

Personalized Robot Tutors that Learn from Multimodal Data

Doctoral Consortium

Samuel Spaulding
MIT Media Lab
Cambridge, Massachusetts

ABSTRACT

As the cost of sensors decreases and ability to model and learn from multi-modal data increases, researchers are exploring how to use the unique qualities of physically embodied robots to help engage students and promote learning. These robots are designed to emulate the emotive, perceptual, and empathic abilities of human teachers, and are capable of replicating some of the benefits of one-on-one tutoring from human teachers. My thesis research focuses on developing methods for robots to analyze and integrate multimodal data including speech, facial expressions, and task performance to build rich models of the user's knowledge and preferences. These student models are then used to provide personalized educational experiences, such as optimal curricular sequencing, or leaning preferences for educational style. In this abstract, we summarize past projects in this area and discuss applications such as learning from affective signals and model transfer across tasks.

KEYWORDS

Human-Robot Interaction; Multimodal Interaction; Social Robotics

ACM Reference Format:

Samuel Spaulding. 2018. Personalized Robot Tutors that Learn from Multimodal Data. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10-15, 2018*, IFAAMAS, 3 pages.

1 INTRODUCTION

My thesis research focuses on designing, building, and evaluating adaptive robot learning companions. These robots interact with diverse human populations, such as young children, and can give students one-on-one practice and feedback in situations where a teacher may not be available. By collecting and analyzing data from interactions with individual students, such robots can become deeply personalized over time, constructing models of an individual's curricular strengths and affective mannerisms, then use these models to develop individually optimal strategies for educational interaction, such as when to introduce or review content or provide emotional support and encouragement [4, 12].

Interactive, educational robotic learning companions could combine the scale and precision of educational software with physical embodiment and sensing, which have been shown to positively impact learning, compliance, and long-term relationships, compared with virtually-embodied agents [1, 2, 9]. In the following sections I will review some important capabilities for adaptive robot learning companions, such as the ability to understand and model human affect, the ability to model individual students' knowledge and

task preferences, and the ability to transfer models from previous interactions to new tasks and scenarios.

2 AFFECT-AWARE ROBOT TUTORS

In order to reach their full potential as learning partners, social robots must be able to "close the affective loop". The affective loop, as defined by Höök and expanded upon by Paiva et al. [5, 11], is a process through which social agents perceive user's affective states, select an appropriate action, observe the effect of the chosen action in their partner's affective response, and learn from and refine their interaction models accordingly. This type of learned socio-emotional behavior is critical to developing agents that exhibit the social and emotional intelligence demonstrated by the best human tutors and is essential to advanced social interaction. Social robots that exhibit appropriate affective responses have been shown to elicit more fluent and prolonged interaction with humans [7] and are less susceptible to novelty effect that has limited long-term human-robot interaction studies [4].

Previous research has shown that tutoring interactions with an embodied social robot lead to improved learning gains, compared to screen-only representations [9]. In my Master's thesis, I demonstrated that embodied social robots effect greater emotional expressivity during educational games with children and that affective data from the interaction can be used to train models that more accurately track student knowledge than models which rely solely on student response data [13]. These results demonstrate the unique value of affective data captured from face-to-face interactions with social robot. This type of data is difficult to collect at scale, but is a crucial resource for developing systems that operate flexibly and empathically with humans in everyday life. As sensing technologies become more widespread and accessible, integrating this data into student models will lead to more natural and effective digital tutoring agents.

Human affect is expressed and sensed across many modalities, including speech, facial expressions, body pose, and gesture. While the performance of low-level classifiers for each modality has improved substantially over the past years [3, 6, 10], the next challenge is for socially interactive agents to understand and interpret these multimodal labels in the context of an interaction. Most robots that respond to human affect do so in a scripted way (e.g. if a smile is detected or the sentiment of a statement exceeds some threshold, execute a pre-defined sequence of behaviors). Due to the difficulties of collecting large amounts of physically sensed data, even those robots that learn affective behavioral models typically learn *generalized* models of affect, based on data collected from interactions across many individuals. If the goal is to learn models that perform well across a wide variety of individuals (as is common for many human-robot collaboration tasks), this is likely the best strategy.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10-15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

But people’s patterns of emotional expression are idiosyncratic. What constitutes an appropriate affective response is highly dependent on situational context and, perhaps more importantly, the interpersonal context between the interactors.

For tasks with a high degree of social interaction such as tutoring or negotiation, personalized affective strategies are crucial to the success of the interacting agent. Therefore, rather than learning a general model of human affect, it may be more fruitful to learn personalized models of affective expression.

3 PERSONALIZED MODELS FOR ROBOT LEARNING COMPANIONS

Good human teachers recognize the importance of personalization across many aspects of tutoring. For autonomous tutoring agents, this personalization often involves personalizing the sequence of curricular content presented to the student to improve learning gains [8]. Typically, this involves tracking student responses and other interactive features to model a student’s knowledge (i.e., learning a student model) and using that student model as an input to select content from the curriculum (i.e. a task model).

For example, I developed a robot system that uses Gaussian Process Regression to quickly build a model of children’s pronunciation from exemplars during gameplay. Based on the learned student model, the robot uses active learning to select words for subsequent rounds of the game. Because Gaussian Processes provide an estimate of the probability of mastery as well as the uncertainty around that estimate, the robot uses the student model as a basis for active learning, and can introduce the word the student is most likely to have not mastered (i.e., optimal for the robot to teach) as well as the word which has the greatest uncertainty in the model (i.e. provides optimal information gain for the robot to learn about the student) [12]. The model quickly learns from a few exemplars and can efficiently estimate a child’s pronunciation ability across a large curriculum of words without prompting the child on each word (as is common in many vocabulary assessments).

There are myriad ways to personalize an educational interaction in addition to curricular personalization. In many interactions the robot and the child are cast as peers rather than a traditional student-teacher relationship (e.g., [4], among others). In these scenarios, the robot has an important role to play by offering appropriate affective responses to the events of the tutoring interaction. Children may prefer different styles of support. Some may be inspired by grand long-term goals, others may prefer incremental encouragement and focus on immediate goals. Similarly, when the robot takes the role of a peer for interactive educational play, it can personalize its play style to an individual. Some children may be motivated by competition, while others may prefer collaboration and teamwork. A robot capable of learning these preferences for individuals and adapting accordingly may be a much more effective tutor.

4 LONG-TERM TEAMING AND TASK TRANSFER

In order to make good on the promise of personalized educational robots, such systems must be able to learn personalized models from personalized data. Long-term interaction with the same individual

over several sessions and tasks is the most straightforward approach to developing such models.

Personalization over single interactions limits the amount and variety of data that can be collected. Thus it is more akin to short-term adaptation, in which the agent learns some regularity in a specific user’s behavior during the course of the interaction. One of the challenges of human-robot interaction is the flexibility and variety of human activity. Data and interaction models, even from the same individual, can change significantly from day to day. A single, short-term interaction will likely not permit learning a representation of the user that captures the richness required for a successful long-term partnership. Even as research on personalized robot tutors is beginning to expand beyond single-session studies, virtually all of the work is conducted on a single task.

In my thesis research, I am exploring methods to transfer models of users across educational tasks, increasing the variety of interaction data the system encounters while preserving the model learning benefits of long-term interaction. Each task has a different interaction paradigm and addresses a different learning goal of a common curriculum; curricular similarities between tasks ensure that data from previous interactions with the same individual can inform the robot’s behavior during the new task.

5 CONCLUSION

My thesis focuses on designing and developing fundamental technologies for autonomous robots capable of engaging learners in educationally productive interactions. In this abstract, I have outlined three important capabilities my thesis research addresses: understanding and modeling human multimodal data (i.e. affect-awareness), personalization of interactive user and task models, and model transfer across tasks to sustain engagement and learn rich user models over long-term interactions.

For primarily interactive tasks like tutoring, understanding human affective behavior is critical to the success of the task. Insight from human tutors and empirical studies confirm that personalized models can yield better results compared to non-personalized models. Long-term interaction across a variety of tasks permits the construction of deeply personalized models to enable long-term teaming. As robust robotics research platforms continue to emerge into people’s homes, research to develop educational, affect-aware robots that personalize over long periods of time is an increasingly important, impactful, and insightful challenge.

REFERENCES

- [1] Wilma A Bainbridge, Justin W Hart, Elizabeth S Kim, and Brian Scassellati. 2011. The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics* 3, 1 (2011), 41–52.
- [2] Cynthia Breazeal. 2004. Social interactions in HRI: the robot view. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* 34, 2 (2004), 181–186.
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2017. Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. In *CVPR*.
- [4] Goren Gordon, Samuel Spaulding, Jacqueline Kory Westlund, Jin Joo Lee, Luke Plummer, Marayna Martinez, Madhurima Das, and Cynthia Breazeal. 2016. Affective Personalization of a Social Robot Tutor for Children’s Second Language Skills. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- [5] Kristina Höök. 2009. Affective loop experiences: designing for interactional embodiment. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364, 1535 (2009), 3585–3595.
- [6] SpeechAce Inc. 2017. SpeechAce. (November 2017). [Online; Accessed November 1, 2015].

- [7] Dan Leyzberg, Eleanor Avrunin, Jenny Liu, and Brian Scassellati. 2011. Robots that express emotion elicit better human teaching. In *Proceedings of the 6th international conference on Human-robot interaction*. ACM, 347–354.
- [8] Daniel Leyzberg, Samuel Spaulding, and Brian Scassellati. 2014. Personalizing robot tutors to individuals’ learning differences. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*. ACM, 423–430.
- [9] D Leyzberg, S Spaulding, M Toneva, and B Scassellati. 2012. The physical presence of a robot tutor increases cognitive learning gains. In *Proceedings of the 34th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- [10] Daniel McDuff, Abdelrahman Mahmoud, Mohammad Mavadati, May Amr, Jay Turcot, and Rana el Kaliouby. 2016. AFFDEX SDK: a cross-platform real-time multi-face expression recognition toolkit. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 3723–3726.
- [11] Ana Paiva, Iolanda Leite, and Tiago Ribeiro. [n. d.]. Emotion modeling for social robots. *The Oxford handbook of affective computing* ([n. d.]), 296–308.
- [12] Samuel Spaulding, Huili Chen, Safinah Ali, Mike Kulinski, and Cynthia Breazeal. 2018. A Social Robot System for Modeling Children’s Word Pronunciation. In *Proceedings of the 2018 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems.
- [13] Samuel Spaulding, Goren Gordon, and Cynthia Breazeal. 2016. Affect-aware student models for robot tutors. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 864–872.