

# Communication and Knowledge Sharing in Human-Robot Interaction and Learning from Demonstration

Nathan Koenig<sup>1</sup> and Leila Takayama<sup>2</sup> and Maja Mataric<sup>1</sup>

nkoenig@usc.edu, takayama@willowgarage.com, maja@usc.edu

<sup>1</sup>Viterbi School of Engineering, Univ. of Southern California, Los Angeles, CA

<sup>2</sup>Willow Garage, 68 Willow Road, Menlo Park, CA

## Introduction

People gain new experiences, knowledge, and skills throughout life, by combining innate abilities [2] and both formal and informal learning [4]. If robots are to become ubiquitous, they too must be able to learn throughout their lives, both from direct programming and informally, from observations and experiences. One promising approach is *learning from demonstration* (LfD) [1], which relies on a human instructor to actively guide a robot student through a task. Observations gathered by the robot provide the necessary data to generate a task policy for the given demonstration.

The success of an LfD system relies on the understanding of all the factors associated with human-robot interaction in the specific teaching context. This paper considers the effect of face-to-face communication, which plays a significant role in human communication [3], on LfD. Specifically, we study the impact that the user's visual access to the robot, or lack thereof, has on teaching performance. Based on the obtained results, we then address how a robot can provide appropriate information to a instructor during the LfD process, to optimize the two-way process of teaching and learning.

## Approach

Two independent studies were implemented to measure the effects of visual obstruction and robot audio feedback on robot LfD. The first study utilized a screen to block visual access to the robot for a subset of instructors while the others had unfettered visual access. The second study incorporated helpful audio cues produced by the robot learner, which provided an additional feedback channel for the instructor. These audio cues consist of three non-verbal sound clips that convey *success*, *error*, and *acknowledgment*. A Mechanical Turk [10] study consisting of 100 participants validated these sounds using an online survey.

In both studies, instructor performance was measured using demonstration duration, the number of commands issued to the robot, and the instructor's perceived level of frustration and effort. The first two measurements were purely quantitative, and provide insight into how well the instructor is giving a demonstration. The final set of measurements are qualitative, and express the degree of difficulty felt by the




Figure 1: The PR2 during the (a) Towers of Hanoi and (b) box sorting demonstrations. The graphical interface on the Nokia N810 used by an instructor is shown in (c). A video of the PR2 completing the Towers of Hanoi puzzle can be viewed here: <http://www.youtube.com/watch?v=k3m0DKSgipI>

instructors and the likelihood they perceived of providing correct and multiple demonstrations.

Both studies took place in the context of LfD and utilized a unique task. The first study used a three-disc version of Towers of Hanoi task shown in Figure 1a, while the second study employed a version of box sorting shown in Figure 1b. A Willow Garage [9] PR2 was used as the robot student, and instructions to the robot were controlled through a graphical user interface (GUI), see Figure 1c.

## Results

The results of these studies show that the way in which we perceive and understand objects in our environment is highly dependent on our sensory information. In the first study, instructors who could see the robot sent more valid commands to the robot ( $M=47.1$ ,  $SD=12.3$ ) than instructors who could not see the robot ( $M=36.8$ ,  $SD=7.7$ ),  $F(1,18)=6.12$ ,  $p < .03$ . The instructors who could see the robot also sent more commands (in total) to the robot ( $M=36.7$ ,  $SD=7.7$ ) than instructors who could not see the robot ( $M=29.9$ ,  $SD=4.0$ ),  $F(1,18)=5.04$ ,  $p < .04$ . This means the visible instructors sent duplicate commands, akin to repeatedly hitting a key, and also deviated more from the optimal Towers of Hanoi solution.

The mean time on task increased for instructors who could see the robot, and also had more variance. Alternatively, instructors who were visually blocked from the robot performed more consistently, and were able to complete the task in a slightly shorter timespan. Robot visibility as an indicator of duration is not significant,  $F(1,18)=2.58$ ,  $p=.13$ ; however these data do suggest a trend. The time on task is closely correlated to the number of commands issued, but it also includes time spent doing other activities such as observing the robot.

The second study followed a similar pattern. When instructors had visual access to the robot, they gave slightly more commands to the robot in total ( $M=27.1$ ,  $SD=2.1$ ) than when they did not have visual access to the robot ( $M=26.4$ ,  $SD=2.2$ ),  $F(1,40)=3.99$ ,  $p=.01$ .

The number of erroneous commands in the second study were marginally predicted by the presence of auditory feedback from the robot. Instructors who heard auditory feedback gave fewer erroneous commands to the robot ( $M=1.9$ ,  $SD=1.5$ ) than people who did not have any auditory feedback ( $M=3.3$ ,  $SD=2.3$ ),  $F(1,40)=4.01$ ,  $p=.056$ .

Neither auditory nor visual feedback from the robot was found to be a significant predictor of time on task, so we ran a regression analysis of robot experience as a predictor. This showed that the more experience people had with robots, the less time they spent on the task,  $\beta=-6.30$ ,  $p < .01$ .

## Experiments

The first study on robot visibility used a 2-level (directly visible vs. visually occluded) between-participants experiment design. This controlled experiment investigated the influences of robot visibility upon human-robot interaction, specifically how human instructors would perform when demonstrating to a robot the Towers of Hanoi puzzle.

Twenty volunteers, six of whom were female, participated in the study, with ages ranging from 22 to 59. The six females were evenly distributed between the visible and non-visible teaching conditions. Each participant was randomly assigned to one of two experimental conditions. Half of the participants were allowed direct observation of the robot while conducting a demonstration. The other participants could not see the robot due to a screen behind which they were placed. These instructors could only rely on the GUI for feedback from the robot.

The second study, incorporating audio cues, used a 2x2 (audio cues enabled vs. disabled, robot directly visible vs. no visualization) experimental design. Audio cues were varied between subjects while visibility was varied within subjects. The goal was to investigate the effects audio cues and student visibility have on interaction in a teaching context.

The second study also had twenty volunteers, with ages ranging from 22 to 60 and an even gender distribution. The men and women were evenly distributed among the test conditions, and the order in which the test conditions were run was randomized. Each participant completed two demonstrations. Robot visibility was changed between each demonstration. The order of the demonstrations and robot visibility

were varied between subjects. Half of the participants received audio cues from the robot, while the other half did not. This resulted in four test cases based on the two conditions: visibility and sound.

## Discussion

The major finding from the two studies is the effect that the visibility of a robot student has on the performance of the instructor in the LfD context. In both tasks, the number of commands issued to the robot increased when the instructor was allowed visual access to the robot. The time on task also tended to increase as well as the number of erroneous commands.

Instructors who could see the robot learner received information from the robot that lacked context. As noted in prior work, merely increasing the amount of information provided does not necessarily increase situational awareness [5]. It is in fact more efficient to improve the display and throughput of information from the robot's sensors to the human in order to decrease confusion [6].

Other work has indicated that showing a humanoid robot to a user automatically causes the user to construct a mental model of the robot [8]. Thus, restricting some instructors from seeing the robot impairs the ability to form such a model and thereby assess the robot's state.

Based on comments from the instructors, collected data, and observations from the Towers of Hanoi study, it was clear that a instructor needs more useful feedback from the robot than just visual observations. This trend has also been noted in [7], where instructors did not wait for the robot to complete an action before providing more input.

Our results also show that sound cues helped to improve the performance of instructors. The additional channel of communication provided instructors with vital information about the internal state of the PR2. Comments from the instructors indicated that the audio cues properly informed them as to when the robot was ready to accept an instruction, experienced an error, and received a new instruction. Visibility of the robot still impacted instructor performance, however the audio cues tempered this effect.

## References

- [1] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, *Robot Programming by Demonstration*, ch. 59, MIT Press, 2008.
- [2] J. Gavin Bremner, *Infancy*, Wiley-Blackwell, Malden, MA, 1994.
- [3] Hebert H. Clark, and S. A. Brennan, *Grounding in communication*, Perspectives on socially shared cognition (L. B. Resnick, J. M. Levine, and S. D. Teasley, eds.), 1991.
- [4] Jay Cross, *Informal Learning: Rediscovering the Natural Pathways That Inspire Innovation and Performance*, Pfeiffer, November 2006.
- [5] M. R. Endsley, *Design and evaluation for situation awareness enhancement*, Proc. of the Human Factors Society 32nd Annual Meeting (Santa Monica, CA), Human Factors Society, 1988.
- [6] Kevin Gold, *An information pipeline model of human-robot interaction*, HRI '09: Proceedings of the 4th ACM/IEEE international conference on Human robot interaction (La Jolla, CA, USA), ACM, 2009, pp. 85-92.
- [7] Elizabeth S. Kim, Dan Leyzberg, and Katherine M. Tsui, and Brian Scassellati, *How people talk when teaching a robot*, HRI '09: Proceedings of the 4th ACM/IEEE international conference on Human robot interaction (La Jolla, CA, USA), ACM, 2009, pp. 23-30.
- [8] Aaron Powers, and Sara Kiesler, *The advisor robot: tracing people's mental model from a robot's physical attributes*, HRI '06: Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction (Salt Lake City, UT, USA), ACM, 2006, pp. 218-225.
- [9] K. Wyrobek, E. Berger, H.F.M. Van der Loos, and K. Salisbury, *Towards a Personal Robotics Development Platform: Rationale and Design of an Intrinsically Safe Personal Robot*, Proc International Conference on Robotics and Automation, 2008.
- [10] Amazon, Mechanical Turk, <https://www.mturk.com>, 2010