# Interactive Robot Task Learning

**Andrea L. Thomaz**
Georgia Institute of Technology
801 Atlantic Dr
Atlanta, GA 30332 USA
athomaz@cc.gatech.edu

**Maya Cakmak**
Georgia Institute of Technology
801 Atlantic Dr
Atlanta, GA 30332 USA
maya@cc.gatech.edu

**Crystal Chao**
Georgia Institute of Technology
801 Atlantic Dr
Atlanta, GA 30332 USA
cchao@gatech.edu

**Nicholas DePalma**
Georgia Institute of Technology
801 Atlantic Dr
Atlanta, GA 30332 USA
ndepalma@cc.gatech.edu

**Michael Gielniak**
Georgia Institute of Technology
801 Atlantic Dr
Atlanta, GA 30332 USA
mgielniak3@gatech.edu


*CHI 2010*, April 10–15, 2010, Atlanta, Georgia, USA.
ACM 978-1-60558-930-5/10/04.

## Abstract
In this paper we provide a brief overview of our research agenda in Human-Robot Interaction and Interactive Learning. We highlight key components to be demonstrated as part of the CHI 2010 Media Showcase.

## Keywords
Human-Robot Interaction, Socially Guided Machine Learning

## ACM Classification Keywords
I.2.9: Robotics – Operator interfaces, I.2.6: Learning – Concept Learning.

## General Terms
Design, Algorithms, Human Factors.

## Introduction
There is currently a surge of interest in having robots leave the labs and factory floors to help solve critical issues facing our society, ranging from eldercare to education [1]. We have many problems to solve before general purpose robots can function in, inherently social, dynamic human environments. A critical issue is that we will not be able to pre-program robots with every skill they will need to play a useful role in society; robots will need the ability to interact and learn new things 'on the job.'

The goal of our research is to enable robots to learn new tasks and skills from everyday people. We focus on the key point that the robot learning by demonstration problem takes place within a social structure that can guide and constrain the learning problem. We believe that addressing this point will be essential for developing systems that can learn from everyday people that are not experts in Machine Learning or Robotics.

## Socially Guided Machine Learning

For years researchers working on robotic and software agents have been inspired by the idea of efficiently transferring knowledge about tasks or skills from a human to a machine. The primary motivation for leveraging human input is often to achieve some learning performance gains for the machine. Alternatively, in Socially Guided Machine Learning (SG-ML), we advocate designing for the performance of the complete, coupled human-machine teaching-learning system [5]. This perspective reframes the Machine Learning problem as an interaction between the human and the machine. This allows us to take advantage of human teaching behavior to construct a machine learning process that is more amenable to the human partner.

*Input Channels*

An SG-ML approach begins with the question: "How do humans want to teach?" In addition to designing the interaction based on what the machine needs to succeed in learning, we need to also understand what kinds of intentions humans will try to communicate in their everyday teaching behavior. We can then change the input portion of the machine learning training process to better accommodate a human partner.

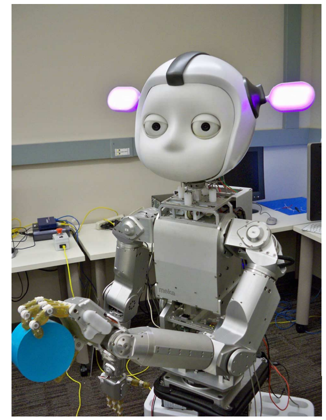

**Figure 1.** The Simon robot platform has two 7-degree of freedom (DOF) arms with 4-DOF hands, and a 13-DOF socially expressive head.

*Output Channels*

An SG-ML approach asks: "How can the output provided by the learning agent improve the performance of the teaching-learning system?" In a tightly coupled interaction, a 'black box' learning process does nothing to improve the quality and relevance of the instructional guidance; however, transparency of the internal state of the machine could greatly improve the learning experience. By communicating its internal state, revealing what is known and what is unclear, the agent can guide the teaching process.

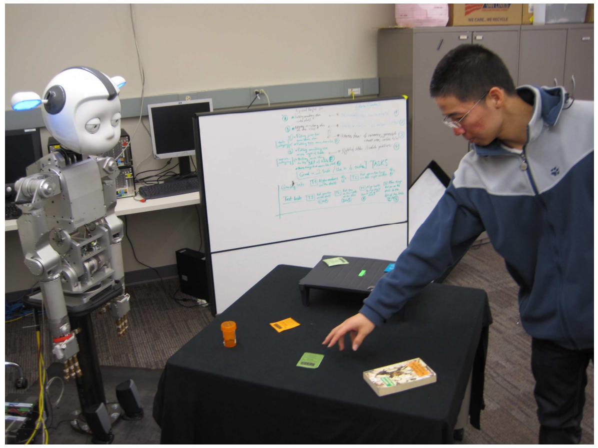

**Figure 2.** Simon pictured with a human partner in the tabletop workspace.

*Input/Output Dynamics*

We also recognize that these input and output channels interact over time, the dynamics of which can change the nature of the input from the human. In particular, the temporal structure of teaching versus performing may significantly influence the behavior of the human. An incremental, on-line learning system creates a very different experience for the human than a system that must receive a full set of training examples before its performance can be evaluated. Iterative feedback allows for on-line refinement; the human can provide another example or correct mistakes right away instead of waiting to evaluate the results at the end of the training process. This may provide a significant benefit to the human's level of engagement and motivation. The sense that progress is being made may keep the human engaged with the training process for a longer period of time, which in turn benefits the learner.

## Simon Robot Platform

The robotic platform for this research is "Simon," an upper-torso humanoid social robot with 7-DOF arms, 4-DOF hands, and a socially expressive head and neck (see Fig. 1 & 2). We are developing the Simon platform specifically for face-to-face human-robot interaction. In our task scenarios, the robot works alongside a human partner at a tabletop workspace. The robot communicates with human partners through both gesture and speech.

Our interactive learning system is implemented within the C6 software system (a derivative of [2]), which follows a traditional sense-think-act architecture and has a specific pipeline for triggering robot actuations from sensory inputs. Objects lying on a table in Simon's workspace are detected through a fixed overhead camera. A number of simple features are computed for each segmented object including (i) height and with of the best fitting ellipse, (ii) the number of corners and area of the simplified polygon contour of the segmented blob and (iii) hue histogram over the segmented blob. The workspace involves various objects people can use to teach tasks to Simon. The robot can learn via Supervised and Active Learning methods [3,4].

## Demonstration Description

In our CHI 2010 demonstration, people will interact with the robot one at a time to teach Simon new tasks with the available objects in the workspace. For example, a person could teach Simon to put objects of a certain type in a particular location (e.g., put the books on the shelf).

The robot learns such tasks over a few demonstrations. These demonstrations can be provided by showing Simon examples of completing the task, or by instructing Simon via speech input to do the appropriate actions to complete the task. Over a few examples, Simon generalizes a representation of the task "goal" and is then able to perform the learned task in the workspace.

Simon exhibits a variety of social cues to facilitate a natural teaching/learning process (e.g., eye gaze and other communicative gestures). Additionally, Simon has the ability to use a simple speech vocabulary to query the human teacher for more information about a task: asking for particular examples, or changes to the current example.

## Conclusion

Our overall research agenda is Socially Guided Machine Learning, investigating ways to enable everyday people to teach new tasks to robots. In this paper we have highlighted key components to be demonstrated as part of the CHI 2010 Media Showcase.

## References Cited

[1] CCC/CRA Robotics Roadmap: http://www.us-robotics.us/

[2] Blumberg, B., Downie, M., Ivanov, Y., Berlin, M., Johnson, M., and Tomlinson, B. Integrated learning for interactive synthetic characters. *In Proc. of the ACM SIGGRAPH,* 2002.

[3] C. Chao, M. Cakmak, and A. L. Thomaz. Transparent active learning for robots. *In Proc. of the International Conference on Human-Robot Interaction (HRI)*, 2010.

[4] C. Chao, M. Cakmak, and A. L. Thomaz. Interactive task learning with discrete and continuous features. *In Proc. of the Robot Workshop at American Association of Artificial Intelligence (AAAI),* 2009.

[5] Thomaz, A.L. Socially Guided Machine Learning. *MIT PhD Thesis,* 2006.