaching task. [Phaijit et al., 2022b] discovered that competitive interactions between the human and the agent allowed for a more enjoyable experience for the novice human during the task of teaching the agent to play Pacman. The participants reported that they felt the agent in the competitive mode was more competent. While this study offered a valuable step towards increasing user engagement in human-robot teaching, it was limited in two key ways: 1) it was a gaming task, which differs from the practical tasks humans typically teach robots in home environments, 2) the learning embodied agent was entirely virtual, which may differ from competitive interactions involving a physical robot.

2.2 Visual representation of the robot learning state

In order for the human teacher to understand the robot's current learning state, there must be robot feedback shown to the user. This can be represented in many forms such as graphical [Liu et al., 2018] or semantic representations [Diehl et al., 2020; Wang and Belardinelli, 2022] of the robot's understanding and learning process for the task. It can be used to "augment human perception of the robot" [Phaijit et al., 2022a] to allow for better human-robot interaction.

Decision trees have been the chosen approach for many applications due to their ease of readability for users. Although these are commonly used by machine learning practitioners, little research has been conducted on the perspective of a non-expert user. In [Phaijit et al., 2022b], the authors mentioned that non-expert participants had difficulties understanding decision trees. For instance, novice users did not understand why the tree may not show all the attributes. In order to build a system to be used by novices, it is crucial that we heavily take their perspectives into consideration.

In contrast to decision trees, information represented in the form of tabular data is common in many non-technical fields. The tabular form requires all records to share the same features and all values must be numerical, categorical, or Boolean [Sahakyan et al., 2021]. Although tabular data, or tables, are widely used amongst people of all backgrounds, they are also commonly used for robot and machine learning [Shwartz-Ziv and Armon, 2022; Gorishniy et al., 2021; Borisov et al., 2022].

Decision trees and tables are some of the most interpretable forms of information. However, there is limited research comparing the two from the perspective of human novices in Learning from Demonstration. This highlights the need for deeper insights into how novices perceive and engage with these different formats.

3 Research Questions

The user experience is comprised of pragmatic and hedonistic qualities [Cansev et al., 2021; Fronemann, 2022]. This means in a robot teaching process, the user experience may be affected by pragmatic components such as how well does the robot execute the task once taught? How slow was the robot at learning? These questions focus on the practical points of view of the experience for the user. On the other hand, questions such as how enjoyable or friendly is the interaction with the robot provide the hedonistic insights to the user experience.

With the current unknowns in how competitive human-robot interactions fare in robot learning, here we seek to evaluate their effects. Although gamification has positive effects in entertainment (such as video games) and education for humans, could the same be applied to human-in-the-loop robot learning where the human takes on the role of a teacher?

We therefore propose to explore the following:

1. Does the competitive human-robot interaction improve the robot learning process in terms of:
   (a) performance?
   (b) user preferences?
2. What are novice users' opinions on decision trees and tabular data on conveying the robot's learning process?

4 Experiment

In this work, we have both decision trees and tabular information to assist the user with the robot teaching task.

4.1 Tasks

The teaching is split into two different tasks: 1) teaching the 7 degree of freedom Kinova Gen3 robot how to grasp objects, and 2) teaching it to sort the objects into the correct bowl. These tasks have been selected as they can be seen as common household chores: picking up objects as a way of tidying up the household, and putting them away in the correct locations. This teaching scenario has been simplified to allow for the robot to be able to fully learn the task from scratch from the novice within the experiment.

Grasping

Firstly, in the grasping task, the user selects an object to teach the robot out of the two choices: a short cylinder or a long cuboid. Upon selection, the participant places the object in front of the robot at a fixed predetermined location, with the object orientation the user wishes to teach the robot to be able to grasp. Once the user is satisfied with the selected object and its orientation, they press confirm in the computer program interfacing with the robot. They are then to select 1) the approach angle, 2) the gripper rotation, and 3) the gripper width (Figure 1) of the robot. For ease of understanding, as the user varies these values, the robot can be seen to physically adjust according to these settings. Once the desired values have been selected, the user can prompt the robot via the computer program to attempt the grasp with the selected values. The robot then performs the taught demonstrations and the user may either press confirm to add this to the list of demonstrations for the robot, or re-do the demonstration.

Sorting

The procedure is similar to the grasping task. However, the user is only required to select the object, out of the 24 objects available, that they wish to teach the robot to sort. Then, the user is required to select the bowl for the object to be placed in: *keep* or *throw away* bowls, according to the shape and colour of the object. The objects have the following shapes: 1) boxes, 2) cups, and 3) cylinder blocks. The colours are: 1) red 2) yellow 3) orange, and 4) purple.

4.2 Competitive vs Non-Competitive Modes

The teaching activity is split into two different modes: competitive and non-competitive. In the competitive mode, after the user has selected the object item as well as its orientation to teach the robot, the program calculates the similarity between this example and the previously taught examples. If the similarity falls below the threshold value, the user gets a point and this is displayed in the score tally. On the other hand, if the similarity exceeds the threshold, the robot receives a point. As the number of demonstrations provided to the robot increases, the dissimilarity between a new demonstration and the existing set of taught demonstrations is expected to decrease. Hence the formula for the threshold was selected such that it would decrease with the increasing number of demonstrations: $y = 2.5e^{-0.9-0.3x} + 0.1$. The threshold values were selected somewhat arbitrarily – it was not too challenging to best the robot, yet selecting a demonstration carelessly could mean the robot would win. The optimisation of the threshold value is out of this work's scope.

For the grasping task, the similarity is calculated using Manhattan distance between the new example's object shape, orientation and position compared to those of the previously taught examples. For the sorting task, the

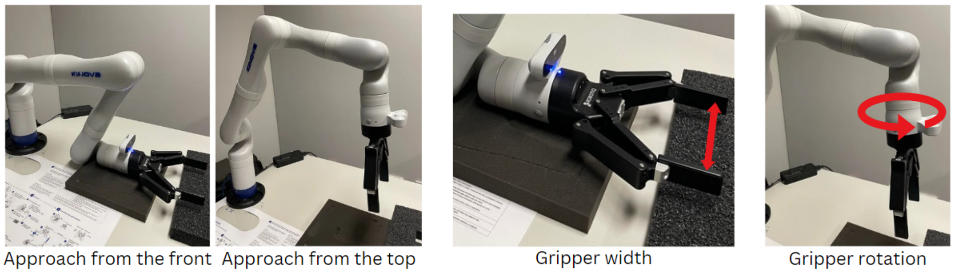

Figure 1: The inputs for the grasping task: 1) the approach angle where the user can choose for the robot to approach the object from the front or from the top, 2) the gripper width, and 3) the gripper rotation

similarity is obtained simply by checking if the exact shape and colour combination have been taught before.

4.3 User Interface Intervention

The two user interface (UI) interventions selected are: decision trees and table. Both data forms are displayed to the user in the interface (Figure 2) throughout the teaching process.

In the grasping task, three decision trees are displayed for making decisions on the approach angle, gripper rotation and gripper width. Each row of the table (e.g. Figure 3) contains information on each demonstration: the parameters of the input (object shape, object position, object orientation, approach angle, gripper rotation and gripper width) and the output (approach angle, gripper rotation and gripper width).

In the sorting task, only one decision tree is shown for making a decision on the bowl the object is sorted into. The table displays the inputs (object shape and colour) and the bowl it is to be sorted into.

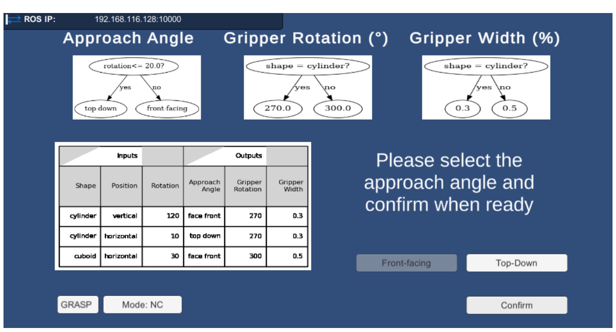

Figure 2: Example screenshot of the robot teaching interface

4.4 System Architecture

Unity Version 2020.3.1f1 was used for developing the computer software that the user would be interacting with the robot via. The robot functionality was built using Robot Operating System (ROS). The ROS package ROS TCP Endpoint was used to bridge the communication between Unity and ROS – enabling the robot and the user interface to communicate (Figure 5). Unique visual tags attached to objects were detected by the robot's RGB-D camera using the *ar_track_alvar* ROS package. The robot's vision module was supported using the *ros_kortex_vision* ROS package. MoveIt Motion Planning Framework was used to assist with the robot's motion planning. The robot uses CART (Classification and Regression Tree) algorithm for its decision-making.

Inputs			Outputs		
Shape	Position	Rotation	Approach Angle	Gripper Rotation	Gripper Width
cylinder	vertical	120	face front	0	0.3
cylinder	horizontal	10	top down	0	0.3
cuboid	horizontal	30	face front	30	0.5
cuboid	vertical	30	face front	30	0.3
cylinder	horizontal	100	top down	100	0.3

Figure 3: Example of the table in the grasping task

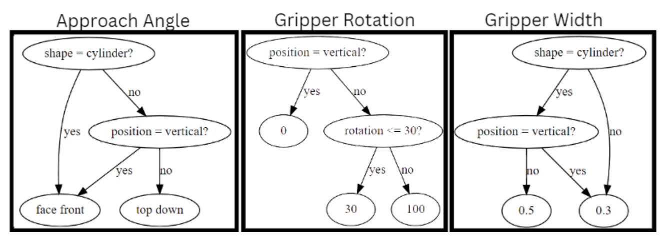

Figure 4: Example of the decision trees in the grasping task

4.5 Experimental procedure

The participants were split into two groups (Figure 6). The first group performed the tasks in the non-competitive mode first, then in the competitive mode. The second group performed tasks in the competitive mode first before the non-competitive mode.

The grasping task was done before the sorting task for both groups as we were not concerned with comparing between grasping and sorting.

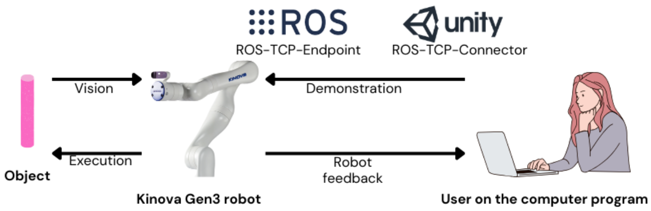

Figure 5: System structure and workflow

Before the experiment, the participants were taught how to read the decision trees and the table such as what the arrows meant and what the different columns of the table were.

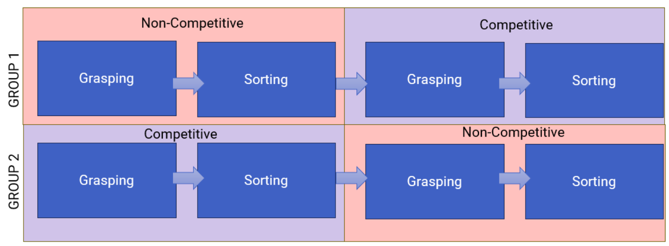

Figure 6: Experimental procedure

The post task experiment involves utilising the logged demonstrations to reproduce the robot performance at each number of demonstrations.

4.6 Robot interpretability assessment

In the grasping task, after four and eight demonstrations, we asked the user to estimate the robot's ability in sampled subtasks – as proposed by [Phaijit et al., 2023] for the robot interpretability assessment. For each object, we sampled eight object orientations such as those shown in Figure 7. In the sorting task, we performed interpretability tests after six and twelve demonstrations. In each test, we asked the user to estimate whether the robot would successfully sort each of the 12 unique objects into the correct bowl.

4.7 Teaching efficiency assessment

To assess the teaching efficiency, we evaluated the number of demonstrations needed to teach the robot the tasks, as suggested by [Sena and Howard, 2020], by collecting the given demonstrations during the experiment. If the robot had not fully learnt by the second round of interpretability assessments, we asked the participant to keep teaching the robots until it had fully learnt the task. After the experiment, we conducted post-task analysis to assess the number of demonstrations required to perform the tasks successfully.

4.8 Participants

Twenty participants (11 M, 9 F) were recruited. The age range was between 22 and 35 (M=27.90, SD=3.84). All participants were screened to self-identify as robotics and programming novices. The participants did not receive any monetary reimbursement. No identifiable data was collected. The research was approved by the University's Human Research Ethics Team (Project Reference Number: iRECS5790).

4.9 Questionnaire

The participants answered the questionnaire before and after the task. In the pre-task questionnaire, we asked demographic questions about age and gender.

For the post-task questionnaire, the participants were asked to answer the questions using a 5-point Likert scale. They were also asked to provide reasoning behind their answers.

- **Q4.** Which mode did you prefer – the competitive mode or the non-competitive mode? Why?
- **Q5.** Which mode made you feel the robot was more competent – the competitive mode or the non-competitive mode? Why?
- **Q6.** Which intervention was easier to understand – the decision trees or the table? Why?
- **Q7.** Which intervention was more effective in helping you determine which demonstrations to teach the robot – the decision trees or the table? Why?
- **Q8.** Which intervention did you prefer to use – the decision trees or the table? Why?

The questionnaire was designed to be comparative to reduce the cognitive load for the participants since they are being asked to assess many variables such as the mode and the user interface. It was also to encourage the participants' qualitative feedback to directly compare between the options instead of rating them individually. For instance, it is irrelevant how much they enjoyed the competitive mode or the non-competitive mode individually, on a Likert scale. This work's scope focuses on whether, given the choice, which option would they would select.

5 Results & Discussion

5.1 Competitive vs Non-Competitive Modes

The Shapiro test confirmed that the competitive vs non-competitive performance data did not follow normal distribution. Hence, the Wilcoxon signed-rank test was conducted to compare the experimental results between the competitive and non-competitive modes.

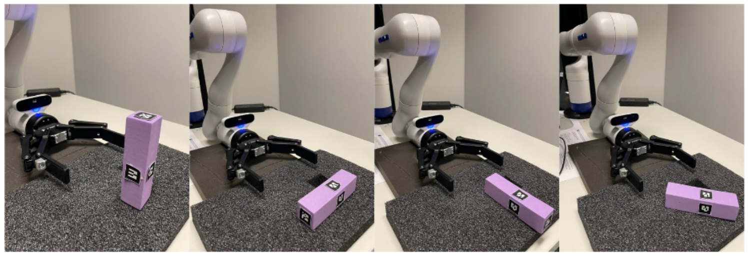

Figure 7: Examples of the sampled object orientations during the robot interpretability assessment

Number of demonstrations required
In the grasping task, there was a significant difference in the number of demonstrations required for the non-competitive mode (M=7.4, SD=1.82) and the competitive mode (M=5.7, SD=1.03); $p = .004$.

Likewise, in the sorting task, there was a significant difference in the number of demonstrations required for the non-competitive mode (M=11.4, SD=2.52) and the competitive mode (M=9.4, SD=1.85); $p = .02$.

Overall, the competitive mode statistically significantly reduced the number of demonstrations required, indicating that it encouraged higher teaching efficiency than the non-competitive mode.

Robot interpretability
In the grasping task, there was no significant difference in the robot interpretability, or the number of correct interpretations, in the first round of interpretation assessment between the non-competitive mode (M=5.65, SD=1.50) and the competitive mode (M=4.75, SD=1.41); $p = .11$. There was also no significant difference for the second round of interpretation assessment between the non-competitive (M=5.65, SD=1.50) and the competitive modes (M=5.65, SD=1.50); $p = .54$.

Similarly, for the sorting task, there was also no significant difference for both the first round of assessment ($p = .87$) and the second round ($p = .64$).

Overall, there was no statistically significant difference in the robot interpretability between the competitive and non-competitive modes.

User Perspective
Preference As can be observed in Figure 8a, the overwhelming majority of the participants preferred the competitive mode to the non-competitive mode. The answers ranged from strongly preferring the competitive mode to not having a preference. The median was moderately preferring the competitive mode. The general consensus was that the participants found that the competitive mode was 1) more transparent – they felt it provided more insights into the robot; "I like knowing what it can and cannot do" and 2) more entertaining/fun as it was like a game; "It makes the process more fun" and "It's more interesting as there are more interactions from the robot...". This was expected as gamification tends to keep the user engaged in the task. Some of the participants also noted that playing a game with a robot while teaching it is new to them. Although gamification is commonly applied to learning processes, it is rarely used for the enjoyment of the teacher.

These results have therefore highlighted how using gamification or competitive interactions may help to make the teaching process more engaging for users and make the users feel there is more transparency from the robot.

Perceived competency The answers ranged from strongly feeling that the robot was more competent in the competitive mode, to moderately feeling the robot was more competent in the non-competitive mode. The median and the most frequent answers were moderately feeling that the robot was more competent in the competitive mode. The participants stated that because the robot was aware of its own capabilities, the participants felt the robot was smarter e.g. "...it knows whether it can succeed, so it seems smarter.". On the other hand, there was a minority of people who felt that knowing the robot did not know how to perform certain tasks, in fact, made them feel that the robot was not very competent. One participant in particular noted that not knowing the robot's ability made them feel like the robot was being secretive and hence more cunning.

Overall, however, participants tended to perceive the robot as more competent in the competitive mode, which could foster greater trust in users. This increased perception of competence may also encourage novice users to engage more frequently with the robot. These findings offer valuable preliminary insights into the potential advantages of competitive human-robot interactions in robot Learning from Demonstration.

5.2 User Interface Interventions

As for the UI interventions, the survey results on the participants' opinions on which of the interventions was easier to understand, which was more effective for determining demonstrations to teach, and which was preferred can be seen in Figure 8b.

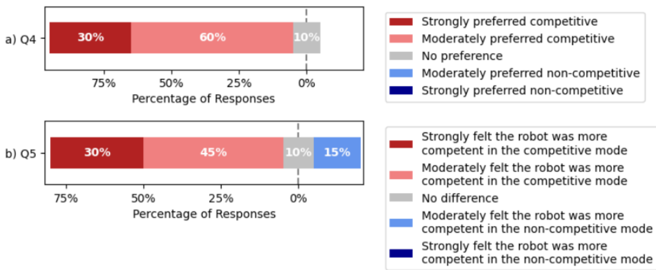

(a) Questionnaire responses on competitive vs non-competitive modes regarding i) which mode participants preferred, and ii) which mode the participants felt the robot was more competent in

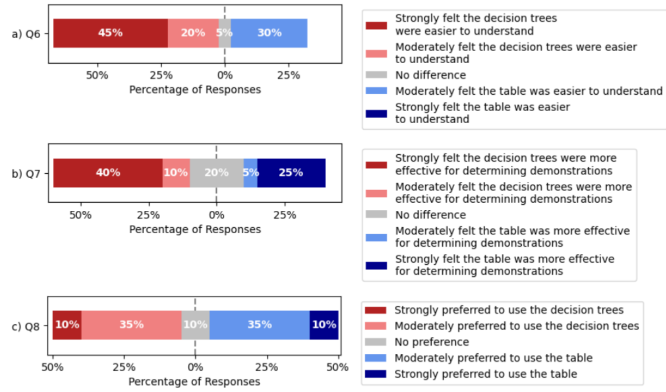

(b) Questionnaire responses on user interface interventions regarding i) which intervention was easier to understand, ii) which intervention was more effective for determining demonstrations, and iii) which intervention was preferred.

Figure 8: Questionnaire results

Ease of understanding In terms of ease of understanding, the most frequent response was *strongly feeling that the decision trees were easier to understand*, with a frequency of nine. The median response was *moderately feeling the decision trees were easier to understand.* Participants generally favoured the decision trees due to their simplicity, as they only required following a single path and making binary choices ("yes" or "no"), whereas the tabular format involved processing multiple columns and rows of information. Some participants indicated that, after receiving an explanation of how decision trees functioned, they found them straightforward to interpret. However, without manuals or instructions, they felt they would have preferred the table due to their familiarity with tabular data. One participant noted that if they were to use the system independently at home, a comprehensive manual would be essential.

Interestingly, the smaller group of participants who found the tabular format easier to understand remarked that they found decision trees harder to follow and preferred the clarity and familiarity of tables.

From these findings, it can be concluded that clear and thorough instruction is essential when introducing decision trees to novice users as a method for representing the robot's learning state.

Effectiveness for determining demonstrations to teach next For the question of which UI intervention was more effective for determining the demonstrations to teach the robot, the median response falls between *moderately feeling the decision trees were more effective* and feeling that there was *no difference.* The most frequent response was *strongly feeling the decision trees were more effective*, with a frequency of eight. Compared to the previous question of ease of understand which favoured the decision trees, this one caused a slight tendency towards selecting the table. Although the decision trees were popular for their simplicity, some of the participants found that they tended to scan the table to find what they had not yet taught the robot. The trees were easy to follow, but ultimately did not contain as much information as the table. Some of the participants preferred to have all the information laid out in front of them compared to only seeing partial information that the decision trees were providing them.

Nevertheless, the most frequent answer was still *strongly feeling the decision trees were more effective* as they were easier to follow, especially when there had been a lot of demonstrations, the table would quickly become congested.

Preference of User Interface Intervention In the final question regarding the UI interventions, participants were asked which intervention they would prefer to use if they could only select one. Responses were evenly split between the two options, with the median response indicating *no preference.* The most common answers were *moderately preferring the decision trees* and *moderately preferring the table*, each with a frequency of 7 out of 20 responses. Despite a tendency towards favouring decision trees in the previous questions, this was not reflected here. The general sentiment was that while decision trees were easier to process, participants were more familiar with tabular data and found the table more comprehensive. As such, if forced to choose, some favoured the familiarity and completeness of the table, while others preferred the simplicity of decision trees.

Though this question was not originally part of the questionnaire, participants were also asked if they would prefer using only one intervention or both. Remarkably, all twenty participants expressed a preference for using *both*. While managing cognitive load is important, the consensus was that having both decision trees and the demonstration table displayed simultaneously was preferable. Participants noted that each UI intervention conveyed distinct information, and having access to both would be beneficial. Although the simplicity of the tasks in this study may have influenced these results, they remain relevant for robot training systems designed with smaller numbers of demonstrations for novice users. Determining at what point the information might overwhelm novice users is beyond the scope of this work. However, in most Learning from Demonstration tasks, such as trajectory learning, fewer than ten demonstrations are typically required [Paraschos et al., 2013; Khansari-Zadeh and Billard, 2011; Ravichandar et al., 2019].

6 Conclusion

In this work, we demonstrated that competitive human-robot interactions can enhance the teaching process by improving the teaching efficiency, the robot's perceived competency and user enjoyment. The competitive elements likely encouraged users to be more deliberate in their demonstrations, resulting in greater teaching efficiency, which in turn contributed to a more positive teaching experience. Additionally, the increased perceived competency of the robot could foster greater user trust, potentially encouraging continued use of robot systems. Overall, our findings underscore the value of incorporating competitive elements into human-in-the-loop robot learning systems.

We also explored the use of decision trees and tables as visual representations of the robot's learning process. While decision trees were generally preferred for their clarity, users expressed a desire to have access to tabular data as well. Future research should investigate whether these preferences persist in more complex teaching tasks, where the increased cognitive load might lead users to favour having only one form of representation rather than both. Moreover, user preferences were found to be influenced by their familiarity with the visualisations and the clarity of the instructions provided, highlighting the importance of clear guidance on interpreting them.

In this experiment, arbitrary thresholds were employed in the competitive mode. The choice of the threshold value plays a critical role in the user's demonstration selection. A threshold that is set optimally, making the competitive mode appropriately challenging, may motivate users to invest effort in selecting diverse demonstrations. However, if the threshold is set too high, with the competition perceived as overly difficult, it could lead to user frustration, potentially deterring continued use of the system. Conversely, a threshold set too low may result in a lack of challenge, failing to encourage users to provide more meaningful or varied demonstrations. Future work should explore the impact of different threshold values on both teaching efficiency and user experience. A more deliberate threshold may also help to improve robot interpretability.

It is worth noting that the positive results observed in this study occurred in relatively short, simple tasks. Longer or more tedious tasks may yield different outcomes, and this remains an area for future exploration. Additionally, this study did not account for participants' backgrounds, which may have influenced their preferences for decision trees or tabular data. Gathering such information in future studies could provide deeper insights into how user characteristics impact interaction preferences.

Acknowledgments

This work was partially supported by the Australian Research Council (Grant No. DE210100858).

References

[Amershi et al., 2014] Amershi, S., Cakmak, M., Knox, W. B., and Kulesza, T. (2014). Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine*, 35(4):105–120.

[Bilal et al., 2024] Bilal, M., Lipovetzky, N., Oetomo, D., and Johal, W. (2024). Beyond success: Quantifying demonstration quality in learning from demonstration. In *2024 IEEE/RSJ International conference on Intelligent Robots and Systems (IROS)*. IEEE.

[Borisov et al., 2022] Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., and Kasneci, G. (2022). Deep neural networks and tabular data: A survey. *IEEE transactions on neural networks and learning systems*.

[Bryant et al., 2020] Bryant, D., Borenstein, J., and Howard, A. (2020). Why Should We Gender?: The Effect of Robot Gendering and Occupational Stereotypes on Human Trust and Perceived Competency. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pages 13–21, Cambridge United Kingdom. ACM.

[Cakmak et al., 2010] Cakmak, M., Chao, C., and Thomaz, A. L. (2010). Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development*, 2(2):108–118.

[Cakmak and Thomaz, 2012] Cakmak, M. and Thomaz, A. L. (2012). Designing robot learners that ask good questions. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 17–24. ISSN: 2167-2148.

[Cansev et al., 2021] Cansev, M. E., Xue, H., Rottmann, N., Bliek, A., Miller, L. E., Rueckert, E., and Beckerle, P. (2021). Interactive Human–Robot Skill Transfer: A Review of Learning Methods and User Experience. *Advanced Intelligent Systems*, 3(7):2000247.

[Caponetto et al., 2014] Caponetto, I., Earp, J., and Ott, M. (2014). Gamification and education: A literature review. In *European conference on games based learning*, volume 1, page 50. Academic Conferences International Limited.

[Chao et al., 2010] Chao, C., Cakmak, M., and Thomaz, A. L. (2010). Transparent active learning for robots. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 317–324. IEEE.

[De Graaf et al., 2017] De Graaf, M., Ben Allouch, S., and Van Dijk, J. (2017). Why do they refuse to use my robot? reasons for non-use derived from a long-term home study. In *Proceedings of the 2017 ACM/IEEE international conference on human-robot interaction*, pages 224–233.

[Diehl et al., 2020] Diehl, M., Plopski, A., Kato, H., and Ramirez-Amaro, K. (2020). Augmented Reality interface to verify Robot Learning. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 378–383. ISSN: 1944-9437.

[Duan et al., 2020] Duan, J., Wang, Q., Pinto, L., Kuo, C.-C. J., and Nikolaidis, S. (2020). Robot Learning via Human Adversarial Games. *arXiv:1903.00636 [cs]*. arXiv: 1903.00636.

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

[Gockley et al., 2005] Gockley, R., Bruce, A., Forlizzi, J., Michalowski, M., Mundell, A., Rosenthal, S., Sellner, B., Simmons, R., Snipes, K., Schultz, A. C., et al. (2005). Designing robots for long-term social interaction. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1338–1343. IEEE.

[Gorishniy et al., 2021] Gorishniy, Y., Rubachev, I., Khrulkov, V., and Babenko, A. (2021). Revisiting deep learning models for tabular data. *Advances in Neural Information Processing Systems*, 34:18932–18943.

[Habibian et al., 2022] Habibian, S., Jonnavittula, A., and Losey, D. P. (2022). Here's what i've learned: Asking questions that reveal reward learning. *ACM Transactions on Human-Robot Interaction*.

[Hamaya et al., 2021] Hamaya, M., Tanaka, K., Shibata, Y., von Drigalski, F., Nakashima, C., and Ijiri, Y. (2021). Robotic Learning From Advisory and Adversarial Interactions Using a Soft Wrist. *IEEE Robotics and Automation Letters*, 6(2):3878–3885. Conference Name: IEEE Robotics and Automation Letters.

[Hancock et al., 2011] Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., and Parasuraman, R. (2011). A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction. *Human Factors*, 53(5):517–527. Publisher: SAGE Publications Inc.

[Khansari-Zadeh and Billard, 2011] Khansari-Zadeh, S. M. and Billard, A. (2011). Learning stable nonlinear dynamical systems with gaussian mixture models. *IEEE Transactions on Robotics*, 27(5):943–957.

[Lewandowsky et al., 2000] Lewandowsky, S., Mundy, M., and Tan, G. (2000). The dynamics of trust: comparing humans to automation. *Journal of Experimental Psychology: Applied*, 6(2):104.

[Lewis et al., 2018] Lewis, M., Sycara, K., and Walker, P. (2018). The role of trust in human-robot interaction. *Foundations of trusted autonomy*, pages 135–159.

[Liu et al., 2018] Liu, H., Zhang, Y., Si, W., Xie, X., Zhu, Y., and Zhu, S.-C. (2018). Interactive Robot Knowledge Patching Using Augmented Reality. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1947–1954. ISSN: 2577-087X.

[Liu et al., 2021] Liu, H., Zhu, Y., and Zhu, S.-C. (2021). Patching interpretable and-or-graph knowledge representation using augmented reality. *Applied AI Letters*, 2(4):e43.

[Muir and Moray, 1996] Muir, B. M. and Moray, N. (1996). Trust in automation. part ii. experimental studies of trust and human intervention in a process control simulation. *Ergonomics*, 39(3):429–460.

[Paraschos et al., 2013] Paraschos, A., Daniel, C., Peters, J. R., and Neumann, G. (2013). Probabilistic movement primitives. *Advances in neural information processing systems*, 26.

[Pekrun et al., 2010] Pekrun, R., Goetz, T., Daniels, L. M., Stupnisky, R. H., and Perry, R. P. (2010). Boredom in achievement settings: Exploring control–value antecedents and performance outcomes of a neglected emotion. *Journal of educational psychology*, 102(3):531.

[Phaijit et al., 2022a] Phaijit, O., Obaid, M., Sammut, C., and Johal, W. (2022a). A taxonomy of functional augmented reality for human-robot interaction. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 294–303. IEEE.

[Phaijit et al., 2022b] Phaijit, O., Sammut, C., and Johal, W. (2022b). 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*, pages 86–94.

[Phaijit et al., 2023] Phaijit, O., Sammut, C., and Johal, W. (2023). User interface interventions for improving robot learning from demonstration. In *Proceedings of the 11th International Conference on Human-Agent Interaction*, pages 152–161.

[Ravichandar et al., 2019] Ravichandar, H., Ahmadzadeh, S. R., Rana, M. A., and Chernova, S. (2019). Skill acquisition via automated multi-coordinate cost balancing. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7776–7782. IEEE.

[Robinson et al., 2019] Robinson, N. L., Turkay, S., Cooper, L. A., and Johnson, D. (2019). Social robots with gamification principles to increase long-term user interaction. In *Proceedings of the 31st Australian Conference on Human-Computer-Interaction*, pages 359–363.

[Sahakyan et al., 2021] Sahakyan, M., Aung, Z., and Rahwan, T. (2021). Explainable artificial intelligence for tabular data: A survey. *IEEE access*, 9:135392–135422.

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

[Shwartz-Ziv and Armon, 2022] Shwartz-Ziv, R. and Armon, A. (2022). Tabular data: Deep learning is not all you need. *Information Fusion*, 81:84–90.

[Wang and Belardinelli, 2022] Wang, C. and Belardinelli, A. (2022). Investigating explainable human-robot interaction with augmented reality.

[Yoon and Nikolaidis, 2020] Yoon, S. H. and Nikolaidis, S. (2020). Robot Learning in Mixed Adversarial and Collaborative Settings. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9329–9336. ISSN: 2153-0866.