Train your robot in AR: investigating user experience with a continuously learning robot

Anna Belardinelli
anna.belardinelli@honda-ri.de
Honda Research Institute Europe
Offenbach, Germany

Chao Wang
chao.wang@honda-ri.de
Honda Research Institute Europe
Offenbach, Germany

Daniel Tanneberg
daniel.tanneberg@honda-ri.de
Honda Research Institute Europe
Offenbach, Germany

Stephan Hasler
stephan.hasler@honda-ri.de
Honda Research Institute Europe
Offenbach, Germany

Michael Gienger
michael.gienger@honda-ri.de
Honda Research Institute Europe
Offenbach, Germany

Figure 1: The AR scenario for the robot training and interaction with the user. The scene is a common office space with a table upon which holographic kitchen items are displayed. Virtual controls (to log out, reinitialize the scene, and initiate a teaching episode) were displayed on the upper visual field of the person wearing AR glasses and moving rigidly with the user's head.

ABSTRACT

Assistive robots that can be deployed in our homes will need to be understandable, operable, and teachable by non-expert users. This calls for an intuitive Human-Robot Interaction approach that is also safe and sustainable in the long term. Still, few studies have looked at repeated, unscripted interactions in a loosely supervised setting with a robot incrementally learning from the user and consequentially expanding its knowledge and abilities. In this study, we set out to test how the user's experience and mental model of the robot evolve when spontaneously teaching it simple tasks in Augmented Reality (AR). Participants could freely access the AR glasses in a common office space and demonstrate physical skills in a virtual kitchen scene, while the holographic robot gave feedback about its understanding and could ask questions to generalize the acquired task knowledge. The robot learned the semantic effects of the demonstrated actions and upon request could reproduce those on observed or novel objects through generalization. Preliminary results show that users find the system engaging, understandable, and trustworthy, but with large variance on the last two constructs. Further analyses will assess how subjective measures can be correlated to user behavior, to evaluate the relation between system understanding and teaching effectiveness.

CCS CONCEPTS

• **Computer systems organization → Robotics**; *Real-time operating systems*; • **Human-centered computing → User studies**.

KEYWORDS

Long-term human-robot interaction, continual learning, learning from demonstration, augmented reality

LEAP-HRI, May 11th 2024, Boulder, Colorado

1 INTRODUCTION

As robots become commonplace also outside laboratories and factories, there is an increasing need for assessing how novice users can intuitively personalize the robot's skills without any technical knowledge. This means that robots supposed to support us in our daily activities, possibly in our private spaces, need to be teachable in an easy way while providing the user with some insight into the robot's perception, reasoning, and action capabilities. Our inner mental model of a novel acquaintance gets refined the more we interact with them. The same happens with complex technical systems: even if instructing material can inform the user on how to operate a machine and what to expect, the more the user uses the system the more proficient they become at understanding the way it processes information and behaves. Nevertheless, because of common embodiment and cognitive structure, we rely on many assumptions about what others see and understand about the world. Yet different robots have different embodiments (sensors, actuators) and might rely on different cognitive architectures (scene understanding, manipulation capabilities,...), all of which might not be directly apparent on a first encounter [18, 19]. Typically, long-term studies have been conducted in the field of social robotics [9], considering acceptance, adaptation, and personalization aspects [1, 4, 8]. Often non-physical tasks are considered, not involving continual learning, or focusing on the robot social cues and appearance [11]. Here, we investigated how users would interact with a system able to learn physical tasks from demonstration and to communicate and show its progress, prioritizing skill learning and explainability over social interaction. We considered that important factors for user experience in the interaction with such a system would be how engaging the whole teaching process is, how understandable would be for the user the way the system processes the collected data and learns, and how trustworthy the learning system appears to be [2, 7]. Moreover, since both the user's mental model of the robot and the robot's capabilities change across multiple interactions, we are interested in characterizing such interdependent temporal evolutions. Our contributions here include the presentation of the system, the setup and procedure for the user study, along with some preliminary results and outlook discussion.

2 SYSTEM DESIGN

We devised a system relying on the integration of multiple modules for semantic learning from demonstration, skill generalization, symbolic planning, and action generation. These communicate via ROS with the AR framework [1].

AR functionalities. The shared environment (Fig. 1), accessible through AR glasses, includes a kitchen table with multiple appliances and food items, while the robot is displayed on the other side as observing the scene and the user. Object poses are continuously tracked and sent to the back-end system, along with the user's head pose and their actions (picking/dropping objects, turning on/off appliances, pouring beverages). Through gesture recognition and grasping/manipulation logic, users interact with virtual objects similarly as with real ones. To increase explainability and give the tutor online feedback about the robot's situation understanding, several virtual design elements [20] are displayed in AR: object and action labels (XAI cues) are popping up whenever the teacher gazes at some object or a manual action is recognized, while action labels are also spoken out verbally (see Fig. 2 top, see also [3, 21, 22]).

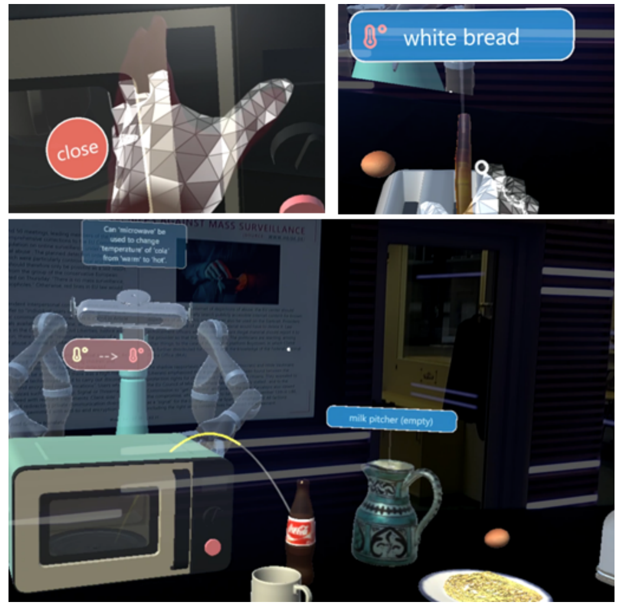

Figure 2: Top: Design elements for online feedback to the tutor about recognized actions and objects. Bottom: after the demonstration, the robot asks questions to generalize the demonstrated skill (here after seeing heating milk in the microwave it asks the tutor *"Can microwave be used to change the temperature of cola from warm to hot?"*).

Learning. For the robot to learn new skills, a two-stage learning concept is used. First, the user demonstrates a skill to the robot, which acquires it using semantic skill learning concepts. The symbolic representation of the skill encodes preconditions, actions, and effects (see [17] for more details). In this study, there are three types of effects the robot can observe as a consequence of the user's actions: the change of temperature (e.g., by putting some food in the fridge or the oven), the change of crispiness (by the toaster), and the change of aroma (by putting the mint tea bag in any liquid)[2]. This is communicated back to the teacher by uttering a sentence in the form `"I've learned that <appliance> can change the <attribute> of <food item> from <predicate 1> to <predicate 2>"`, with $appliance \in \{oven, refrigerator, microwave, toaster, teabag\}$, while $attribute \in \{temperature, crispness, aroma\}$, and $fooditem \in \{bread\ slice, bagel, pizza\ slice, apple, egg, spaghetti, coke, milk, water\}$. As to the predicates of the attributes, temperature can be one of $\{cold, warm, hot\}$, crispiness in $\{soft, crispy\}$, while aroma can vary in $\{no\ aroma, mint\}$. In the second stage, the robot asks the user questions about the demonstrated skill, to generalize to similar objects. For example, after seeing a demonstration on how to use the microwave to heat milk, the robot can

[1] In the present study, all physical objects and the robot itself are presented as holograms for logistic simplicity, still, they all are virtual replicas of real counterparts that can be used as anchors for Augmented Reality visualizations.

[2] We assume the robot relies on hypothetical sensors to detect such changes. In practice, such events are simulated in the virtual scene by changing related elementary attributes of the food when turning on an appliance or if the tea bag is in a liquid, respectively.

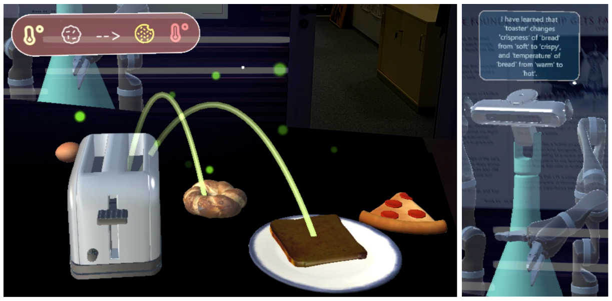

Figure 3: Graphical elements used to explain to the user that the robot learned that food items in the "bread" category can be heated up in the toaster (left). The same message is also provided by text (right) and speech.

ask whether a similar food (see Fig. 2 bottom) can be heated in the microwave or in some other appliance it knows has the same effect (e.g., the oven). Across demonstrations and questions, the system generalizes to categories higher up in the WordNet hierarchy [5] (e.g., learning that the toaster can make "bread" hot, which is the category encompassing white bread and bagel, see Fig. 3).

Planning. The use of elementary, domain-agnostic attributes helps generalizing skills, allowing the robot to plan to reproduce the observed tasks with the same or similar food items. Beside the elementary attributes introduced above, the system makes use of spatial attributes for $content \in \{inside, outside\}$ and $distance \in \{close, distant\}$, and of device-related attributes for $openness \in \{open, close\}$ and $activity \in \{on, off\}$. In this way, the system can represent also the effects of manipulation actions modifying spatial and device attributes, such as *put in, get out, approach with, pour, open, close, switch on, switch off.* The actions that change spatial predicates are mostly independent of elementary predicates, and the actions that change device-specific predicates are fully independent of spatial and elementary ones. Such multi-level representation relying on different attribute spaces allows a symbolic planner to iteratively plan according to different views, each containing ever more composite sets of attributes, first considering just which objects are present and the predicates of the elementary attributes, up to including spatial relations and device-related predicates.

3 STUDY DESIGN

The study was designed to be conducted semi-"in-the-wild" in an office setting, with participants taking part spontaneously, without experimenter supervision (if they wished so). This was meant to allow users to interact with the system as long and as often as they wanted, reinforcing the perception of a 24/7 functioning system. The study was configured as a loose repeated interaction study, with users asked to engage in multiple teaching and planning sessions to monitor the evolution of their behavior as the system increasingly learns and of their subjective experience.

3.1 Procedure and methods

The study took place in a lobby connecting different office spaces and featuring a long table upon which the virtual scene was calibrated. Participants were invited via mail to the study, explaining the purpose of the study and providing a link to a video tutorial of the experimental procedure, so that participants knew what to expect. Participants were also reminded by a virtual board in the AR of the basic functions of the system and that they could ask the experimenters in a nearby room, in case they needed any assistance (Fig. 1 on the right). The informed consent for the study was provided on the table, along with a paper form asking for demographic data (gender and age) and the Affinity for Technology Interaction (ATI) scale [6]. These paper forms were inserted in sealed paper boxes, before the very first interaction. On the table, participants could find the HoloLens and a laptop to fill out the scale questionnaires after each session. Their answers to the paper and online questionnaires and their interactions with the systems were tracked via a pseudo-ID. After logging into the HoloLens, users were presented with the virtual scene in Fig. 1 and could play around with objects to get acquainted with manipulating virtual objects. The robot would ask the user to teach it something or give it a task to execute. Each robot utterance was also displayed as a speech bubble on top of the robot's head until a new utterance was spoken out, helping participants remember what the robot said last. Participants were told that their task was to train a personal robot, by demonstrating the use of kitchen appliances to prepare some food, answering related questions, and, to assess the learning progress by asking the robot to make something "hot/cold/crispy" or to change the aroma to "mint". A teaching episode was started by pressing the "Start learning" button in the user's upper field of view. Just one skill (e.g., toast bread in the toaster) or multiple skills (e.g., heat water in the microwave and then make tea by putting the teabag in the cup) could be demonstrated. As soon as the user pressed the "Stop learning" button, the robot would give a summary of the (last) learned skill (see Fig.3). Subsequently, a loop would start with the robot asking permission to ask a question and upon affirmative answer, asking the user a yes/no question (see Fig.2, bottom). A planning episode could be triggered by the user looking at food item and uttering their request. If the robot could find a plan to achieve the desired effect, it would state "This is my plan". Then a transparent avatar of the robot would appear on the user's side of the table (Fig. 4) and demonstrate the plan while announcing each step. The plan could reflect the user's demonstration or deviate from it, in case the robot had generalized (e.g., it could use the microwave instead of the toaster). If no plan was found, this was communicated to the user. The system behavior is handled by a state machine with three states: a default playground state to play with the scene, renew it, or log out; a teaching state triggered by the learning button press; a planning state, triggered by a plan request.

To nudge the user into teaching sensible skills and answering questions diligently, we introduced a teaching score, shown as a number and a progress bar above the robot in the playground state. When logging in for the first time, a zero score is displayed which is updated after each teaching episode. The score is computed as an F2 score on a predefined ground truth, using correct generalizations as true positives – covering the correct objects with the generalization –, wrong generalizations as false positives – covering incorrect objects –, and missing generalization as false negatives – not covering objects that should be covered. This gamification element aims also at giving participants some measure of how their teaching expands the robot's knowledge and enticing them to keep

Figure 4: The robot avatar executing a requested plan.

training the system, also in competition with other users [12, 13]. It was specified in fact that the pseudo-IDs of the best robot trainers would be ranked and made known at the end of the study.

After logging out of the system, participants were finally prompted to fill out the online questionnaires on the laptop. The questionnaires are established and validated scales, aimed at assessing the user experience through the proxy constructs of engagement, understanding, and perceived trust: the User Engagement Scale (UES, 9 items on a 5 points Likert-scale) [10], the Subjective Information Processing Awareness Scale (SIPAS, 6 items, 6 points Likert-scale, α = .88) [15], and the Trust Perception Scale - HRI (TP-HRI, 14 items, 0−100 % rating scale) [14]. We reasoned indeed that in the context of long-term HRI such constructs would play a large role, besides the actual utility and efficacy of the robot. The user study went on for about 13 weeks (with season holiday interruption). The whole framework was active and available during working hours 5 days/week. In the end, the data collected amounted to the subjective measures (the questionnaires) and behavioral data collected through the system and stored in a database. These latter comprise for each user: the trained model, the latest score achieved, timestamped teaching episodes in terms of learned skills, questions and answers, and timestamped planning episodes in terms of user request and related plan or negative answer.

3.2 Participants

Participants were recruited among associates and PhD students at our institute. In total, 37 people took part in the study, of which 3 did not fill out the demographics and ATI questionnaires and 1 never filled out the online questionnaires. Thus, we analyzed the data from 33 participants (26 male, 7 female; age between 26 and 59, M= 36.82). The median ATI score was 4.44 (M= 4,40; SD=0.84; α = .88). Thus, participants were non-experts as to AR or the specific functioning of the system, but tech-savvy and often with an engineering background.

Table 1: Descriptive statistics for three used scales (after the first interaction). For each scale the number of Likert-scale points or rating range is reported in parentheses.

SCALE	MEDIAN	MODE	RANGE
UES (5 points)	3.67	3.67	1.77
SIPAS (6 points)	4.0	4.0	4.0
TP-HRI (0-1)	0.69	0.73	0.66

4 PRELIMINARY RESULTS

We are examining the collected data, but we have inspected only questionnaire answers thus far. These revealed that 33 people interacted with the system at least once, 16 interacted twice, and only 2 people used the system in three or more sessions. When considering only the first interaction, the descriptive statistics regarding the UES, SIPAS, and TP-HRI are presented in Tab. 1 as recommended in [16]. A one-tailed one-sample t-test confirmed that the mean value was significantly above the "neutral" point for UES (coded as 3, $t(32) = 9.06, p < .001$) and the border between "slightly disagree" and "slightly agree" for SIPAS (coded as 3.5, $t(32) = 3.12, p = .004$). Similarly, the percent of time users would generally trust the robot was higher than 50% for the TP-HRI scale ($t(32) = 6.72, p < .001$). No significant correlation was found with the ATI scores (all $p > .05$).

Looking at the data from users who interacted at two different time points, differences in UES, SIPAS, and TP-HRI scores between the first and the second interaction were all not significant (all $p > .05$), but all three scores slightly declined (UES: $M_{t2-t1} = -0.08, SD = 0.06$; SIPAS: $M_{t2-t1} = -0.34, SD = 1.30$; TP-HRI: $M_{t2-t1} = -0.06, SD = 0.20$).

5 DISCUSSION AND OUTLOOK

In this study, we put forward a robot training system, integrating continual learning, task planning, and explainable HRI in AR, and investigated how such framework is perceived by novel users invited to repeatedly interact.The closed-world and limited ontology used in the scenario allowed users to teach and test the system in short episodes. Preliminary results from scale data show that the robot is positively perceived across dimensions concerning engagement, trust, and understanding, with such perception staying stable across first and second interaction. Such data still offer rather limited insight into the varied experiences users had with the system: the amount of time spent in the system was not prescribed, hence a longer first interaction might give a better idea of the system than two short interactions. Future analyses thus are planned to correlate subjective measures to behavioral ones (e.g., number of demonstrated skills, answered questions, requested plans) and to the performance of the two partners (e.g., teaching score and number of successfully generated plans) through multiple regression, hopefully shedding light on which factors contribute to a better perception by the user and result in more effective teaching. While future robots will come with numerous pre-trained skills, teaching specific tasks will help users personalize robots to comply with their own needs and preferences, but this will need to be made possible in an easy and intuitive way for long-term interaction.

REFERENCES

[1] Muneeb Imtiaz Ahmad, Omar Mubin, and Joanne Orlando. 2017. Adaptive social robot for sustaining social engagement during long-term children–robot interaction. *International Journal of Human–Computer Interaction* 33, 12 (2017), 943–962.

[2] Beatrice Alenljung, Jessica Lindblom, Rebecca Andreasson, and Tom Ziemke. 2019. User experience in social human-robot interaction. In *Rapid automation: Concepts, methodologies, tools, and applications*. IGI Global, 1468–1490.

[3] Anna Belardinelli, Chao Wang, and Michael Gienger. 2023. Explainable Human-Robot Interaction for imitation learning in Augmented Reality. In *16th International Workshop on Human-Friendly Robotics (Springer Proceedings in Advanced Robotics)*. Springer, LNCS, LNAI, LNBI.

[4] Maartje MA De Graaf, Somaya Ben Allouch, and Jan AGM van Dijk. 2016. Long-term evaluation of a social robot in real homes. *Interaction studies* 17, 3 (2016), 462–491.

[5] Christiane Fellbaum. 2010. WordNet. In *Theory and applications of ontology: computer applications*. Springer, 231–243.

[6] Thomas Franke, Christiane Attig, and Daniel Wessel. 2019. A personal resource for technology interaction: development and validation of the affinity for technology interaction (ATI) scale. *International Journal of Human–Computer Interaction* 35, 6 (2019), 456–467.

[7] Soheil Habibian, Antonio Alvarez Valdivia, Laura H Blumenschein, and Dylan P Losey. 2023. A Review of Communicating Robot Learning during Human-Robot Interaction. *arXiv preprint arXiv:2312.00948* (2023).

[8] Bahar Irfan, Nathalia Céspedes Gomez, Jonathan Casas, Emmanuel Senft, Luisa F Gutiérrez, Monica Rincon-Roncancio, Marcela Munera, Tony Belpaeme, and Carlos A Cifuentes. 2020. Using a personalised socially assistive robot for cardiac rehabilitation: a long-term case study. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. Ieee, 124–130.

[9] Iolanda Leite, Carlos Martinho, and Ana Paiva. 2013. Social robots for long-term interaction: a survey. *International Journal of Social Robotics* 5 (2013), 291–308.

[10] Heather L O'Brien, Paul Cairns, and Mark Hall. 2018. A practical approach to measuring user engagement with the refined user engagement scale (UES) and new UES short form. *International Journal of Human-Computer Studies* 112 (2018), 28–39.

[11] Giulia Perugia, Maike Paetzel-Prüsmann, Madelene Alanenpää, and Ginevra Castellano. 2021. I can see it in your eyes: Gaze as an implicit cue of uncanniness and task performance in repeated interactions with robots. *Frontiers in Robotics and AI* 8 (2021), 645956.

[12] Nicole L Robinson, Selen Turkay, Leonie AN Cooper, and Daniel Johnson. 2019. Social robots with gamification principles to increase long-term user interaction. In *Proceedings of the 31st Australian Conference on Human-Computer-Interaction*. 359–363.

[13] Michael Sailer, Jan Ulrich Hense, Sarah Katharina Mayr, and Heinz Mandl. 2017. How gamification motivates: An experimental study of the effects of specific game design elements on psychological need satisfaction. *Computers in human behavior* 69 (2017), 371–380.

[14] Kristin E Schaefer. 2016. Measuring trust in human robot interactions: Development of the "trust perception scale-HRI". In *Robust intelligence and trust in autonomous systems*. Springer, 191–218.

[15] Tim Schrills, Mourad Zoubir, Mona Bickel, Susanne Kargl, and Thomas Franke. 2021. Are users in the loop? Development of the subjective information processing awareness scale to assess XAI. In *Proceedings of the ACM CHI Workshop on Operationalizing Human-Centered Perspectives in Explainable AI (HCXAI 2021), Upol Ehsan, Q. Vera Liao, Martina Mara, Mark Riedl, Andreas Riener, Marc Streit, Sandra Wachter, and Philipp Wintersberger (Eds.). The Internet.*

[16] Mariah Schrum, Muyleng Ghuy, Erin Hedlund-Botti, Manisha Natarajan, Michael Johnson, and Matthew Gombolay. 2023. Concerning trends in likert scale usage in human-robot interaction: Towards improving best practices. *ACM Transactions on Human-Robot Interaction* 12, 3 (2023), 1–32.

[17] Daniel Tanneberg and Michael Gienger. 2023. Learning type-generalized actions for symbolic planning. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2010–2017.

[18] Sam Thellman, Maartje de Graaf, and Tom Ziemke. 2022. Mental state attribution to robots: A systematic review of conceptions, methods, and findings. *ACM Transactions on Human-Robot Interaction (THRI)* 11, 4 (2022), 1–51.

[19] Sam Thellman and Tom Ziemke. 2021. The Perceptual Belief Problem: Why Explainability Is a Tough Challenge in Social Robotics. *ACM Transactions on Human-Robot Interaction* 10, 3 (2021), 1–15. https://doi.org/10.1145/3461781

[20] Michael Walker, Thao Phung, Tathagata Chakraborti, Tom Williams, and Daniel Szafir. 2022. Virtual, augmented, and mixed reality for human-robot interaction: A survey and virtual design element taxonomy. *arXiv preprint arXiv:2202.11249* (2022).

[21] Chao Wang and Anna Belardinelli. 2022. Investigating explainable human-robot interaction with augmented reality. In *5th International Workshop on Virtual, Augmented, and Mixed Reality for HRI (VAM-HRI)*.

[22] Chao Wang, Anna Belardinelli, Stephan Hasler, Theodoros Stouraitis, Daniel Tanneberg, and Michael Gienger. 2023. Explainable Human-Robot Training and Cooperation with Augmented Reality. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–5.