Interactive Personalization for Socially Assistive Robots

Caitlyn E. Clabaugh
University of Southern California
3710 McClintock Ave
Los Angeles, CA 90089-0781
clabaugh@usc.edu

ABSTRACT

In this work, we seek to define a new problem of interactive personalization in the context of socially assistive robotics. We analyze a robotic tutor's elicitation of learning-sensitive information to be leveraged by interactive machine learning methods for personalized education. Our results, evaluated using a variety of subjective measures, demonstrate that a humans-in-the-loop approach positively benefits the human-robotic tutor interaction, while minimizing the computational complexity of personalization.

Keywords

human robot interaction; personalization; user modeling

1. INTRODUCTION & RELATED WORK

Human-robot interactions (HRI) can be designed to improve robotic systems through input from human researchers or end users. These inputs can be formulated as feedback, directives, or examples given by the human to the robot, or elicited from the human by the robot. This approach – broadly known as Interactive Machine Learning (IML) [5] or Learning from Demonstration (LfD) [3] – applied to HRI, leverages human users as untapped resources for additional information that can guide the robotic system's learning of its environments, tasks, or users themselves.

Socially Assistive Robots (SAR) aims to supplement human care, coaching, and companionship and improve quality of life through social, goal-oriented HRI [8]. *Personalization* of interaction and instruction is a long-standing goal of HRI and SAR [7], the benefits of which are well-supported by research across the disciplines of psychology [4] and human-computer interaction (HCI) [1]. However, individuals vary greatly in terms of their personal abilities, special needs, learning styles, and preferences [2], and the state of human users is non-deterministic and partially observable at best; such computational problems are theoretically and empirically hard.

The computational challenges of personalized SAR can be alleviated by including humans-in-the-loop, guiding the robot's personalization process through interactive feedback. It is crucial, however, to consider that humans can and often do provide mis-

information [9], and its effects on the user's overall experience, acceptance, and attachment.

In this work, we consider the case in which a robot may directly and socially elicit input from a human user to better personalize its assistive, tutoring interaction. We define this problem of *Interactive Personalization*, or the process through which a SAR may increase its assistive understanding and performance with respect to its human interaction partner. We explore measures and methods in Interactive Personalization, seeking interdisciplinary insights into the computational and social challenges and benefits thereof.

2. INTERACTIVE PERSONALIZATION

We define **Interactive Personalization** as the process by which an intelligent agent adapts to the needs and preferences of an individual user through eliciting information directly from that user about his or her state [6]. Certain information about a human learner may only be observable through direct input by the learner. However, such information may also be sensitive in nature as it may relate to the learner's identity, expose the learner's weaknesses, or create phenomena like survey fatigue if elicited too frequently. We, thus, consider an extensive set of social and computational trade-offs below as the parameterization of Interactive Personalization.

2.1 Sensitivity

Learning-sensitive information includes any measurable data about a human subject that effects his or her ability or style by which to achieve some learning goal [6]. Thus, we define **sensitivity** as the qualitative connection between the input elicited by the robot and the human learner's identity.

2.2 Frequency

We define **frequency** of elicitation as the number of learning-sensitive questions q posed by a robotic tutor over the total number of learner tasks K – $f = \frac{q}{K}$ [6]. In the following section 3, we describe our initial exploration into this parameter.

2.3 Scope

Questions posed by the robot can have varying levels of **scope** defined as binary, ordinal, and open-ended. Questions with binary or ordinal scope may be easier to quantify; however, open-ended questions may provide more information to the robot and more interactivity.

2.4 Relevance

Relevance is defined as the probabilistic correlation or dependency between the learner's response to the robot's question q and the learner's outcome or performance in task k . For example, for a set of possible outcomes $[x, y, z]$, if $P(k = x | q = x) = 1$, then q is a

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

HRI '17 Companion, March 6–9, 2017, Vienna, Austria.

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4336-7/17/03..

DOI: <http://dx.doi.org/10.1145/3029798.3034813>

highly relevant question to task k , whereas if $P(k = x|q = x) = 0$, then q is a highly irrelevant question to task k .

2.5 Transparency

When to robot elicits information from the human learner, it may be phrased as an act of assistance to the robot. The robot may make an appeal to the human learner that it could use the information or feedback to better understand how to assist him or her. We define this as the binary feature of **transparency**. If $t = 0$, the robot asks the learner for feedback or input without exposing the reason for that need, or with zero transparency, and when $t = 1$ the robot asks appeals to the human learner for feedback or input that may be mutually beneficial to their joint goal.

2.6 Empathy

When to robot elicits information from the human learner, it may phrase the question with **empathy** or sympathetically to the learner's struggle or cause. Defines as the binary feature for simplicity, empathy $e = 0$ means the robot asks the learner for feedback or input without appeal to the learner's emotional state or identity, and when $e = 1$ the robot asks with a synthetic understanding of the learner's plight.

2.7 Presence

HCI and HRI research studies agents of varying embodiments; agent **presence** can significantly weight the outcomes of the above parameters. For instance, it is more challenging to express empathy as a disembodied agent than one with affective actuators.

3. STUDY ON FREQUENCY

In our previous work, we have explored some limited aspects of Interactive Personalization. Specifically, we studied the relationship between *frequency of elicitation* and the human learners' impressions of the SAR tutoring interaction. The higher the frequency, the more input is being asked of the human learner.

3.1 Experimental Design

In a small-scale study, conducted with 38 undergraduate and graduate students, we tested two hypotheses relative to frequency of elicitation of learning-sensitive information and the interaction and impression of a robotic tutor. We hypothesized that 1. Interactive Personalization would provide social gains in addition to computational gains, i.e., a robotic tutor that elicits learning-sensitive information ($f > 0.0$) would be considered more social, interactive, and intelligent than one that does not ask the human learner about his or her state, and 2. too much elicitation ($f > 0.5$) would cause discomfort or annoyance to the human learner with minimal computational gains.

We designed and tested three discrete conditions: 1. The robot never elicits learning-sensitive information from the participant, $f = 0.0$. Thus, the interaction consists of a series of questions, followed by the robot's response as to whether the participant answered correctly or incorrectly. 2. The robot elicits learning-sensitive information 25% of the time, $f = 0.25$, evenly distributed. Thus, in our short interaction, the robot asks the user a learning-sensitive question every 5 problems. 3. The robot elicits learning-sensitive information 50% of the time, $f = 0.5$, evenly distributed. Thus, in our short interaction, the robot asks the user a learning-sensitive question every other problem.

3.2 Results

Among our three conditions, we found that participants felt they were more stimulated ($p < 0.05$) in the third condition ($f = 0.5$),

in which the robotic tutor elicited learning-sensitive information more frequently. Additionally, participants reported that they communicated significantly more ($p < 0.005$) with the robotic tutor in the second ($f = 0.25$) and third ($f = 0.5$) conditions, and that they felt the robotic tutor was generally more interactive in the second ($f = 0.25$) and third ($f = 0.5$) conditions ($p < 0.05$). Lastly, participants were more likely to report that the robotic tutor had feelings ($p < 0.05$) and that they were more likely to trust the robot's advice ($p < 0.01$) in the case in which the robot posed the most learning-sensitive questions ($f = 0.5$).

4. DISCUSSION & FUTURE WORK

These results confirm the interest of Interactive Personalization as a means by which to increase the robot's performance with greater information about the human learner as well as improve the social components of the interaction. Through our previous research, we are motivated to explore more aspects of Interactive Personalization, some of which we have defined in this abstract.

We intend to explore the scaffolding necessary for a robotic tutor to elicit more sensitive, and perhaps difficult to observe, information from a human learner. Ultimately, we intend to leverage insights about Interactive Personalization gained with convenience populations and apply them to assistive domains such as young children with special needs.

5. ACKNOWLEDGMENTS

This work is supported by the National Science Foundation Expeditions in Computing grant for "Socially Assistive Robotics" (NSF IIS-1139148).

6. REFERENCES

- [1] J. R. Anderson, C. F. Boyle, A. T. Corbett, and M. W. Lewis. Cognitive modeling and intelligent tutoring. *Artificial intelligence*, 42(1):7–49, 1990.
- [2] L. W. Anderson, D. R. Krathwohl, and B. S. Bloom. *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*. Allyn & Bacon, 2001.
- [3] B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- [4] B. S. Bloom. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational researcher*, pages 4–16, 1984.
- [5] J. Canny. Interactive machine learning. *University of California, Berkeley*, 2014.
- [6] C. Clabaugh and M. J. Matarić. Exploring elicitation frequency of learning-sensitive information by a robotic tutor for interactive personalization. In *Robot and Human Interactive Communication (RO-MAN), 2016 25th IEEE International Symposium on*, pages 968–973. IEEE, 2016.
- [7] K. Dautenhahn. Socially intelligent robots: dimensions of human-robot interaction. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362(1480):679–704, 2007.
- [8] D. Feil-Seifer and M. J. Matarić. Socially assistive robotics. *IEEE Robotics & Automation Magazine*, 18(1):24–31, 2011.
- [9] W. B. Knox, B. D. Glass, B. C. Love, W. T. Maddox, and P. Stone. How humans teach agents. *International Journal of Social Robotics*, 4(4):409–421, 2012.