Advances in Cognitive Systems (2015) 1-4 Submitted 04/2015; published 05/2015

Case-Based Skill Transfer in Robotic Agents

Tesca Fitzgerald TESCA.FITZGERALD@CC.GATECH.EDU
School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA 30332 USA

Abstract

Imitation learning is an effective method for interactively teaching a robot learner to complete a task. We address the problem of *transfer* for robotic agents that learn tasks from demonstrations, where a robot is asked to adapt a learned task to be repeated in a related, but unfamiliar, environment. We take a case-based approach to transfer, where the robot learner stores individual task demonstrations in a *case memory* such that they can be used at a later time for adaptation and reuse in new, target environments. We describe our ongoing work to enable transfer for robots that imitate task demonstrations.

1. Introduction

Imitation learning is a skill essential to human development and cognition (Tomasello, Kruger, & Ratner, 1993; Piaget & Cook, 1952). In imitating the actions of a teacher, a cognitive agent learns the demonstrated action such that it may perform a similar action later and achieve a similar goal. Thus, we expect that a cognitive system that learns from imitation would reuse what it has learned from one experience to reason about addressing related, but different, problem scenarios.

Imitation learning has become a topic of focus for robotics research as well, particularly in kinesthetic *learning from demonstration*, a well-studied approach to interactive robot learning in which a human teacher physically guides the robot to complete a task or skill (Argall et al., 2009; Chernova & Thomaz, 2014). We seek to enable imitation for a robotic agent such that, when provided with a single demonstration for completing a task, it can reuse the demonstration to address a related environment that differs in features such as the location, size, or appearance of objects.

Case-based reasoning provides a cognitively-inspired account for storing and reusing experiences individually in memory as *source* cases, before addressing an unfamiliar problem (known as the *target* problem) as follows: (i) retrieve the most relevant *source case* experience from memory, (ii) create a mapping which outlines the differences between the retrieved source case and the target problem, and (iii) use this mapping to adapt the source case such that it can be used to address the target problem (Kolodner, 1993; Thagard, 2005). Using a case-based approach allows us to represent demonstrations as individual experiences in the robot's memory, and provides us with a framework for identifying, transferring, and executing a relevant *source case* demonstration in an unfamiliar, *target environment*. Case-based reasoning and learning from demonstration have been integrated in domains such as RoboCup soccer (Floyd, Esfandiari, & Lam, 2008; Floyd & Esfandiari, 2011; Ontañón et al., 2007; Ros et al., 2009). However, to our knowledge, there is no current work that enables *transfer* for tasks learned by demonstration.

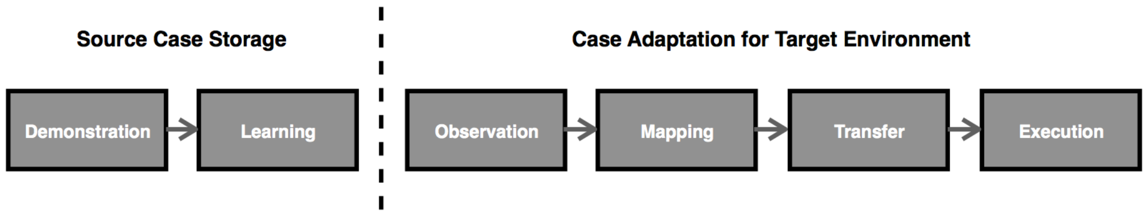

Figure 1: Case-Based Process for Task Demonstration Transfer

2. Completed Work

We have defined a case-based approach to transfer for enabling imitation in robotic agents, consisting of two separate processes (as shown in Figure 2): one in which the robot receives demonstrations of a task and stores each demonstration as a *case* in memory, and a separate system which is used at a later time when the robot is asked to repeat a task in a target environment.

2.1 Case Storage

We have implemented the first process of *case storage*, where the robot records and stores each task demonstration as a *case* in memory. We define each case as the tuple $S =< D,\ T,\ O,\ L >$ where:

- D represents a set of sub-skill models, each representing a segment of the demonstration. The demonstrated motion trajectory is segmented such that each sub-skill corresponds to a primitive action. As an example, a demonstration of scooping pasta into a bowl would be segmented into sub-skills corresponding to the *scoop*, *move-to-target*, and *pour* primitive actions, which can then be executed in sequence to reproduce the full task demonstration. Each sub-skill is represented as a Dynamic Movement Primitive (Schaal, 2006), 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, and is defined as $T =<< {}_xt_0, {}_yt_0, {}_zt_0 >, ..., < {}_xt_n, {}_yt_n, {}_zt_n >>$

- 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 Case Adaptation

At a later time, the robot may be asked to repeat a learned task in an unfamiliar target environment. Using this framework, the robot may address a target environment by (i) observing the target environment, (ii) retrieving the most related source case demonstration from memory, (iii) identifying

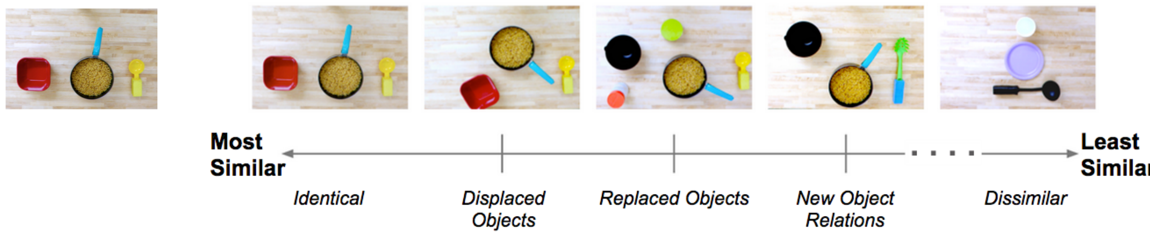

(a) Source Case (b) Potential Target Environments

Figure 2: Spectrum of Similarity Between Source and Target Environments

a mapping which encodes the differences and similarities between the target problem and the retrieved, source case demonstration, (iv) using this mapping to adapt the source case demonstration to be reused in the target problem, and then (v) executing the adapted action in the target environment. We have implemented the last two steps of this process, transfer and execution, and thus manually provide the robot with the most relevant source case demonstration and a mapping between objects in the source and target environments.

2.2.1 Similarity in Transfer

We take a similarity-based approach to transfer, where we consider the similarity between the source case and target environments when defining transfer processes. Figure 2a depicts an overhead view of an environment in which the robot was trained to complete a scooping task. Figure 2b depicts a set of overhead views of 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 ends 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 case demonstration. The remaining three images represent more common and realistic transfer problems, each of which corresponds to a separate level of similarity and is addressed using a separate transfer method.

- *Retargeting Transfer Approach*: The environment shown in the *displaced objects* scene in Figure 2b is very similar to the source environment shown in Figure 2a, differing only the location of each object. Thus, all elements of the source case 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 case representation. 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 case.

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. Furthermore, future work will involve defining a process for identifying and retrieving an appropriate source case demonstration that is most applicable to a given transfer problem.

References

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

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

Floyd, M. W., & Esfandiari, B. (2011). A case-based reasoning framework for developing agents using learning by observation. *Tools with Artificial Intelligence (ICTAI), 2011 23rd IEEE International Conference on* (pp. 531–538).

Floyd, M. W., Esfandiari, B., & Lam, K. (2008). A case-based reasoning approach to imitating robocup players. *FLAIRS Conference* (pp. 251–256).

Kolodner, J. (1993). *Case-based reasoning*. Morgan Kaufmann.

Ontañón, S., Mishra, K., Sugandh, N., & Ram, A. (2007). Case-based planning and execution for real-time strategy games. In *Case-based reasoning research and development*, 164–178. Springer.

Piaget, J., & Cook, M. T. (1952). The origins of intelligence in children.

Ros, R., Arcos, J. L., Lopez de Mantaras, R., & Veloso, M. (2009). A case-based approach for coordinated action selection in robot soccer. *Artificial Intelligence*, *173*, 1014–1039.

Schaal, S. (2006). Dynamic movement primitives-a framework for motor control in humans and humanoid robotics. In *Adaptive motion of animals and machines*, 261–280. Springer.

Thagard, P. (2005). *Mind: Introduction to cognitive science*. MIT press.

Tomasello, M., Kruger, A. C., & Ratner, H. H. (1993). Cultural learning. *Behavioral and brain sciences*, *16*, 495–511.