Grant agreement No: 101017008

Harmony
Assistive robots for healthcare

Enhancing Healthcare with Assistive Robotic Mobile Manipulation

(HARMONY) | H2020-ICT-2018-20| RIA

Start of the project: 01.01.2021
Duration: 42 months

Deliverable Number	D6.4
Deliverable Name	Human-robot social implications of learning from demonstration
WP Number	6
Lead Beneficiary	UT
Dissemination Level	Public
Internal Reviewer	UEDIN
Due Date	30.12.2023
Date of Submission	22.12.2023
Version	2.0

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101017008

Revision History

Version	Date	Author(s)	Comments
0.1	18-12-2023	Bob Schadenberg (UT)	First draft
0.2	19-12-2023	Mohamamdreza Kasaei, Zhibin(Alex) Li (UoE)	Reviewed
1.0	21-12-2023	Bob Schadenberg (UT), Mohamamdreza Kasaei, Zhibin(Alex) Li (UoE)	Final version
2.0		Camille Sallaberry	Updated version

Table of Contents

Summary

In this deliverable, we report on a study regarding the social implications of providing staff demonstrations. Specifically on the impact of the teacher of the robot on the worker who is to work in a cooperative task with the robot.

Within the Harmony use-case, having laboratory workers program the robot could be utilised to deal with changes in the shapes of the test samples and the boxes that contain them or to readjust the robot's effort to a different task within the laboratory that requires additional labour. Given that Learning from Demonstration (LfD) relies on demonstrations which are provided by a person, the paradigm is sensitive to who is teaching the robot. In human-robot cooperative or collaborative tasks, the role of the teacher in the execution of the task may be exacerbated, as that worker may not have been the one who taught the task to the robot. Because a continuously changing cobot may be problematic, we need to understand the effect who is teaching the robot has on the task execution by the workers, including the teachers, thereafter.

In the study reported in this deliverable, we specifically investigate the effect of who taught the robot on the task performance, trust, and fluency of the task. Furthermore, we look at the effects of a robot failing at a subtask after training and the attribution of that failure to who taught the robot. This will give us insight into how learning from demonstration may be employed effectively in the hospital laboratory.

At the time of writing this deliverable, the study still has to be executed due to the integration of the Harmony robot taking more time than anticipated. We will update this deliverable in early 2024 with the results of the study.

1. Introduction

The Harmony robot is a mobile manipulation platform that is designed for specific tasks within the bioassay sample flow. These test samples need to be unpacked, registered, sorted and pre-processed so that they are ready for the respective clinical test (Hawker, 2007). This is a task that is associated with causing work-related musculoskeletal disorders due to the repetitiveness of the task, as well as the potential for coming into contact with the, potentially hazardous, contents of the test samples (Schadenberg et al., 2023). For instance, when a lid is not closed correctly or there is a spill. While a robot could alleviate these issues, it is currently a difficult task to automate fully. Cooperation between worker and robot could also address the issues experienced by the laboratory workers, where the robot carries out the repetitive tasks and hands over the test samples for further inspection and sorting on priority by the worker. Within the Harmony use-case, having laboratory workers program the robot could be utilised to deal with changes in the shapes of the test samples and the boxes that contain them or to readjust the robot's effort to a different task within the laboratory that requires additional labour (Schadenberg et al., 2023).

To enable more flexible use of collaborative robots (cobots in short), such as the Harmony robots, they often feature programming interfaces that require little knowledge of programming. A worker with little programming knowledge could then program the robot without the intervention of the robot's supplier. For instance, to program the robot to perform a new (sub)task, or to adapt the cobot to the preferences of that specific worker. Workers then can delegate subtasks to the robot that they find physically or mentally stressful, whilst doing the subtasks that cannot be done by the robot, they find enjoyable, or that give meaning to their work.

There are a variety of methods for workers with limited knowledge of programming robots to program one. One such way is through the learning from demonstration (LfD[1]) paradigm, where a person provides demonstrations to the robot on the task from which the robot then derives a policy to execute the task itself (Argall et al., 2009; Billard et al., 2008; Schaal, 1999). LfD is particularly compelling for tasks that are difficult to script or frame as an optimisation problem but can be demonstrated (Ravichandar et al., 2020). The LfD techniques broadly fall within one of three categories (Ravichandar et al., 2020). Firstly, *kinesthetic teaching*, where the person physically moves the robot through the desired motions. To optimise the motion trajectories, the robot can then smoothen the motion or use other post-processing techniques. Secondly, the *teleoperation* of the robot through an external device (e.g., a joystick). And lastly, through passively *observing* a person do the task.

Given that LfD relies on demonstrations by a person, the paradigm is sensitive to who is teaching the person. One teacher may be better at teaching the robot than another, but

[1] Also referred to as apprenticeship learning, programming by demonstration, learning by demonstration, behavioural cloning, or imitation learning.

teachers may also encode their preferences in how the robot should execute the task. Elor et al. (2022) investigated how teaching the robot influences the teacher's trust and believability in the robot, compared to observing the robot learn by itself. The robot had to learn a reaching task, and the authors found that teaching a robot improves the teacher's trust in the robot and they considered the robot to have more humanlike motions compared to when observing the robot learn. When a robot that has been taught by a person fails at its task, Hedlund et al. (2021) found that those people then have a lower impression of both the robot and themselves, and trust themselves and the robot less as well as believe that others will trust them less.

In human-robot cooperative or collaborative tasks, the role of the teacher in the execution of the task may be exacerbated, as that worker may not have been the one who taught the task to the robot. The preferences of the teacher then can have a direct effect on how other workers will need to perform the task with the robot. This can result in physical mismatches where the robot's motions are relative to the teacher's bodily dimensions, the speed at which the task needs to be performed, or the order of subtasks. In particular in tasks that allow for the encoding of preferences (i.e., for which there is no single optimal way to perform the task).

A technological solution could be that each worker trains the robot on their own so that it knows their preferred way of working. Or that the robot learns the preferences (e.g., working speed of a worker) of individual workers throughout the interaction and adapts accordingly. However, from a practical point of view, changes in robot behaviour are problematic for risk assessment and are currently the largest barrier to the adoption of cobots (Kopp et al., 2021). While the cobot itself may still be safe to use, the process in which it's embedded may need to be reevaluated. For instance, while the robot may not severely harm a person directly (e.g., stop when it comes into contact with a person), it could drop an object that then hurts a person, interact with its surrounding environment to create a hazard or cause a failure in the process down the chain. A risk assessment is a time-consuming and complex matter that may in turn undermine the benefits of a LfD system.

Because a continuously changing cobot may be problematic, we need to understand the effect who is teaching the robot has on the task execution by the workers, including the teachers, thereafter. In the study reported in this deliverable, we specifically investigate the effect of who taught the robot on the task performance, trust, and fluency of the task. Furthermore, we look at the effects of a robot failing at a subtask after training and the attribution of that failure to who taught the robot. At the time of writing this deliverable, the study still has to be executed due to the integration of the Harmony robot taking more time than anticipated. We will update this deliverable in early 2024 with the results of the study. For now, this deliverable contains the method that we will use for the study, described in the next Section.

2. Materials and Methods

2.1 Participants

We plan to recruit 20 participants who work at hospital laboratories in the Netherlands for this study.

2.2 Materials

For this study, we use the Franka Emika Research 3 robotic arm, equipped with the Franka Emika Hand; a two-fingered gripper (see Figure 1). This robot enables kinesthetic teaching out of the box. Custom code to support the experiment was coded in C++.

Figure 1. The Franka Emika Research 3 arm with two-fingered gripper.

2.3 Experiment Design

The study is set up as a 2 by 2 between-participants design. The independent variables are the *teaching*, which will either be done by the participants themselves or by somebody else and the robot making a *mistake* whilst performing the task, which either does or does not happen. For every participant who will teach the robot how to perform the cooperative task, another participant will not teach the robot but instead work with the robot's behaviour taught by the other participant. Together, these two participants form a pair for whom the robot performs the task in the same manner. Each participant performs the task twice, each in a different session on a different day.

2.4 Experimental Setup

2.4.1 Kinesthetic Teaching Setup

Participants will need to perform a sorting task, resembling the sorting of test samples (see Figure 2). There are three boxes with trays inside them and an empty box. The goal of the cooperative task is that the robot hands the participant the boxes and that the participant then sorts the samples inside the trays (inside the boxes). The sorted trays then have to be placed inside the empty box.

The initial position of the boxes is fixed to ensure robust picking up of the boxes. The participants can teach the robot how and where to place the boxes. Each motion trajectory is stored in the order they are taught.

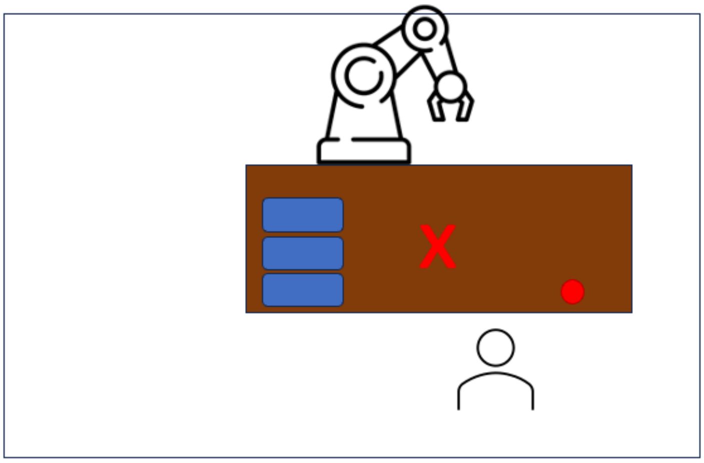

Figure 2. Experimental setup. The blue rectangles resemble the boxes and the red X is the location of where to drop them approximately. The red circle is the button.

2.4.2 Task Execution

For the two sessions that involve performing the task, participants will encounter the same setup as during teaching. By pressing a large button next to them. By doing so, the robot will execute the next motion trajectory in the queue and return to a neutral position afterwards. During the second time the participant carries out the task, the robot will make a mistake. This is done by playing a different motion trajectory than was programmed by the participant (but a trajectory with the same goal).

2.5 Evaluation Metrics

2.5.1 Task Performance

As a measure of performance, we measure the time it takes for participants to complete the task. The time starts when they press the button for the first time, until the last time it is pressed. From this number, we subtract the time it takes for the robot to complete its motions, to account for any motions that may take longer for certain participants. We expect that when somebody else taught the robot, the participant will take more time to complete the task (more hesitant).

2.5.2 Performance Trust

We measured the performance trust through the Multi-Dimensional Measure of Trust (MDMT) v2 (Ullman & Malle, 2023), where higher scores reflect higher performance trust in the robot. This was measured before training, after training, and after executing the task. The trust measure before training the robot is a baseline and reflects the expectations of the participant. After training, they will have experience with the robot and better understand whether the robot lived up to their initial trust. Then, once the participant finishes the task sessions, which may include a mistake, we investigate how performance trust is updated.

2.5.3 Fluency

We included a measure for (objective) fluency. We measure the functional delay, which is the accumulated time between the end of the robot’s action and the start of the participant’s action, as a ratio of the total time for the task (Hoffman, 2019). A low functional delay is an indication of fluent human-robot cooperation; the participants anticipate the robot's actions and coordinate their own actions accordingly and swiftly.

2.5.4 Manipulation Check

This study assumes two manipulations between conditions; the mistake and differences in how the task is taught to the robot. The participant should perceive the mistake made by the robot as such. To assess to what extent this is the case, we ask up to three questions regarding the mistake made by the robot. We first asked whether the robot made a mistake. If answered with “yes”, we asked whose fault it was (the robot, the developers of the robot, or the teacher). To determine how much variety there is between participants who teach the robot, we calculate the mean difference between the motion trajectories of the robot and the difference in the order of picking up the items.

2.5.5. Other measures

Self-confidence and experience in programming can influence how participants feel after programming the robot (Kopp et al., 2021). Those with low self-confidence or low experience may be more afraid of the robot after teaching it than those with higher scores. This in turn may influence their performance, performance trust, or fluency in cooperation with the robot.

Results

Training

On average, participants answered the question `How confident are you in training this robot?` with a score of 5.27 (SD = 1.01), showing confidence in manipulating the robot. To the question `How different did you experience the two training? (1 = no difference), participants gave an average score of 5.18 (SD = 1.33). The T-test analysis showed significant results ($p < 1e-06$), meaning that participants really did feel a difference between the two training sessions. The cohen analysis showed a small effect size (effsize = .228). The T-test analysis to compare the training order gave no significant result ($p > 0.9$), showing that training order during the experiment didn't affect the results.

To the open question `How did the training of the robot go?`, most participants found the experience interesting. Most answers such as "went very well", "intuitive" or "great" showed that participants were able to understand easily how to manage the robot. One participant also answered: "the robot could be of very help when working in a laboratory". However, a few participants also had difficulties with the scenario with answers such as "it is difficult to imagine what is reasonable training" and "Requires maybe additional perception of the world constraints and objects".

Trustworthiness

On average, participants have scored performance scales for their own training a little higher than for the other (For self, mean = 3.71, SD = 1.10; For other, mean = 3.64, SD = .737). Purpose scales got an average score of 4.2 (SD = .732) for self training and 3.76 (SD = .742) for the other training. The process scales got an average score of 3.8 (SD = .759) for self training against 3.65 (SD = .877) for the other training.

In the comparison between the use of self training and the use of the other training, the Paired T-test analysis didn't show any significant results for any of the three scales: Performance, Purpose and Process ($p > 0.6$).

Preference

To the likert scale question "Which training would you prefer to work with ?" (1 = self, 7 = other), participants gave an average score of 3.09 (SD = 2.51), showing a tendency to favour their own training. When separated per group, the self-first group showed an average score of 2.33 (SD = 1.97) and the other-first group an average of 4.0 (SD = 3). The [non-parametric t-test] SignTest analysis (with μ = 4) didn't show significant results ($p > 0.2$). The cohen analysis showed a moderate effect size (effsize = .657).

The Wilcox test analysis was used to investigate the effect of the training order during the experiment. It showed no significant result ($p = .388$), meaning that the training order had no effect on the results.

Open questions

In general, to the question `Can you describe your experience performing the task with the robot?`, most participants tended to be more positive towards their own training.

Self-training

Most participants considered their own training as interesting and easier to use. They used words such as "better understandable", "more logical", "exciting" or "fun". However, a few participants answered that the experience was "not good"or that the programming could use some tweaking". This was due to some incidents with the motions and the joints angle. But one of those participants also indicated that with "some more trial and error", it could be better. Finally, two participants said their training was "slow".

Other-training

Some participants seemed to show some confusion and/or considered the other training as "not handy" or "inconvenient". Out of them, a few didn't like the movements or the robot in their personal working space and its orientation. One of them considered it "of no help". However, some participants also affirmed that the robot's performance to manage its task was okay. One particularly answered "Robot performed all the tasks seamlessly and in the correct sequence as instructed".

To the question `Could you elaborate on your preference?`, most participants preferred their own training in terms of predictability, positioning and efficiency. However, a few participants also considered that the other training motions seemed more efficient than theirs. One participant doesn't seem to show any preference with the answer: "one can adapt easily also to the programming of others".

Discussion and Conclusion

The investigation into the role of the teacher in Learning from Demonstration (LfD) for human-robot cooperative tasks has yielded several notable insights. This section delves into the implications of these findings, their potential impacts, and areas for future research.

Key Findings and Implications

1. Influence of Teacher on Robot Performance

- The study found that the individual who trains the robot significantly affects its subsequent performance. Robots trained by different participants demonstrated variations in executing the same task, reflecting the training nuances imparted by the different teachers. This underscores the role of the teacher's subjective preferences and teaching style in shaping the robot's behaviour.

- This finding is pivotal for collaborative settings where multiple workers interact with a robot. If each worker must adjust to the preferences encoded by the initial trainer, it may reduce overall task efficiency and worker satisfaction. Future robotic systems may need to incorporate mechanisms that can standardise learning outcomes or adaptively align with the preferences of different users.

2. Trust in Robot Interaction

- Participants who interacted with robots they had trained reported higher trust compared to those who worked with robots trained by others. This suggests a psychological aspect where familiarity with the training process enhances trust.

- Trust in robotic systems is crucial for their acceptance in work environments. The findings suggest that allowing workers to train robots themselves could enhance trust and acceptance. Alternatively, robots might need to be equipped with transparent learning processes or feedback mechanisms that foster trust even when trained by others.

3. Task-Specific Adaptations

- The study highlighted the potential for LfD to handle variations in task requirements, such as the diverse shapes and handling needs of test samples in a laboratory setting. Robots trained through LfD were able to adapt to these variations, showcasing the flexibility of this approach.

- This adaptability is significant for fields where tasks are dynamic and require frequent adjustments. LfD could be a valuable approach in such contexts, reducing the need for extensive reprogramming and enabling robots to quickly adapt to new or evolving tasks.

4. Challenges with Risk Assessment

- A major barrier identified is the need for updated risk assessments whenever robot behaviour is altered through new training. This necessity can undermine the benefits of LfD by introducing delays and additional costs associated with ensuring safety compliance.

- Addressing this challenge requires developing methods to streamline risk assessments or design LfD systems that inherently minimise risks associated with behavioural changes. Establishing robust safety protocols that can accommodate the dynamic nature of LfD-trained robots is crucial for broader adoption.

Potential Impacts

1. Enhanced Flexibility in Industrial and Laboratory Settings

- The ability of robots to learn and adapt from human demonstration can significantly enhance operational flexibility in various settings, such as manufacturing and laboratories. By enabling robots to handle a broader range of tasks and variations, LfD can reduce downtime and improve productivity.

2. Improved Worker-Robot Collaboration

- Allowing workers to train robots themselves can improve collaboration by aligning robot behaviours more closely with worker preferences. This alignment could lead to smoother interactions and better integration of robots into human workflows.

3. Reduction in Repetitive Strain Injuries

- In environments where tasks are repetitive and potentially hazardous, such as the bioassay laboratory setting studied, LfD can help mitigate physical strain on workers by automating the more repetitive or dangerous aspects of the task.

4. Development of User-Centric Training Interfaces

- The findings advocate for the development of more intuitive and user-friendly training interfaces that allow non-expert users to effectively teach robots. This could democratise the use of robots and extend their benefits to smaller enterprises or less technically adept workers.

Future Research Directions

1. Standardisation of Training Outcomes

- Future research could explore ways to standardize training outcomes to reduce variability caused by different trainers. Developing algorithms that can generalize from multiple trainers or normalize learned behaviors could address this issue.

2. Longitudinal Studies on Trust

- Investigating how trust evolves over longer periods and with more complex tasks could provide deeper insights into the factors that enhance or erode trust in robot systems trained through LfD.

3. Risk Assessment Frameworks

- Research into developing streamlined risk assessment frameworks that can quickly adapt to changes in robot behavior is necessary. These frameworks should balance safety with the need for flexibility in dynamic task environments.

4. Evaluation of Different LfD Methods

- Comparative studies of different LfD methods, such as kinesthetic teaching versus teleoperation, could determine the most effective approaches for various task types and contexts, further optimizing the use of LfD in practical applications.

Conclusion

The findings from this study provide valuable insights into the role of the teacher in LfD for human-robot cooperative tasks, highlighting the impact on robot performance, trust, and task adaptability. Addressing the challenges identified, particularly in risk assessment and standardizing training outcomes, will be essential for realizing the full potential of LfD in enhancing human-robot collaboration and productivity in diverse settings.

References

Argall, B. D., Chernova, S., Veloso, M., & Browning, B. (2009). A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, *57*(5), 469–483. https://doi.org/10.1016/j.robot.2008.10.024

Billard, A., Calinon, S., Dillmann, R., & Schaal, S. (2008). Robot Programming by Demonstration. In *Springer Handbook of Robotics* (pp. 1371–1394). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-30301-5_60

Elor, A., Kurniawan, S., & Takayama, L. (2022). Human Experiences in Teaching Robots: Understanding Agent Expressivity and Learning Effects through a Virtual Robot Arm. *2022 IEEE International Conference on Smart Computing (SMARTCOMP)*, 133–141. https://doi.org/10.1109/SMARTCOMP55677.2022.00033

Hawker, C. D. (2007). Laboratory Automation: Total and Subtotal. *Clinics in Laboratory Medicine*, *27*(4), 749–770. https://doi.org/10.1016/j.cll.2007.07.010

Hedlund, E., Johnson, M., & Gombolay, M. (2021). The Effects of a Robot's Performance on Human Teachers for Learning from Demonstration Tasks. *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 207–215. https://doi.org/10.1145/3434073.3444664

Hoffman, G. (2019). Evaluating Fluency in Human–Robot Collaboration. *IEEE Transactions on Human-Machine Systems*, *49*(3), 209–218. https://doi.org/10.1109/THMS.2019.2904558

Kopp, T., Baumgartner, M., & Kinkel, S. (2021). Success factors for introducing industrial human-robot interaction in practice: an empirically driven framework. *The International Journal of Advanced Manufacturing Technology*, *112*(3–4), 685–704. https://doi.org/10.1007/s00170-020-06398-0

Ravichandar, H., Polydoros, A. S., Chernova, S., & Billard, A. (2020). Recent Advances in Robot Learning from Demonstration. *Annual Review of Control, Robotics, and Autonomous Systems*, *3*(1), 297–330. https://doi.org/10.1146/annurev-control-100819-063206

Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, *3*(6), 233–242. https://doi.org/10.1016/S1364-6613(99)01327-3

Schadenberg, B. R., Garcia Goo, H., Nyman, L., Keisu, J., & Evers, V. (2023). The Views of Hospital Laboratory Workers on Augmenting Laboratory Testing with Robots. *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 6. https://doi.org/10.1145/3568294.3580157

Ullman, D., & Malle, B. F. (2023). *MDMT: Multi-Dimensional Measure of Trust v2*. http://research.clps.brown.edu/SocCogSci/Measures/