

Developing Adaptive Social Robot Tutors for Children

Aditi Ramachandran and Brian Scassellati

Yale University
Department of Computer Science
New Haven, CT 06511

Introduction

There has been a large body of research demonstrating that students that receive one-on-one tutoring perform, on average, significantly better than students learning via conventional classroom instruction when tested on the same material (Bloom 1984; VanLehn 2011). During tutoring, the teacher has the ability to tailor the instruction and support to the individual learner, creating a personalized learning environment for each student. Research involving robotic agents as tutors indicates that the physical presence of a robot tutor can increase cognitive learning gains (Leyzberg et al. 2010). Further research shows that a robot tutor employing relatively simple personalization strategies can benefit the learner (Leyzberg, Spaulding, and Scassellati 2014). This motivates the need to more deeply investigate robotic tutoring systems as an effective method of instruction.

In order to build a robotic tutoring system that adapts to an individual child, many key research questions need to be addressed. A student's knowledge level and preferences related to how he or she learns best is known to be extremely difficult to model. Even with a student model that captures relevant aspects of a student's knowledge and preferences, it is not straightforward to plan what personalized actions the robot tutor should select for a given individual to maximize engagement or learning gains. Furthermore, children can be unpredictable in learning environments, emphasizing the need for the robot to have the capability to respond in real-time to their behavior, including the presence of learning-centric affective states, such as confusion or boredom.

This paper describes an ongoing human-robot interaction study aimed at understanding how children seek help during a tutoring interaction with a social robot, as well as how simple adaptive strategies can be used to foster productive behavior and learning over multiple tutoring sessions.

Robot Tutors that Provide Adaptive Support

Our work focuses on one salient aspect of a tutoring interaction that human tutors are well suited for, which is providing help to the student at the right time (Merrill et al. 1992). While there has been much research demonstrating




Figure 1: Child interacting with a NAO robot in a tutoring scenario

that on-demand help is useful in interactive learning environments (ILEs) and can positively affect learning, students only benefit from this help when it is used effectively (Alevin et al. 2003). Two established patterns of unproductive help-seeking behavior in ILEs involve learners who are help-averse, and learners who game the system (Alevin et al. 2004; Baker et al. 2008). In this context, gaming the system has been defined as “attempting to succeed in a learning environment by exploiting properties of the system rather than by learning the material and trying to use that knowledge to answer correctly”, and examples of this in a tutoring context include systematic guessing and rapid hint requests (Baker et al. 2006). As social robots are novel learning companions, it is important to understand what types of help-seeking strategies seen with ILEs transfer to interactions with social robot tutors. We seek to design interactions that contain adaptive support strategies to combat sub-optimal help-seeking behaviors. This will ultimately help students engage with their learning environment more effectively, and cause higher learning gains.

In order to foster long-term learning gains, we are interested in both the student's learning progress throughout the interaction, as well as sustaining the student's level of engagement over time. There has been much work demonstrating the inter-related nature of emotions and learning (Pekrun et al. 2002). Therefore, it is crucial that these social robot tutors simultaneously monitor both affective state and progress on the learning task to provide a personalized interaction, thereby maximizing the potential for the child to learn.

Our research aims to build adaptive tutoring robots for one-on-one tutoring interactions for children. The robot

should be able to provide support (such as hints, or social support) specific to each individual. This support should adapt to the child over time as well as change based on the child's behavior in real-time. The specific contributions of this paper include understanding why productive help-seeking behavior is an important component of tutoring, outlining a data collection procedure, and detailing how to use the collected data to examine the effects of having a robot employ simple adaptive strategies that impact a child's help-seeking behavior over time.

Help-seeking Differences in Children

In order to build an adaptive system for a robot providing help to a child during a learning interaction, we first must understand what types of help-seeking behaviors children engage in when interacting with a tutoring robot. We design an HRI study that allows children (approximately fifth-grade aged) to interact with the robot while completing math problems on a tablet device (Figure 1). Each child completes a series of fractions problems in four one-on-one sessions with a NAO robot. In each of the four sessions, the child has the option to ask for a hint from the robot three times by pressing a button at the bottom of the tablet. We will collect metrics such as the number of hints requested and time between hint requests for each child during each session. To understand various help-seeking profiles, we focus on session one, in which we allow children to ask the robot for help whenever they would like, as long as a hint is available for the given math problem. In this session, the robot will provide the hint whenever it is requested by the child. Collecting all transactions between the learner, the robot, and the tablet will allow us to use unsupervised learning methods, specifically clustering techniques, to understand various help-seeking behaviors we see the children exhibit. We would like to verify whether groups such as help-averse learners and learners who game the system emerge in tutoring interactions with a social robot. Understanding the nature of these help-seeking behaviors will be the first step in designing an adaptive help system for a social robot during a learning interaction. As each child will complete four sessions, we will examine how these behaviors change over time for each individual student.

Before the first interaction, each child will fill out a questionnaire called the Self-Regulation Questionnaire - Academic (SRQ-A), which will assess their level of academic self-regulation (Ryan and Connell 1989). As there has been established connections between self-regulated learners and academic performance, behaviors, and willingness to ask for help, we want to understand the correlations between a child's level of academic self-regulation and the help-seeking behaviors he or she employs during these tutoring interactions. This may provide a basis for designing personalized supportive behaviors for a social robot.

This HRI study will contain two groups of children (see study design in Figure 2). In sessions two through four, the adaptive group will receive robot behavior that employs two simple strategies promoting productive help-seeking behavior. The robot will automatically provide a hint if the child does not request one after multiple incorrect attempts to ad-

dress help-aversion and require an attempt from the child before providing many consecutive hints to address gaming the system. The control group will interact with the robot in sessions two through four in the same way they interacted during session one; in this condition, the robot will provide a hint to the child whenever the child requests one. Because both groups contain an identical session one where the robot provides a hint whenever asked, we can use session one to assess each child's baseline help-seeking behaviors. Baseline help-seeking behavior can be characterized by events from the tutoring session, including features such as total number of hints requested, number of sub-optimal behaviors exhibited, and percentage of problems that a hint was requested. We can use this in comparison to each child's behavior in session four to understand whether these simple strategies were able to shape more productive help-seeking behaviors for the adaptive group. We will assess this by looking at the number of times the adaptive robot behavior was triggered (or would be triggered for the control group) and how this number changes between session one and four. We will also measure the child's perception of the robot using various Likert scale questions as well as measure learning gains over time using a pretest and posttest. We are currently in the process of running this HRI study and collecting this data.

Building a User Model of Adaptive Support

After thoroughly analyzing all data from the previously described HRI study designed to understand help-seeking differences in children as well as how they change across sessions, we will use this information to create a more robust user model of how to adaptively provide support during a tutoring interaction. Our previous studies will help us understand how can we identify extreme help-seeking behaviors and shape these behaviors into being more productive over several interactions. Our work aims to build a more generalizable user model of when a social robot tutor should intervene during a learning interaction. This includes monitoring and responding to behaviors other than just help-aversion and gaming the system, and therefore requires the robot to continuously adapt online in order to maintain productive help-seeking behavior for each student throughout several interactions.

We will use association rule mining and sequential pattern mining techniques to analyze data from our previous study to identify which behaviors necessitate action from the robot and build a system that can identify these relevant behaviors. Understanding these trends will allow us to design a more robust user model that will shape productive help-seeking behavior based on the individual user, rather than a rough characterization of the user upfront. Some examples of adaptive support the robot will provide will be knowing when to automatically provide a hint, rate-limiting hints when exploitative behavior is detected, as well as providing social support and encouragement when needed. The adaptive component of the user model requires the robot to maintain the state of each user over time. Typically, this is very computationally challenging in an interactive learning environment due to the vast number of features that can be col-

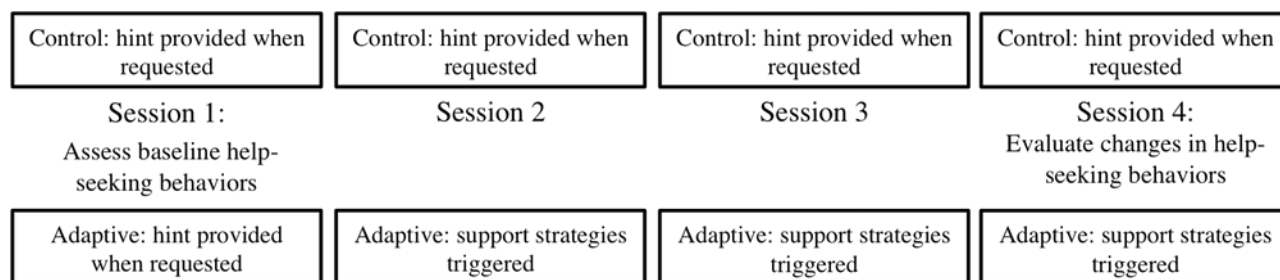


Figure 2: Experimental design for four-session tutoring HRI study

lected and processed continuously throughout each session. We will develop a user model for this context that maintains relevant, salient features for the each user across multiple sessions, including context-sensitive variables such as learning performance, which needs to be updated each time the student answers a question during the learning interaction.

Support Based on Multiple Types of Feedback

Rather than just observing the child's progress through the learning task at hand and their exhibited help-seeking behaviors, detecting the child's affect will allow the robot to better personalize to the individual student as children often behave differently on various days, depending on their mood or affective state.

We will use results from the previous HRI studies conducted to outline a reinforcement learning framework to account for affective feedback from the child, thereby allowing the adaptive robot to use this information to further tailor its support strategies. This approach relies on detecting affective states that can be used as part of a reward function in learning what support strategies should be used for a given child at a given time. Understanding that detecting affective states reliably is extremely challenging, we hypothesize that allowing the robot to use real-time information to learn optimal actions by observing various rewards based on learning progress and affective feedback will lead to a more engaging experience for each user, as well as more effectively shape productive help-seeking behaviors in children during tutoring interactions.

References

Aleven, V.; Stahl, E.; Schworm, S.; Fischer, F.; and Wallace, R. 2003. Help seeking and help design in interactive learning environments. *Review of Educational Research* 73(3):277–320.

Aleven, V.; McLaren, B.; Roll, I.; and Koedinger, K. 2004. Toward tutoring help seeking. In *Intelligent Tutoring Systems*, 227–239. Springer.

Baker, R. S.; Corbett, A. T.; Koedinger, K. R.; Evenson, S.; Roll, I.; Wagner, A. Z.; Naim, M.; Raspat, J.; Baker, D. J.; and Beck, J. E. 2006. Adapting to when students game an intelligent tutoring system. In *Intelligent tutoring systems*, 392–401. Springer.

Baker, R. S.; Corbett, A. T.; Roll, I.; and Koedinger, K. R. 2008. Developing a generalizable detector of when students

game the system. *User Modeling and User-Adapted Interaction* 18(3):287–314.

Bloom, B. S. 1984. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher* 13(6):4–16.

Leyzberg, D.; Spaulding, S.; Toneva, M.; and Scassellati, B. 2010. The physical presence of a robot tutor increases cognitive learning gains. In *34th Annual Conference of the Cognitive Science Society (Cog. Sci.'12)*.

Leyzberg, D.; Spaulding, S.; and Scassellati, B. 2014. Personalizing robot tutors to individuals' learning differences. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-robot Interaction, HRI '14*, 423–430. New York, NY, USA: ACM.

Merrill, D. C.; Reiser, B. J.; Ranney, M.; and Trafton, J. G. 1992. Effective tutoring techniques: A comparison of human tutors and intelligent tutoring systems. *The Journal of the Learning Sciences* 2(3):277–305.

Pekrun, R.; Goetz, T.; Titz, W.; and Perry, R. P. 2002. Academic emotions in students' self-regulated learning and achievement: A program of qualitative and quantitative research. *Educational psychologist* 37(2):91–105.

Ryan, R. M., and Connell, J. P. 1989. Perceived locus of causality and internalization: examining reasons for acting in two domains. *Journal of personality and social psychology* 57(5):749.

VanLehn, K. 2011. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist* 46(4):197–221.