Robotics Program: Enabling Intelligence through Middleware: Papers from the 2010 AAAI Workshop (WS-10-09)

Interactive Learning by Demonstration with the Simon Robot

Crystal Chao, Michael Gielniak, Jae Wook Yoo, and **Andrea L. Thomaz**
School of Interactive Computing
Georgia Institute of Technology
Atlanta, Georgia 30332, USA
{cchao, mgielniak3, jyoo}@gatech.edu, athomaz@cc.gatech.edu

Abstract

In this paper we highlight the components of our entry in the AAAI 2010 Learning by Demonstration (LbD) Challenge. Our overall research agenda is Socially Guided Machine Learning, investigating ways to enable everyday people to teach new tasks to robots. Thus, our focus on the LbD problem is that of making it socially interactive. The challenge task was block sorting. Our robot Simon was on display for three days and learned several kinds of sorting tasks from human partners.

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. 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 tasks and skills "on the job." This is the motivation for Learning by Demonstration.

Our approach to the problem of Learning by Demonstration focuses specifically on goal learning. Humans have a propensity to interpret actions based on goals and intentions rather than motion trajectories or literal configurations of the world (Woodward, Sommerville, and Guajardo 2001; Baldwin and Baird 2001). The ability to infer goals forms a common ground during interactions and is a basis for human communication. This principle is important for robots that learn from human teachers, as they need to assume that their human partners are goal-oriented. The robots task is to infer the underlying goal that the teacher is communicating. We aim to build robots that similarly frame the learning problem as goal inference so that they are natural to teach and meet people's expectations for a learning partner. In this paper we present an overview of our entry in the AAAI 2010 Learning by Demonstration Challenge event.

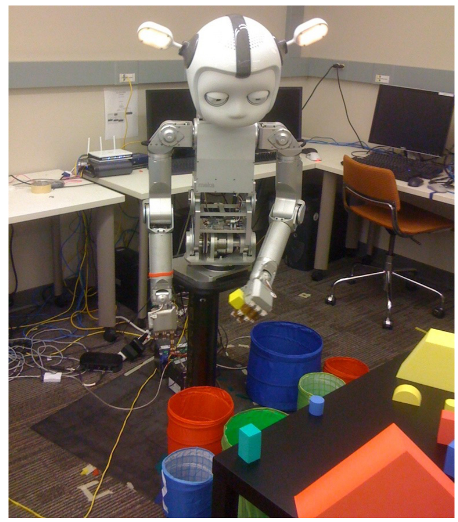

Figure 1: Simon learns to sort the Learning by Demonstration Challenge objects into six different bin locations.

Learning by Demonstration Challenge

The Learning by Demonstration Challenge task was to sort colored foam blocks. We augmented the task description by considering different sized blocks, as well as target locations with both color and size features. Thus Simon could learn to sort the blocks into the bins however the human teacher liked. For example, learning that all green blocks go in the large blue bin, or that small blocks go in large bins, etc. In this section we describe details of implementation.

Platform

The robotic platform used for the challenge was "Simon," an upper-torso humanoid social robot with 7-DOF arms, 4-DOF hands, and a socially expressive head and neck (Fig. 1). For safety in human-robot interaction, the robot has series elastic actuators at every joint in both arm chains. For vision, Simon has a Point Grey Firefly camera in each eye. Our learning system is implemented within the C6 software system, which provides an architecture for managing perceptual data and actions (Blumberg et al. 2002).

Table 1: Learning task features

Feature	Values
Object Color	Red, Green, Blue, Yellow
Object Size	Large, Small
Target Bin Color	Red, Green, Blue
Target Bin Size	Large, Small

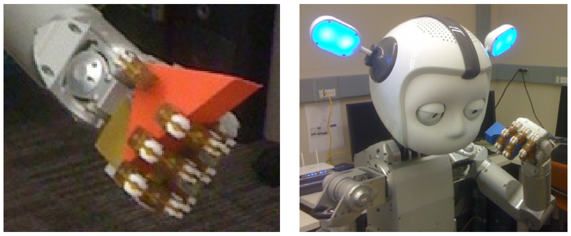

(a) Size: Position feedback during grasp. (b) Color: Closer inspection before eye camera.

Figure 2: Perception of features for learning task.

Feature Space and Perception

The features considered for the task were color and size for objects and target bin locations. The discrete values for these features are shown in Table 1.

We placed six different target bins in Simon's workspace, which represented all combinations of possible color and size features. Simon perceived bin features and locations using a routine that allowed him to scan the ground for colored blobs. Size was determined by comparing the relative sizes of the blobs. This would update his internal representation of which bins were at which location. The routine could be triggered by telling him, "Simon, I moved the bins."

During the interaction, the human teacher holds an object out to Simon and says, "Take this." This triggers Simon to reach out his arm. The robot power grasps the object when it detects torque changes of the joints in his arm chain, indicating that the teacher has placed an object in his hand. Simon perceives the object features after he has a firm grasp. Size is determined from the position feedback in the thumb and the other finger joints, shown in Figure 2(a); the more extended his fingers are, the likelier it is that the object is large. Simon then waves the object in front of one eye camera, shown in Figure 2(b). Optical flow is used to segment the region of moving pixels, which is used to determine the color. Once Simon determines the color of the object, he changes his ear LED color matching the detected color as feedback for the human teacher.

Supervised Learning

In the supervised learning mode, Simon learns interactively with a human teacher from examples with positive or negative labels. This proceeds in the interaction as follows. After allowing Simon to determine the features of the object, a human teacher tells him, "This goes in bin n," where $n \in [1, 6]$. This provides a positive label to the example vector containing that particular combination of object and bin features.

Simon learns on the labeled 4-feature vectors described in the previous section by using them to form a hypothesis space consisting of clusters of such vectors. Each cluster is a group of vectors that covers the entire object space to represent the task of sorting all the objects. Learning occurs by refining the version space over these hypotheses. The target concept is the cluster that is most consistent with the data, which is the maximum likelihood hypothesis. More details about the learning system can be found in (Chao, Cakmak, and Thomaz 2010).

Active Learning

In Machine Learning theory, the problem of active learning is one of selecting the optimal or seminal examples to be labeled. In the context of a robot active learner, there are Human-Robot Interaction considerations as well. The question is not just what to query but when and how. In our demonstration, the teacher explicitly solicits queries from Simon by saying "Simon, do you have any Questions?" This method was found to be favorable in previous work (Cakmak, Chao, and Thomaz 2010). Simon then gazes at the objects on the table before him and points at most informative object among those present while asking, "Does this object go in bin n?" The human teacher can then provide a positive or negative label using speech.

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 that were demonstrated in the AAAI 2010 Learning by Demonstration Challenge. Simon was on display for three days in the conference exhibit hall and learned several kinds of sorting tasks from human partners.

References

Baldwin, D., and Baird, J. 2001. Discerning intentions in dynamic human action. *Trends in Cognitive Sciences* 5(4):171–178.

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

Cakmak, M.; Chao, C.; and Thomaz, A. L. 2010. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development* 2(2):108–118.

Chao, C.; Cakmak, M.; and Thomaz, A. L. 2010. Transparent active learning for robots. In *Proceedings of the 2010 ACM Conference on Human-Robot Interaction (HRI)*.

Woodward, A. L.; Sommerville, J. A.; and Guajardo, J. J. 2001. How infants make sense of intentional actions. In Malle, B.; Moses, L.; and Baldwin, D., eds., *Intentions and Intentionality: Foundations of Social Cognition*. Cambridge, MA: MIT Press. chapter 7, 149–169.