

Fostering Learning Gains Through Personalized Robot-Child Tutoring Interactions

Aditi Ramachandran
Yale University
51 Prospect St.
New Haven, CT
aditi.ramachandran@yale.edu

Brian Scassellati
Yale University
51 Prospect St.
New Haven, CT
brian.scassellati@yale.edu

ABSTRACT

Social robots can be used to tutor children in one-on-one interactions. It would be most beneficial for these robots to adapt their behavior to suit the individual learning needs of children. Each child is different; they learn at their own pace and respond better to certain types of feedback and exercises. Furthermore, being able to detect various affective signals during an interaction with a social robot would allow the robot to adaptively change its behavior to counter negative affective states that occur during learning, such as confusion or boredom. This type of adaptive behavior based on perceived signals from the child (such as facial expressions, body posture, etc.) will create more effective tutoring interactions between the robot and child. We propose that a robotic tutoring system that can leverage both affective signals as well as progress through a learning task will lead to greater engagement and learning gains from the child in a one-on-one tutoring interaction.

Categories and Subject Descriptors

I.2.9 [Artificial Intelligence]: Robotics

Keywords

Human-robot interaction, adaptive systems, tutoring

1. BACKGROUND AND MOTIVATION

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 [2, 7]. During tutoring, the teacher has the ability to tailor the instruction to the individual learner, creating a personalized learning environment for each student. Due to research that shows that the physical presence of a robot tutor can increase cognitive learning gains, social robots are a natural option to explore when searching for methods of



Figure 1: Robot-Child Tutoring Interaction

instruction that may emulate the benefits of one-on-one human tutoring [5]. While there are many aspects of a one-on-one tutoring interaction that could be personalized, we are interested in customizing the pace of the interaction. Questions that are too easy for a student may cause boredom, while questions that are too difficult may lead to confusion. Both are negative affective states that could lead to disengagement [3]. Simultaneously, the student should still be making satisfactory progress through the learning goals of a tutoring interaction. We can see here that there are multiple aspects of a tutoring interaction that need to be monitored by the tutor.

In order to foster long-term learning gains, we are equally 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 [6]. Therefore, it is crucial that these social robot tutors simultaneously monitor both affective and cognitive states to maximize the potential for the child to learn.

Expert human tutors are able to fluidly transition between addressing a student's affective state and cognitive state when necessary through the course of a tutoring interaction [4]. In our search for tutoring agents that can mirror these expert human tutors, we aim to create sophisticated robotic tutors that can adaptively control the pace of the interaction by using both learning gains and affective feedback from the user, ultimately creating an engaging environment that increases long-term learning gains.

2. MODELING APPROACH

Within a tutoring interaction, we ask the following questions: Given a specific student's history, should the next question the robot provides be harder, easier, or of the same difficulty? Does the student need a break? The robot would

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). Copyright is held by the author/owner(s).

HRI'15 Extended Abstracts, March 2–5, 2015, Portland, OR, USA.

ACM 978-1-4503-3318-4/15/03.

<http://dx.doi.org/10.1145/2701973.2702721>.

adaptively choose which questions to present to the student based on the relative difficulties of the questions. We would like to frame this as a contextual bandit problem; given some state, we want to learn the actions that have the greatest reward over time [1]. We take the following approach:

For $t = 1, \dots, T$:

1. Given state $x_t \in X$
2. Choose 1 of K actions:
 $a_t \in 1, \dots, K$
3. Receive reward $r_t(x_t, a_t) \in [0, 1]$

In this approach, state refers to aspects of the interaction such as the percentage of questions answered correctly, the number of consecutive questions answered correctly, and the difficulty level of the current question. Examples of actions to choose from include presenting a harder question, an easier question, or an equally difficult question as compared to the previous question, providing praise, or taking a small break, which could involve a fun activity. Lastly, the reward $r_t(x_t, a_t)$ depends on both when the student answers the question correctly and whether the student is engaged. One of the existing challenges of this approach that we are working on is figuring out exactly how multiple reward signals (dependent on both learning progress and affective state) should be structured. The following are examples of actions that could be learned for an individual through this approach in a tutoring interaction:

- Answered many consecutive questions correctly: present harder question (boredom detected in reward)
- Answered similar questions correctly: present harder question (engagement detected in reward)
- Answered consecutive questions incorrectly: present easier question (engagement detected in reward)
- Answered similar questions incorrectly, frustration detected: take break (frustration detected in reward)

3. PROJECT COMPONENTS

While an integrated system involving a social robot that can adapt to an individual based on its perception of a user's affect is the longer-term goal for this project, we must first tackle some of the individual components of such a system. Leveraging tools to automatically extract various features from a child interacting with a social robot in real-time will be a crucial component to this line of research.

3.1 Detecting user affect

Specific nonverbal behaviors (e.g. smiling and posture change) are often more informative than prototypical affective states or self-reported assessments of affective state. In order to detect whether a child is exhibiting learning-centric affective states such as boredom or confusion, we first need to identify what type of information can be reliably detected. From a prior HRI study done with children in our research group, we have a set of 40 children interacting with social robots. We are using this data corpus to systematically evaluate various commercial software suites that perform affect detection to discern which will be most effective to use with children in real-time. While many of these commercial tools detect general affective states such as engagement, or valence, we also want to monitor individual physiological signals that likely correspond to affect, such as smiling, eyebrow raising and furrowing, and change in body posture. We are currently building tools and testing the feasibility of vari-

ous sensors (e.g. Kinect version 2) to deliver this type of information to a robot in real-time.

3.2 Proposed Study

Our goal is to gradually learn actions for the robot that best suit learners over time using the described approach. Once we can identify relevant affective states in the user as well as monitor a student's progress through the learning goals, we also need to be able to evaluate whether an adaptively-paced robot tutor is effective. We propose a study in which two groups of children (fifth graders) interact with a social robot in a one-on-one educational interaction, practicing math problems with the robot. The control group will have users receiving a random ordering of questions and breaks, and the second group will receive an adaptive pace based on both pedagogical progress and perceived engagement level. We hypothesize that the group of students receiving the adaptive tutoring from the robot will show larger learning gains as compared to the control group.

This study will allow us to evaluate the effectiveness and feasibility of creating such an adaptive system. Further analysis of this data corpus will elucidate what types of feedback signals happen at specific points during the tutoring interaction for students in various states to inform our use of reward signals in the modeling approach described above. Lastly, we will rely on surveys and questionnaires to understand how users perceived the changing behavior of the robot and how this may have impacted their tutoring experience overall. What we learn from analysis of this initial study will further inform how we balance multiple reward signals in our proposed adaptive model in future studies that focus on personalizing individual tutoring interactions over a longer period of time.

4. REFERENCES

- [1] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- [2] B. S. Bloom. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6):4–16, 1984.
- [3] S. D'Mello and A. Graesser. Dynamics of affective states during complex learning. *Learning and Instruction*, 22(2):145–157, 2012.
- [4] M. R. Lepper and M. Woolverton. The wisdom of practice: Lessons learned from the study of highly effective tutors. *Improving academic achievement: Impact of psychological factors on education*, pages 135–158, 2002.
- [5] D. Leyzberg, S. Spaulding, M. Toneva, and B. Scassellati. The physical presence of a robot tutor increases cognitive learning gains. In *34th Annual Conference of the Cognitive Science Society (Cog. Sci. '12)*, 2010.
- [6] R. Pekrun, T. Goetz, W. Titz, and R. P. Perry. Academic emotions in students' self-regulated learning and achievement: A program of qualitative and quantitative research. *Educational psychologist*, 37(2):91–105, 2002.
- [7] K. VanLehn. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4):197–221, 2011.