

2016-03-11

Providing a Robot with Learning Abilities Improves its Perception by Users

Senft, E

<http://hdl.handle.net/10026.1/12112>

10.1109/HRI.2016.7451832

Proceedings of the 2016 ACM/IEEE Human-Robot Interaction Conference

IEEE

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Providing a Robot with Learning Abilities Improves its Perception by Users

Emmanuel Senft, Paul Baxter, James Kennedy, Séverin Lemaignan and Tony Belpaeme
Centre for Robotics and Neural Systems

Plymouth University, U.K.

{emmanuel.senft, paul.baxter, james.kennedy, severin.lemaignan, tony.belpaeme}@plymouth.ac.uk

Abstract—Subjective appreciation and performance evaluation of a robot by users are two important dimensions for Human-Robot Interaction, especially as increasing numbers of people become involved with robots. As roboticists we have to carefully design robots to make the interaction as smooth and enjoyable as possible for the users, while maintaining good performance in the task assigned to the robot. In this paper, we examine the impact of providing a robot with learning capabilities on how users report the quality of the interaction in relation to objective performance. We show that humans tend to prefer interacting with a learning robot and will rate its capabilities higher even if the actual performance in the task was lower. We suggest that adding learning to a robot could reduce the apparent load felt by a user for a new task and improve the user’s evaluation of the system, thus facilitating the integration of such robots into existing work flows.

I. INTRODUCTION

This paper presents a study exploring the impact of providing a robot with learning capabilities on the interaction preferences and robot performance evaluations by users.

Two main approaches are reported in the literature to study human preferences about robots. The first one involves the administration of surveys where participants are asked robot-related questions with or without priming. For example, in [1] 240 subjects are asked questions about tasks that could be replaced by robots and about general attitudes toward robots, without trying to influence the participants *a priori*. Priming can also be a useful means of educating the participants before administering a questionnaire, allowing them to imagine a more constrained and plausible scenario than they otherwise would. This approach has been followed by Coeckelbergh et al. [2], who surveyed the attitudes of participants toward Robot Assisted Therapy (RAT) for children with autism spectrum disorder. The participants answered more positively, in contrast to previous studies conducted without priming, when they were first exposed to a one minute video presenting the state of the art of robotics in RAT.

The second main approach is administering a questionnaire to participants after an actual interaction with a robot. Using this method, the responses are grounded in the context of their interaction: this can diminish the generalisability of the results, but makes them more reliable. This method has been applied to explore how elderly people react to a robot with learning abilities [3].

This paper follows the real robot interaction approach, and presents additional results gathered in the experiment

presented in [4]. In this study, participants interacted with a robot both with and without learning capabilities, and this manuscript reports their interaction preference and their relative performance evaluation of the two robots.

II. METHODOLOGY

The study (and therefore the methodology) is the same as in [4] where we introduced Supervised Progressively Autonomous Robot Competencies (SPARC), a means for the robot to learn from the interaction to improve its capabilities. This previous paper also reported the impact of SPARC on the performance and workload of a robot’s human supervisor in a scenario inspired by RAT for children with autism spectrum disorder. In classical RAT, the robot is interacting with the child and is often controlled using the Wizard of Oz (WoZ) paradigm, i.e. fully tele-operated. This often implies a high workload on the therapist, which could be reduced by providing the robot with a supervised autonomy. As this study focuses on the interaction between the *wizarded-robot* and its supervisor, we replace the child with a robot interacting in his place (the *child-robot*) to produce consistent experimental conditions (fig. 1).




Fig. 1. Installation used for the study. The *child-robot* stands on the left, performing the task on the touchscreen, and facing the *wizarded-robot* on the centre-right. The human supervisor can control the action about to be executed by the *wizarded-robot* using the GUI on the right.

The child-robot is interacting with a touchscreen, and performs a categorisation task where it has to classify images of face as either happy or sad, with the aim of improving its performance. The wizarded-robot can execute actions (e.g. giving positive or negative feedback, waiting, or prompting the child to act), aiming to help the child-robot in its task.

The participants have to control the wizarded-robot to make it execute the correct actions to help the child-robot. This is however context dependent: actions can either improve or worsen the performance of the child-robot based on its current state. A Graphical User Interface (GUI) allows the users to control the wizarded-robot in a WoZ inspired scenario involving supervised autonomy. At specific times, the wizarded-robot makes suggestions to the supervisor who can either not react and let the action execute, or use a button to force the wizarded-robot to execute another action. A habituation phase allows the participants to become familiar with the interface and action set. If the suggestion of the robot is correct, the supervisor does not need to act to have this action executed.

The participants interacted with two systems. In the first system, the actions proposed by the wizarded-robot are random, so we expect the user to correct them most of the time. This system simulates a classical WoZ setup, which we denote the *non-learning robot*. The second system uses SPARC and includes a learning algorithm based on a Multi-Layer Perceptron using noisy observation of states as inputs and a winner-take-all on the actions as output. This system is referred to as the *learning-robot*. It is important to note that in both systems, these terms relate to the capabilities of the wizarded-robot, not the child-robot (which had constant behaviour in both systems).

The study involved ten participants (7M/3F, age $M=29.3$, 21 to 44, $SD=4.8$ years) taken from a robotic research group, as typical WoZ users are technical. Each participant interacted for 10 minutes with both systems, with the order counterbalanced. In the LN condition, participants interacted first with the learning robot then with the non-learning one, and the order is reversed in the NL condition. This paper presents and analyses the responses from the participants to the questions:

- Which *wizarded-robot* was better able to perform the task?
- Which *wizarded-robot* did you prefer supervising?

III. RESULTS AND DISCUSSION

Overall, the participants preferred supervising the learning-robot (6 out of 10) and found it better able to perform the task (8 out of 10). Despite the limitations of the small sample size, these results suggest that providing a robot with learning capabilities can improve its perception by users and also make the users prefer supervising it. These results are consistent with previous results [4], which showed that providing a robot with learning capabilities can decrease the number of interventions required to achieve a similar performance compared to a robot without learning. The reduction in the number of interventions needed might explain the results observed here.

Breaking the results down by ordering condition (LN vs. NL) provides a more detailed perspective (fig. 2). From these separated results, the learning capability is not the only effect influencing the preferences of, and the evaluation by, the participants: the order of interaction also plays an important role. On average, the second robot is the preferred one to supervise (7 of 10) and rated as better able to perform the task (7 of 10). This ordering effect was probably due to the




Fig. 2. Results for supervisory preference and rating of ‘preferred to supervise’ and ‘better to perform the task’ for the two conditions. The vertical bars represent the number of times that the learning robot was selected and the horizontal dotted line denotes chance (i.e. 50%).

complexity of the system that the participants interacted with. The participants had to get used to a GUI displaying a large volume of information, and to the time constraints.

Additionally, in 4 of the 5 cases when participants interacted with the learning robot first, they achieved a better performance during the second interaction than during the first one. Three of these participants also rated the learning robot as better able to perform the task even when it had a lower performance. This could indicate that participants can distinguish between the robot’s abilities and the performance achieved (depending also on their abilities). It could also be a reflection of the natural propensity of humans to adapt and learn through interaction. Viewed in this way, the results could be interpreted as showing that interaction with the learning robot first better equips the human to interact with the non-learning robot than vice-versa, leading to higher performance, and hence preference ratings, for the non-learning robot in the LN condition. While another potential benefit of learning robots, this interpretation will require further empirical investigation.

In this paper we presented results showing a trend towards the addition of learning capabilities to a robot helping users to cope with a new or complex task, and improving the rating of their performance by their supervisor. This is an important point for design, especially when there is a heavy workload on users such as in RAT when therapists have to use WoZ to continuously control the robot.

IV. ACKNOWLEDGEMENTS

This work is funded by the EU FP7 project DREAM (grant 611391).

REFERENCES

- [1] C. Ray *et al.*, “What do people expect from robots?” in *Int. Conf. on Intelligent Robots and Systems*. IEEE/RSJ, 2008.
- [2] M. Coeckelbergh *et al.*, “A Survey of Expectations About the Role of Robots in Robot-Assisted Therapy for Children with ASD: Ethical Acceptability, Trust, Sociability, Appearance, and Attachment,” *Science and Engineering Ethics*, pp. 1–19, 2015.
- [3] J. Hoefinghoff *et al.*, ““Yes Dear, that Belongs into the Shelf!”-Exploratory Studies with Elderly People Who Learn to Train an Adaptive Robot Companion,” in *Proc. of the 7th Int. Conf. on Social Robotics*, 2015.
- [4] E. Senft *et al.*, “SPARC: Supervised Progressively Autonomous Robot Competencies,” in *Proc. of the 7th Int. Conf. on Social Robotics*, 2015.