

A Case-Based Approach To Imitation Learning in Robotic Agents

Tesca Fitzgerald, Ashok Goel

School of Interactive Computing
Georgia Institute of Technology, Atlanta, GA 30332, USA
{tesca.fitzgerald,goel}@cc.gatech.edu

Abstract. Learning by imitation is an essential process in human cognition. Recently, imitation learning has also become important in robotics research. We address the problem of learning by imitation in interactive, robotic agents using case-based reasoning. We describe two tasks for which case-based reasoning may be used: (i) interpretation, in which the robot interprets new skill demonstrations as being related to previous observations, and (ii) imitation, in which a robot seeks to use previously learned skills to address new problem scenarios. We present a case-based framework for imitation and interpretation in a robotic agent that learns from observations of a human teacher.

Keywords: Case-Based Reasoning, Intelligent Agents, Learning by Imitation, Human-Robot Interaction

1 Motivation and Goals

Learning by imitation is a fundamental process of human learning [16] and a central construct of social cognition [5]. Human children exhibit learning by imitation as infants and toddlers [14], as do animals of at least some other species [17]. Thus, it is only natural and logical that learning by imitation has also come to acquire an important place in social robot learning [3, 4, 6].

Robot learning by imitation is a type of observational learning in which the imitated behavior is novel to the robot, and the robot uses the same strategy to accomplish the same goal as the human teacher. Many current approaches to imitation learning in interactive robotics adopt the following methodology: (i) the human teacher demonstrates several instances of a new skill, (ii) the robot generalizes models over the instances, and (iii) the robot transfers the generalized model to new problems. In contrast, we seek to explore a case-based approach to imitation learning in interactive robotics so that (a) learning is incremental, (b) learning is task-specific in that the robot learns an abstraction from the most similar case to the current problem, and (c) learning is “lazy,” meaning that the robot learns the abstraction only when needed.

We seek to address problems such as the following. A human teacher demonstrates a new skill as simple as closing a book (requiring only that the book cover is contacted at some point) and as complex as pouring coffee from a pot into a

cup on a dining table (requiring precision in grasping the handle and tilting the coffee pot). The robot observes the human demonstration. Later, the robot can imitate the human in performing the same skill or similar skills. This process applies to other pairs of similar skills, such as closing a laptop computer after observing a teacher closing a book. There is no other communication between the human and the robot. However, if needed, the human may demonstrate the skill again. Clearly, this task requires us to define processes such as case acquisition, representation, retrieval, adaptation, and execution. Additionally, this raises the difficult issues of defining the robot’s background knowledge, representation language, memory, abstraction, and transfer processes. In this paper, we describe our long-term goals for a case-based learning agent, as well as present preliminary work [9, 10] toward developing such an agent.

2 Background

2.1 Learning by Imitation in Human Cognition

Although learning by imitation has been studied quite extensively, it is not yet a well-understood cognitive or behavioral phenomenon. Indeed, the term “learning by imitation” itself appears ill-defined. Tomasello, Kruger & Ratner [17] consider “true” learning by imitation to entail not only mimicking a sequence of actions, but also learning about the intentions of the teacher or the leader (sometimes called the model). It is unclear whether animal species other than human have the capacity for “true” learning by imitation. Meltzoff [14] found that learning by imitation is common among human infants and toddlers. He also found that for some tasks human infants can understand the teacher’s goal for a sequence of actions without being able to themselves imitate the action sequence. It is unclear how human infants can understand the goal for a sequence of actions without being able to imitate the actions themselves. Byrne & Russon [7] suggest a hierarchical organization of learning by imitation: At one level, the learner learns about action sequences from the teacher, but at another level, the learner learns an abstract “program” similar to hierarchical planning.

Heyes [13] distinguishes between two kinds of theories of learning by imitation: transformational and associative. According to the transformation theory, humans already have the cognitive representations and processes for recreating any behavior; the associative theory suggests that the knowledge needed to replicate a behavior is acquired from the environment. Nehaniv & Dautenhahn [8] analyze the correspondence problem in learning by imitation. The correspondence problem pertains to deciding what sensors and effectors in the learner’s body correspond to what sensors and effectors in the teacher’s body, respectively. This problem is especially acute if the teacher and the learner have different sets of sensors and/or effectors, as typically happens when a robot is learning by imitating a human teacher. In sum, while cognitive and behavioral accounts of learning by imitation in humans have inspired similar work in interactive robotics, at present they provide little guidance on how to actually enable robots to learn by imitation.

2.2 Case-Based Approaches to Observational Learning

Case-based approaches to observational learning have been studied in related work, mostly in the context of games. For example, Ontanon et al. [15] have studied learning by demonstration and online case-based planning in real-time strategy games. However, their focus has been on online case-based planning, and particularly, adapting plans observed from game logs of expert demonstrations. While an important domain, games do not offer the challenges of perception and action to the same degree that interactive robots immediately pose. Perhaps the most relevant work is that of Floyd, Esfandiari & Lam [12], who describe a case-based scheme for learning soccer team skills by observing spatially distributed soccer team plays. More recently Floyd & Esfandiari [11] describe a preliminary scheme for separating domain-independent case-based learning by observation from domain-dependent sensors and effectors on a physical robot. While clearly there is growing interest in case-based approaches to learning by imitation, it seems fair to say that this line of research is still in early stages, especially in the context of interactive robots.

2.3 Learning by Imitation in Robotic Agents

The ability to learn from imitation would allow a robotic agent to learn skills quickly, while enabling the human teacher to convey the skill more intuitively. One focus of Human-Robot Interaction research is to enable robots to interface with end-users who may not be experienced with working with robots [4]. In particular, Learning from Demonstration allows a human teacher to demonstrate a skill via a series of interactions, rather than directly program the robot [3, 4]. The robot learns a representation of the demonstrated skill, which it can then use to mimic the skill at a later time.

There are several methods of interactively demonstrating a skill, including shadowing, teleoperation, and wearing a sensor suit that records the teacher's actions [3]. An additional form of demonstration is pure observation, in which the human teacher demonstrates the skill in front of the robot without guiding the robot to perform the skill. An example of pure observation is a demonstration of pouring a cup, where the teacher demonstrates pouring without physically interacting with the robot. As a result, the robot cannot record a physical action model for the demonstration, but is still able to model the teacher's actions by recording the visual state of the demonstration at frequent intervals.

An intuitive demonstration method is kinesthetic teaching, in which the teacher guides the robot's end-effector hand or gripper, and manually demonstrates how the robot should complete the skill [2, 1]. This method resembles how an adult may guide a young child in completing a skill such as stacking a set of blocks. During each interaction, the robot records its physical state at frequent intervals. The physical state is represented as the angle of each of the robot's joints at a given timestamp. We refer to the robot's *action* model as the trajectory of its end-effector over the course of the demonstration. A series

of demonstrations can be given such that several instances of the skill are addressed. For example, when teaching a robot to perform the skill of closing a box, the teacher may provide demonstrations of the skill using boxes of different shapes, sizes, or locations.

3 Exploring Case-Based Reasoning in Learning by Imitation

We distinguish between two tasks in learning by imitation: interpretation and imitation. We refer to interpretation as a task that labels a skill demonstration based on case memory, and imitation as a task which retrieves and adapts a source case demonstration to accomplish a goal. While a skill model learned over a set of demonstrations can cover a limited range of variations from the demonstrated problem, it is not necessarily applicable to problems that vary more radically from the demonstrated problem. For example, if the robot is taught how to pour a mug, it may be able to apply its skill model for “pouring” to cases in which similar objects are used. However, this skill model may not be applicable to a case in which a different kind of object such as a tablespoon must be poured, despite the two problems being examples of the “pouring” skill. Additionally, a skill model alone cannot be used to identify or interpret new demonstrations as being related to previously learned skills. We posit that both imitation and interpretation in a robotic agent can be addressed by case-based reasoning.

We refer to a *source case* as a skill demonstration that has been provided to the robot and is stored in the robot’s memory. We use the terms *demonstration*, *skill demonstration*, and *case* interchangeably. Each demonstration is defined as $d = \langle p, a \rangle$, where p encodes the *problem* description of the demonstration and a encodes the demonstrated action. We focus on representing demonstrations that illustrate the completion of a single task, such that the problem description contains one action relation r .

A skill demonstration consists of the following elements:

- The problem description $p = \langle r, o, f, v \rangle$, where r is the action relation (“pouring”, “opening”, etc), o is the list of observed objects, f is the list of observed object features (color, size, etc), and v is a set of parameters (object locations, initial position, etc)
- The action model $a = \{j_0, j_1, \dots, j_n\}$, where j_n encodes the robot’s end-effector trajectory consisting of a list of end-effector positions at each recorded time interval

3.1 Categorizing Transfer Problems

Before describing how case-based reasoning can be used to address variations between demonstration problems, we must define how different types of variations in a skill demonstration affect how a skill should be transferred. Particularly,

the action representation associated with the source case must be transferred to address variations in the current problem. The type of variation between the source and target problems defines how this transfer should occur. We refer to two types of problem variations, within-domain and cross-domain, which each require a distinct transfer process [9].

Within-Domain Adaptation Within-domain adaptation is needed when the source and target problems involve similar objects and the same skill. As a result, the source and target problems contain the same set f of object features. Thus, the variation between the two cases is minimal, and only a simple modification of the source action is required to address the target problem. Transferring a “closing” action model from a source case involving a blue box to a target problem involving a red box only requires a simple modification.

Cross-Domain Transfer In contrast, cross-domain transfer is needed when the source and target problems differ significantly. Whereas a within-domain target problem may involve similar objects, a cross-domain target problem may involve an entirely different object than in the source. As a result, the source and target problems contain different sets of object features. An example of this may be transferring a “pouring” skill learned in the context of pouring the contents of a coffee pot into a mug to a target problem in which contents of a tablespoon must be poured; the same pouring skill is applicable to both instances, but must be altered to address the new object’s feature set f . Additional examples include transferring a “closing” skill from a problem involving closing a box to closing a laptop, and transfer “stacking” from stacking blocks to stacking plates. To perform cross-domain transfer, a mapping between the source and target problems’ object feature sets is necessary to transfer the action model in the source problem to the target problem.

4 Case-Based Interpretation

We first investigate the use of case-based reasoning in interpreting skill demonstrations. In interpreting a skill demonstration, the robot identifies a new demonstration as being related to a previously observed demonstration. We primarily focus on skill demonstrations provided via kinesthetic teaching.

Case-based interpretation provides the following benefits to a robotic learner:

1. It can enable classification of skill demonstrations, thus inferring the new demonstration’s action relation based on previous demonstrations.
2. It can be used to generalize action labels, such as incrementally expanding the “pouring” label as demonstrations of the skill are provided using new objects.
3. If the new demonstration was only observation based, rather than kinesthetic, then the robot does not have a skill model which could be used to execute the newly observed skill. By interpreting it in terms of a known skill with a skill model, then the model could be reused to imitate the new skill.

4.1 Case-Based Approach

The main objective of case-based interpretation is to infer a new demonstration’s action relation r from a source case. We approach this task using the process shown in Figure 1. In case-based interpretation, the human teacher interacts with the robotic learner as follows. First, the robotic agent observes the teacher demonstrating the skill via kinesthetic teaching. Then, the agent retrieves a source case with a similar action representation and similar object features from its case memory. Next, the agent adapts the knowledge from the source case, such as the action relation, to the target problem. Finally, the agent stores the target case in the case memory.



Fig. 1: Case-Based Process for the Interpretation Task

Observation In case-based interpretation, the robot is given a kinesthetic demonstration of the skill, which provides the action model a via a *trajectory* or *keyframe* demonstration. In a trajectory demonstration, the robot observes the skill by recording the entire action from start to finish in terms of the positions of each of its joints over time. Another demonstration method involves recording only *keyframes*, which are points during the skill that the teacher specifies as being important to correct completion of the demonstration [2, 1]. The observation then consists of only the robot’s physical state at the keyframes specified by the teacher. The benefit of this demonstration method is that it provides a sparse representation of the skill and allows the demonstration to encode only what is important to successfully completing the skill.

Elements o , f , and v of the problem description can be determined using an overhead camera situated above the robot’s workspace. Since the teacher does not verbally specify the action being performed, the action relation r is initially unknown.

Abstraction Once the skill demonstration has been observed, it is encoded and abstracted. In previous work, we have described two approaches to building skill representations: (i) encoding both the robot’s end-effector trajectory and goal of the demonstration, and (ii) encoding only the goal of the demonstration at an abstracted level [9]. By abstracting the demonstration, the demonstrated problem is represented such that problem constraints are relaxed at each abstraction level. Thus, we can encode the problem description as follows:

- $p_0 = \langle o, f, v \rangle$
- $p_1 = \langle o, f \rangle$
- $p_2 = \langle o \rangle$

Each tuple represents the problem at a higher level of abstraction. Since the action relation is not yet known, only elements o , f , and v are encoded. Likewise, the action model is also abstracted as follows:

- $a_0 = \{j_0, j_1, \dots, j_n\}$ where j_n is the robot’s physical state represented in joint space at time interval n
- $a_1 = \{o_0, o_1, \dots, o_n\}$ where o_n is the robot’s physical state represented in end-effector space, relative to the object position at time interval n
- $a_2 = \{p_0, p_1, \dots, p_n\}$ where p_n is the robot’s end-effector position (excluding end-effector orientation) relative to the object position at time interval n

Again, each abstraction level represents the action model with fewer constraints. Recording the robot’s physical state in seven degree-of-freedom joint space presents more constraints than recording the robot’s state in six degree-of-freedom end-effector position, which records the robot’s end-effector location and orientation along the x, y, and z axis. Likewise, recording the three degree-of-freedom end-effector position (excluding end-effector orientation) relative to the manipulated object presents the fewest constraints.

Memory Retrieval, Transfer, and Storage Finally, the robot uses the abstracted representations of the problem and demonstration to recall a source case from memory, with the goal of inferring the action relation r from the source case with the most similar action representation. One method of source case retrieval may involve choosing the case with a similar action representation at the most specific abstraction level. We describe another method of source case recall in previous work, in which the visual transformation observed in the skill demonstration is compared to visual transformations observed in potential source case demonstrations [10]. The action relation of the chosen source case is then transferred to the new case.

5 Case-Based Imitation

Rather than use source cases to interpret new skill demonstrations, case-based imitation uses source cases to address similar problems, thereby enabling the robot to attempt to address new situations for which an exact demonstration has not yet been provided. For example, if the robot has been taught how to pour using a tablespoon in the problem case shown in Figure 2(a) and now has to pour using a cup in the problem case shown in Figure 2(b), the robot can transfer the “pour” action from the source case by adapting the action to use the cup rather than the tablespoon.

Case-based imitation provides the following benefits to a robotic learner:

1. The robot can still attempt to complete the new task despite the new problem having a different set of parameters and/or objects.
2. The robot can use the source case to bootstrap its learning of the new skill. If it is able to successfully transfer and execute part of the pouring skill in the new problem context, simple corrections from the human teacher can guide the robot to executing the skill correctly in the new problem context.

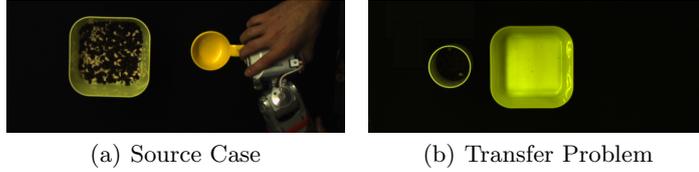


Fig. 2: Variations of a “Pouring” Problem

5.1 Case-Based Approach

During case-based imitation, the human teacher’s interaction with the robotic learner occurs as follows. First, the teacher instructs the robot to complete a specific problem. Given the problem, which indicates the type of action and the object of the action, the robot retrieves the most similar source case. Then, the robot transfers the action from the source case to the target problem and executes the action. The primary goal in case-based imitation is to infer the action model a for the problem from the source case. Figure 3 illustrates the case-based process for this task.



Fig. 3: Case-Based Process for the Imitation Task

Abstraction and Retrieval The imitation process begins when the teacher makes a request such as “pour the blue cup”. Here, only the problem description p is provided, and the robot needs to infer the corresponding action model a from a source case. Both the source case and target problem are abstracted as follows:

- $p_0 = \langle r, o, f, v \rangle$
- $p_1 = \langle r, o, f \rangle$
- $p_2 = \langle r, o \rangle$
- $p_3 = \langle r \rangle$

Since the action relation (e.g. “pour”) is provided in the teacher’s request, r is included in the problem representation. The source case’s action model is abstracted as described previously in the case-based interpretation method. The recall step queries the case memory for the appropriate source case to be retrieved based on p , recalling a similar source case at the most specific abstraction level.

Transfer and Execution Once a source case is retrieved, the action model in the case needs to be transferred to the target problem, where the transfer process depends on the type of similarity between the source case and target problem. The following Table 1 indicates for each type of similarity the (i) elements of p that are similar between the two cases and (ii) what changes to the source case’s action model are necessary for transfer.

A target problem that is identical to the source case (type 1) does not require any adaptation of the source case’s action model. A target problem falls under

	1. Identity	2. Adaptation	3. Transfer	4. Creativity	5. Disparate
Similar Elements	$\langle r, o, f, v \rangle$	$\langle r, o \rangle$	$\langle r \rangle$	$\langle o \rangle$	$\langle \rangle$
Differing Elements	$\langle \rangle$	$\langle f, v \rangle$	$\langle o, f, v \rangle$	$\langle r, f, v \rangle$	$\langle r, o, f, v \rangle$

Table 1: Types of Similarity in Transfer Problems

type 2 or type 3 depending on whether it contains within-domain or cross-domain variations from the source case, respectively. The fourth category, creativity, is manifested when the robot takes inspiration from one or more previous cases to perform a distinct action. An example of creative adaptation would be referencing a source case of opening a box when attempting to close the same box. We do not yet describe how this creative process would occur, but leave this as a topic for future work. Finally, category 5 represents a class of problem cases which are dissimilar from any source case and are beyond the scope of any adaptation or transfer process, and thus cannot be addressed using case-based reasoning.

Once the level of abstraction has been chosen, a new action model is created for the target case which adheres to the constraints. The rest of the action model is created to address the transfer problem. The resulting representation is then converted into a trajectory in terms of the robot’s joint space and executed.

6 Conclusions and Future Work

In this paper, we presented an initial analysis of imitative learning for robotic agents from the perspective of case-based reasoning. We have described a case-based approach to imitation and interpretation in a robotic agent that learns from demonstration. Thus far, our work has focused on the implementing the interpretation task on an interactive robot [10, 9]. Future work will involve implementing the imitation task. To do this, we must further explore how the transfer step should be implemented in the imitation process. Our initial focus will be on defining a process for adaptation in the first three types of transfer problems listed in Table 1.

7 Acknowledgements

We thank Andrea Thomaz for many helpful discussions about the robotics aspects of this work. This material is based upon work supported by the United States’ National Science Foundation through Graduate Research Fellowship Grant #DGE-1148903 and Robust Intelligence Grant #1116541. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NSF.

References

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

2. Akgun, B., Cakmak, M., Wook Yoo, J., Thomaz, L.A.: Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective. In: ACM/IEEE Intl. Conference on Human-robot interaction (HRI). pp. 391–398 (2012)
3. Argall, B.D., Chernova, S., Veloso, M., Browning, B.: A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57(5), 469–483 (2009)
4. Atkeson, C.G., Schaal, S.: Robot learning from demonstration. In: ICML. vol. 97, pp. 9–15 (1997)
5. Bandura, A.: *Social foundations of thought and action*. Englewood Cliffs, NJ Prentice Hall. (1986)
6. Breazeal, C., Scassellati, B.: Robots that imitate humans. *Trends in cognitive sciences* 6(11), 481–487 (2002)
7. Byrne, R.W., Russon, A.E.: Learning by imitation: a hierarchical approach. *Behavioral and brain sciences* 21(05), 667–684 (1998)
8. Dautenhahn, K., Nehaniv, C.L.: *The correspondence problem*. MIT Press (2002)
9. Fitzgerald, T., Goel, A.K., Thomaz, A.L.: Representing skill demonstrations for adaptation and transfer. In: *AAAI Symposium on Knowledge, Skill, and Behavior Transfer in Autonomous Robots* (2014)
10. Fitzgerald, T., McGregor, K., Akgun, B., Goel, A.K., Thomaz, A.L.: A visual analogy approach to source case retrieval in robot learning from observation. In: *AAAI Workshop on Artificial Intelligence and Robotics* (2014)
11. Floyd, M.W., Esfandiari, B.: A case-based reasoning framework for developing agents using learning by observation. In: *Tools with Artificial Intelligence (ICTAI), 2011 23rd IEEE International Conference on*. pp. 531–538. IEEE (2011)
12. Floyd, M.W., Esfandiari, B., Lam, K.: A case-based reasoning approach to imitating robocup players. In: *FLAIRS Conference*. pp. 251–256 (2008)
13. Heyes, C.: Transformational and associative theories of imitation. In: *Imitation in animals and artifacts*. pp. 501–523. MIT Press (2002)
14. Meltzoff, A.N.: Imitation and other minds: The like me hypothesis. *Perspectives on imitation: From neuroscience to social science* 2, 55–77 (2005)
15. Ontañón, S., Mishra, K., Sugandh, N., Ram, A.: Case-based planning and execution for real-time strategy games. In: *Case-Based Reasoning Research and Development*, pp. 164–178. Springer (2007)
16. Piaget, J., Cook, M.T.: *The origins of intelligence in children*. (1952)
17. Tomasello, M., Kruger, A.C., Ratner, H.H.: Cultural learning. *Behavioral and brain sciences* 16(03), 495–511 (1993)