A Similarity-Based Approach to Skill Transfer

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

ABSTRACT

We address the problem of transferring skills learned from interactive demonstrations. Learning from Demonstration is an effective method for quickly learning to complete a task; however, task demonstrations are often provided within a particular environment, and thus cannot be immediately repeated in unfamiliar environments that differ in object configurations and features. We describe our ongoing work in implementing a transfer processes which represents task demonstrations according to the level of similarity between the original environment and the new, unfamiliar environment.

1. INTRODUCTION

Learning from Demonstration is well-researched approach which aims to enable imitation by having the robot receive a demonstration of a human teacher completing a task [2, 3, 4]. This demonstration can be provided via several means such as *kinesthetic teaching*, in which the human teacher physically guides the robot's hand to complete the task [2, 1]. At a later time, the robot may be asked to repeat the demonstrated task in a new workspace which may vary in object configurations or features. Since demonstrations are provided manually, the robot may have been provided demonstrations within only a single environment (referred to as the *source* environment), and thus cannot immediately repeat the task in a new, *target*, environment.

We discuss our ongoing work to enable *transfer* for robots that learn tasks from kinesthetic demonstrations. Transfer learning has been addressed in related work via methods such as reinforcement learning [5, 8] and case-based reasoning [6] in domains such as RoboCup soccer. However, to our knowledge, there has been no prior work which integrates transfer learning and Learning from Demonstration.

Our eventual goal is to enable transfer processes such as the following. A human teacher guides the robot to complete a task such as scooping the contents of one container into another. During the demonstration, the robot records the demonstrated trajectories and object features. At a later time, the robot is asked to repeat the *scooping* task, but in a new, *target* environment. Thus, the robot must use a different set of object features to parameterize and execute the *scooping* task than those observed in the original, *source* environment. Next, the robot transfers its representation of the *scooping* task to accommodate for the differences between the source and target environments. The transferred task representation is then executed in the target environment.

2. APPROACH

We separate the task transfer process into two stages: (i) the task learning and representation stage, occurring when the robot receives a task demonstration, and (ii) the task transfer stage, occurring when the robot is later asked to repeat the task in the target environment.

2.1 Task Learning and Representation

The human teacher first interacts with the robot to demonstrate a *task*, which may consist of several *sub-skills* that are to be completed in sequence. As an example, the task of scooping pasta from one bowl into another consists of several sub-skills including moving the scoop to the pasta bowl, executing the "scoop" action, and so forth. The teacher demonstrates the full *task* trajectory, using voice commands to segment the task trajectory into a sequence of *sub-skills*. Meanwhile, the robot records and learns a skill model for each sub-skill trajectory. The robot stores this demonstration in memory as a *source demonstration* containing the record $R =< D,\ T,\ O,\ L >$, where:

- D represents a set of sub-skill models. Each sub-skill is represented as a Dynamic Movement Primitive [7], which allows the robot to later reproduce a motion trajectory that is similar to the original demonstration, but with modified starting and ending point locations.
- T is the set of object relations that express the end point location of each sub-skill in relation to the locations of objects in the robot's environment, e.g. the offset between the robot's hand and a bowl in its environment after completing the "scoop" sub-skill. This is defined as $T = <<_x t_0,\ _y t_0,\ _z t_0>,\ ...,\ <_x t_n,\ _y t_n,\ _z t_n>>$
- O is the set of objects observed in an overhead view of the robot's environment, defined as $O = <o_0,\ ...,\ o_i>$ where o_i lists a single object's ID.
- L is the set of object locations, and is defined as $L = <l_0,\ ...,\ l_i>$ where l_i contains the x, y, z coordinates of a single object.

2.2 Transfer and Execution

At a later time, the robot must be able to transfer the source demonstration to address a new, target environment. An approach to transfer must be able to account for a wide range of transfer problems, as Figure 1 illustrates. We take a similarity-based approach to transfer, where we consider the similarity between the source demonstration and target environments when defining transfer processes. Figure 1a

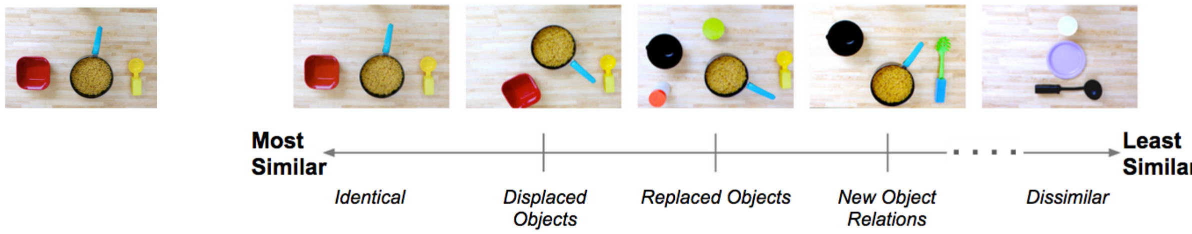

(a) Source Environment

(b) Potential Target Environments

Figure 1: Spectrum of Similarity Between Source and Target Environments

depicts an overhead view of a source environment in which the robot was trained to complete a scooping task. Figure 1b depicts a set of overhead views of potential target environments in which the robot may be expected to repeat this task, arranged left-to-right according to their decreasing similarity to the source environment. The two extremes of this spectrum represent environments that are identical and dissimilar to the source environment, respectively, and thus should not be addressed by adapting the source demonstration. The remaining three images represent more common and realistic transfer problems, each of which corresponds to a separate level of similarity. We address each class of transfer problems using a separate transfer method, each differing in the amount of information that can be transferred from the source demonstration:

- *Retargeting Transfer Approach*: The environment shown in the *displaced objects* scene in Figure 1b is very similar to the source environment shown in Figure 1a, differing only the location of each object. Thus, all elements of the source demonstration representation described in Section 2.1 can be transferred, except for the locations of objects in the target environment. Once these object locations are updated, the sub-skill models can be retargeted to account for the new object locations.

- *Mapping Transfer Approach*: The target environment depicted in the *replaced objects* scene can be addressed by transferring the sub-skill models and targeting relation elements of the source demonstration. However, the robot must additionally be provided with a mapping between objects in the source and target environments, which we currently provide manually.

- *Relational Transfer Approach*: Finally, the target environment shown in the *new object relations* scene can be addressed by transferring the same sub-skill models as in the source demonstration. However, by changing the size of the scoop, the relation between the robot's hand and objects in the environment must be adjusted such that the robot's hand is higher above the pasta bowl prior to scooping. Thus, to address this transfer problem, the robot must be provided with an updated list of object locations, a mapping between objects in the source and target environments, and a new set of targeting relations that redefine the relation between the robot's actions and the location of objects to account for the change in scoop size.

3. FUTURE WORK

We have implemented three approaches to transfer, each addressing transfer problems occurring at a different level of similarity. Preliminary experiments have evaluated each method under the assumption that we select the approach to be used for a given transfer problem. Future work will integrate all three methods of transfer, such that the robot can autonomously select the approach that best addresses a given transfer problem. Additionally, the current implementation assumes that we manually provide the robot with a mapping between objects that are equivalent between the source and target environments. We plan to identify a method for autonomously determining this object mapping.

4. ACKNOWLEDGMENTS

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] S. Chernova and A. L. Thomaz. Robot learning from human teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(3):1–121, 2014.

[5] 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.

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

[7] 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.

[8] 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.