Skill Demonstration Transfer for Learning from Demonstration

Tesca Fitzgerald, Andrea Thomaz
School of Interactive Computing
Georgia Institute of Technology
Atlanta, Georgia
{tesca.fitzgerald, athomaz}@cc.gatech.edu

ABSTRACT

Learning from Demonstration is an effective method for interactively teaching skills to a robot learner. However, a skill learned via demonstrations is often learned within a particular environment and uses a specific set of objects, and thus may not be immediately applicable for use in unfamiliar environments. Transfer learning addresses this problem by enabling a robot to apply learned skills to unfamiliar environments. We describe our ongoing work to develop a system which enables transfer learning by representing skill demonstrations according to the level of similarity between the source and target environments.

Categories and Subject Descriptors

I.2.9 [**Artificial Intelligence**]: Robotics

Keywords

Learning from demonstration; Transfer learning

1. INTRODUCTION

Learning from Demonstration (LfD) seeks to enable end-users to teach robots quickly and intuitively. Rather than "teach" a skill by programming the robot manually, LfD enables an end-user to provide a skill demonstration that guides the robot to complete the intended skill [2, 3]. However, these skill demonstrations are often provided using a limited set of objects, and are often expected to be repeated within the same environment. Suppose that a robot is taught to perform a *pouring* skill using a tablespoon object and is later asked to perform the pouring skill in an environment containing a mug; the robot may not succeed in performing the skill using the new object, despite continuing to have the same goal of *pouring*. As such, skill transfer may be necessary for robots to adapt skill demonstrations such that they can be repeated in unknown environments.

Without this transfer ability, a robot would require additional time and training to relearn each skill in the new environment. Leveraging previously learned skills to bootstrap the learning of new skills would significantly reduce the amount of time necessary to learn a new skill, and would require less effort of the skill teacher. Additionally, transfer would enable a robot to attempt to complete a task in an unknown environment, and then request additional information or demonstrations from the human teacher when necessary. We refer to the *source environment* as the original environment in which the skill demonstration was provided, and the *target environment* as the new scenario which the robot is expected to address.

Related work has addressed the problem of transfer learning via reinforcement learning [4, 7] and case-based reasoning [5] in domains such as RoboCup soccer. However, to our knowledge, there has been no work thus far that integrates LfD and transfer learning.

1.1 Problem Definition

Skill demonstrations can be provided using various means, including kinesthetic learning and teleoperation [2]. Our focus is on skill transfer for skills taught via kinesthetic learning, in which the teacher demonstrates the skill by physically moving the robot's arm [1]. Particularly, we address transfer problems in which few source skill demonstrations have been provided for the requested skill. We seek to address problems such as the following. A human teacher demonstrates a skill such as pouring a cup over a target or scooping coffee beans from one container to another. Meanwhile, the robot records the demonstrated trajectories and object locations. At a later time, the robot is asked to repeat the *pouring* task, but using a different set of objects to parameterize and execute the *pouring* skill than those used in the original demonstration. The robot then transfers its representation of the *pouring* demonstration to address the differences between the source and target environments. Finally, the transferred skill model is executed in the target environment.

The goal of our work is to enable this full interaction to occur. However, our main focus is on the *transfer* step of this process (Step 4), which requires the robot to (i) select the appropriate source skill demonstration that most resembles the requested skill, (ii) create a mapping between objects in the target environment and the selected source skill demonstration's environment, (iii) use this mapping to adapt the source demonstration to address changes in the target environment, and (iv) execute the adapted skill representation in the target environment. We approach this step by proposing a set of skill demonstration representations, described in Section 2, that enable mapping, transfer, and execution.

HRI'15 Extended Abstracts, March 2–5, 2015, Portland, OR, USA.
ACM 978-1-4503-3318-4/15/03.
http://dx.doi.org/10.1145/2701973.2702728.

(a) Source Environment (b) Similar Trajectory (c) Similar Sub-Tasks (d) Similar Goal (e) Dissimilar

Figure 1: Overhead Views of Source and Potential Target Environments

2. APPROACH

An approach to transfer learning must be able to account for a wide range of transfer problems, as Figure 1 illustrates. Figure 1a represents an overhead view of the environment in which a demonstration of a *scooping* skill was provided. The remaining images in Figure 1 each represent an overhead view of a potential transfer environment to which the robot is expected to transfer the skill demonstration. Figure 1b represents the "easiest" transfer problems in which there is little difference between the source and target environments. Figures 1c and 1d represent problems in which the target environment becomes progressively more dissimilar as objects are moved, replaced, and removed from the scene, until the target environment is too dissimilar to transfer the learned skill (Figure 1e). To accommodate this range of transfer problem difficulty, we approach transfer learning by defining a set of skill demonstration representations, each of which enable transfer at a different level of similarity.

To address the target environment shown in Figure 1b, we represent the skill demonstration as a set of Dynamic Movement Primitives (DMP) [6], each of which represents a subtask of the scooping task (e.g. grasping the scoop, scooping the pasta, moving the scoop to the target location). The representation is then transferred by adjusting the goal location of each DMP to account for changes in object locations. Our previous work has included a preliminary implementation of this demonstration representation for use in a pick-and-place type skill, and indicated that this representation can be used to address transfer problems in which object locations vary slightly from the source demonstration.

The target environment shown in Figure 1c must be addressed using the same sub-tasks as provided in the demonstration. However, the object identifiers used to determine which objects to target cannot be transferred. As a result, the demonstrated trajectories cannot be directly reused without first mapping the differences between the source and target environments, and then using this mapping to define the goal location of each DMP. Thus, a representation that addresses this type of transfer problem must encode the demonstrated sub-tasks and a mapping between objects in the source and target environments. Our implementation for this representation is currently in development.

Finally, the target environment shown in Figure 1d represents a transfer problem in which only the overall goal (pouring the contents of one container into another) is shared between the source and target cases. Thus, the trajectories demonstrated by the human teacher are not transferred to the target problem; rather, only the goal(s) of the teacher's actions is transferred. We address this type of transfer problem with a representation that encodes the demonstration as a goal end-effector position and velocity. We have completed a preliminary implementation of this representation for use in a pendulum-targeting skill, in which the robot is taught to move a pendulum at one of two demonstrated speeds. In this example, the robot's end-effector reached the pendulum at the demonstrated goal location and velocities, but by planning a trajectory that was dissimilar from the demonstrated trajectory. This illustrates how a goal-centric demonstration representation can be applied to transfer problems such as that shown in Figure 1d, where the goal cannot be achieved in the target environment using the demonstrated trajectory.

3. PROPOSED WORK

Current work is focused on implementing the second representation described in Section 2. Future work will integrate all three representations in a single system, such that the robot can autonomously select a single representation when attempting to solve a transfer problem. Additionally, future work may investigate whether the robot should further involve the human teacher by asking for assistance in solving more difficult transfer problems.

4. ACKNOWLEDGMENTS

We thank Ashok Goel for many discussions on defining the described skill representations. This work was supported by NSF Graduate Research Fellowship DGE-1148903.

5. REFERENCES

[1] B. Akgun, M. Cakmak, K. Jiang, and A. L. Thomaz. Keyframe-based learning from demonstration. *Int. Journal of Social Robotics*, 4(4):343–355, 2012.

[2] B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.

[3] C. G. Atkeson and S. Schaal. Robot learning from demonstration. In *ICML*, volume 97, pages 9–15, 1997.

[4] F. Fernández and M. Veloso. Policy reuse for transfer learning across tasks with different state and action spaces. In *ICML Workshop on Structural Knowledge Transfer for Machine Learning*, 2006.

[5] M. W. Floyd, B. Esfandiari, and K. Lam. A case-based reasoning approach to imitating robocup players. In *FLAIRS Conference*, pages 251–256, 2008.

[6] S. Schaal. Dynamic movement primitives-a framework for motor control in humans and humanoid robotics. In *Adaptive Motion of Animals and Machines*, pages 261–280. Springer, 2006.

[7] M. E. Taylor, S. Whiteson, and P. Stone. Transfer via inter-task mappings in policy search reinforcement learning. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, page 37. ACM, 2007.