Special Issue on Robot Learning from Demonstration

Sonia Chernova · Andrea L. Thomaz

Published online: 8 August 2012
© Springer Science & Business Media BV 2012

Algorithms for robot Learning from Demonstration (LfD) seek to enable human users to expand the capabilities of robots through interactive teaching instead of explicit programming. This special issue includes a collection of articles that span some of the issues and challenges that arise when LfD is done in the context of a social interaction with a human partner.

The importance of designing algorithms and interactions with non-expert end-users in the loop is highlighted in the article by Suay et al. In one of the first comparative user studies of its kind, non-experts used and evaluated three different learning by demonstration algorithms/interactions in the same domain. Their results point out the non-trivial challenges that arise in getting learning input from a non-expert human teacher.

One common problem domain tackled with Learning from Demonstration is that of learning low-level motion control policies to achieve some particular task or skill. Several of the articles in this special issue are working in this domain, touching on various important questions that come to bare when this type of LfD is done in a social context with a human teacher.

In their article, Grollman et al. address an important topic for social LfD, learning from failed demonstrations. When demonstrations are to come from naïve non-expert humans,

it can be expected that not every demonstration will be successful. Their experiments show that failure data does have useful information for the learner, and can be used to discover ways to successfully achieve the target skill.

Akgun et al. similarly assume that the teacher is a naïve user, and present a unified LfD framework that supports three types of motion demonstration inputs—trajectory, keyframe and hybrid—allowing users to have maximal flexibility in teaching new skills. The paper presents techniques for converting all demonstration inputs into a keyframe representation, and the resulting Keyframe-based LfD system is shown to perform on par with state of the art methods while significantly broadening the range of user interactions.

Dong and Williams study an alternate variant of the above problem, in which the robot learns a library of activities from user demonstrations and uses it to recognize an action performed by an operator in real time. The authors contribute a novel probabilistic flow tube representation that can intuitively capture a wide range of motions, a method for identifying the relevant features of a motion, and ensuring that the learned representation preserves these features in new and unforeseen situations. Experimental results using a robotic arm show a 49 % improvement over prior work in recognition rate for varying environments.

The work of Mohammad and Nishida alternately assumes that the task expert is not explicitly aware of their role as a teacher at all. In their approach, the learner has to take the initiative and decide for itself which of the observed actions performed by the teacher should be learned. The authors present a novel fluid imitation engine that enables the learner to automatically segment and learn observed behaviors without the help of the demonstrator.

While still in the problem domain of learning low-level skills, three of the issue's articles look at mechanisms for allowing humans to provide some high-level corrective feedback or advice to the learner in addition to demonstrations.

S. Chernova (✉)
Robotics Engineering Program, Worcester Polytechnic Institute,
Atwater Kent Labs 125, 100 Institute Road, Worcester,
MA 01609, USA
e-mail: soniac@cs.wpi.edu

A.L. Thomaz
School of Interactive Computing, Georgia Institute
of Technology, 801 Atlantic Dr., Atlanta, GA 30332, USA
e-mail: athomaz@cc.gatech.edu

Argall et al. look at allowing a human operator to give advice and feedback on a learned policy during its execution using advice operators. This allows the operator to give corrections to state action pairs during execution. Their implementation and experiments in the domain of mobile robot motion control demonstrate the utility of learning with this kind of directed feedback versus simply giving more demonstrations of the task.

The work of Tommaso et al. seeks to improve shared task understanding between the robot and teacher by introducing an interactive active sensing interface for the assessment of robot skills acquisition. The presented interface provides a visually augmented operating space shared between the learner and the teacher by graphically superimposing task features onto the physical operating space. The resulting system enables the user to incrementally visualize and assess the learner's state and, at the same time, focus on the skill transfer without disrupting the continuity of the teaching interaction.

Knox et al. seek to understand how people naturally teach and to design learning algorithms that match, or even enhance, natural human tendencies. Their paper describes two experiments that examine how differing conditions affect a human teacher's feedback frequency and the computational agent's learned performance. The results indicate that the robot's action selection may not only be limited to exploration or policy exploitation, but could also be designed to indirectly, and subconsciously, alter the behavior of the human teacher.

In a different kind of high-level feedback, Mercili et al. look at the scenario of a human demonstrator providing feedback to a robot learning motion control tasks (e.g., walking). Their novel twist is in using the human to learn a policy of when to pay attention to what. Through interaction with the user, the agent learns the appropriate resolution of state to use at various points in the task in order to represent the optimal action in that state.

A second class of problem domain that is tackled by LfD researchers is learning high-level tasks as opposed to low-level motor skills. This approach is represented in the article

by Jakel et al., where they learn task models for dexterous manipulation tasks from observing human demonstrations. The learned task models use a feature space of automatically generated contact constraints which are used by a planning algorithm when it comes time to execute the learned task.

All of the above articles address the situation of a single user interacting with a single robot in a co-located space. However, the article by Osentoski et al. envisions a future of remote experimentation with robots learning from any human on the web. Their article details the design and implementation of the PR2 Remote Lab that was used in the AAAI 2011 Learning by Demonstration challenge. The experience of the teams that used the remote lab during this event points out the important challenge problems for the future of robot learning in "the cloud".

In conclusion, this special issue presents a broad cross-section of Learning from Demonstration research, particularly as it relates to social robotics. Presented approaches range from low level motion planning to high level tasks, and highlight the intricate role that human users play in this robot learning paradigm. The diversity of articles underscores the breadth of research questions posed by learning through social interaction with everyday humans, as well as the many challenges we face in broadening the range of interactions between robots and naïve users.

Sonia Chernova is an Assistant Professor at Worcester Polytechnic Institute (WPI), where she directs the Robot Autonomy and Interactive Learning group. Her research interests lie in interactive machine learning, adjustable autonomy, crowdsourcing and human-robot interaction. Her work focuses on the development algorithms that enable robots to learn through social interaction with humans. Prof. Chernova received her Ph.D. from Carnegie Mellon University in 2009 and completed a postdoc at the MIT Media Lab before joining WPI.

Andrea L. Thomaz is an Assistant Professor in the School of Interactive Computing, Georgia Institute of Technology. She received a B.Sc. degree from the University of Texas, Austin, in 1999, and M.Sc. and Ph.D. degrees from the Massachusetts Institute of Technology in 2002 and 2006. Thomaz directs the Socially Intelligent Machines Lab, where her work focuses on social robots and socially guided machine learning. She is published in Robotics, Human-Robot Interaction, and Artificial Intelligence.