

## A survey of robot learning from demonstration

Brenna D. Argall<sup>a,\*</sup>, Sonia Chernova<sup>b</sup>, Manuela Veloso<sup>b</sup>, Brett Browning<sup>a</sup>

<sup>a</sup> Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA

<sup>b</sup> Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA

### ARTICLE INFO

#### Article history:

Received 23 May 2008

Received in revised form

15 October 2008

Accepted 25 October 2008

Available online 25 November 2008

#### Keywords:

Learning from demonstration

Robotics

Machine learning

Autonomous systems

### ABSTRACT

We present a comprehensive survey of robot *Learning from Demonstration (LfD)*, a technique that develops policies from example state to action mappings. We introduce the LfD design choices in terms of demonstrator, problem space, policy derivation and performance, and contribute the foundations for a structure in which to categorize LfD research. Specifically, we analyze and categorize the multiple ways in which examples are gathered, ranging from teleoperation to imitation, as well as the various techniques for policy derivation, including matching functions, dynamics models and plans. To conclude we discuss LfD limitations and related promising areas for future research.

© 2008 Elsevier B.V. All rights reserved.

### 1. Introduction

The problem of learning a mapping between world state and actions lies at the heart of many robotics applications. This mapping, also called a *policy*, enables a robot to select an action based upon its current world state. The development of policies by hand is often very challenging and as a result machine learning techniques have been applied to policy development. In this survey, we examine a particular approach to policy learning, *Learning from Demonstration (LfD)*.

Within LfD, a policy is learned from *examples*, or demonstrations, provided by a teacher. We define examples as sequences of state–action pairs that are recorded during the teacher’s demonstration of the desired robot behavior. LfD algorithms utilize this dataset of examples to derive a policy that reproduces the demonstrated behavior. This approach to obtaining a policy is in contrast to other techniques in which a policy is learned from *experience*, for example building a policy based on data acquired through exploration, as in Reinforcement Learning [1]. We note that a policy derived under LfD is necessarily defined only in those states encountered, and for those corresponding actions taken, during the example executions.

In this article, we present a survey of recent work within the LfD community, focusing specifically on robotic applications. We segment the LfD learning problem into two fundamental phases:

*gathering* the examples, and *deriving* a policy from such examples. Based on our identification of the defining features of these techniques, we contribute a comprehensive survey and categorization of existing LfD approaches. Though LfD has been applied to a variety of robotics problems, to our knowledge there exists *no* established structure for concretely placing work within the larger community. In general, approaches are appropriately contrasted to similar or seminal research, but their relation to the remainder of the field lies largely unaddressed. Establishing these relations is further complicated by dealing with real world robotic platforms, for which the physical details between implementations may vary greatly and yet employ fundamentally identical learning techniques, or vice versa. A categorical structure therefore aids in comparative assessments among applications, as well as in identifying open areas for future research. In contributing our categorization of current approaches, we aim to lay the foundations for such a structure.

For the remainder of this section we motivate the application of LfD to robotics, and present a formal definition of the LfD problem. Section 2 presents the key design decisions for an LfD system. Methods for gathering demonstration examples are the focus of Section 3, where the various approaches to teacher demonstration and data recording are discussed. Section 4 examines the core techniques for policy derivation within LfD, followed in Section 5 by methods for improving robot performance beyond the capabilities of the teacher examples. To conclude, we identify and discuss open areas of research for future work in Section 6 and summarize the article with Section 7.

\* Corresponding author. Tel.: +1 4122689923.

E-mail addresses: [bargall@cs.cmu.edu](mailto:bargall@cs.cmu.edu) (B.D. Argall), [soniac@cs.cmu.edu](mailto:soniac@cs.cmu.edu) (S. Chernova), [mveloso@cs.cmu.edu](mailto:mveloso@cs.cmu.edu) (M. Veloso), [brettb@cs.cmu.edu](mailto:brettb@cs.cmu.edu) (B. Browning).

### 1.1. Support for demonstration learning

The presence of robots within society is becoming ever more prevalent. Whether an exploration rover in space, robot soccer or a recreational robot for the home, successful autonomous robot operation requires robust control algorithms. Non-robotics-experts may be increasingly presented with opportunities to interact with robots, and it is reasonable to expect that they have ideas about what a robot should do, and therefore what sort of behaviors these control algorithms should produce. A natural, and practical, extension of having this knowledge is to actually develop the desired control algorithm. Currently, however, policy development is a complex process restricted to experts within the field.

Traditional approaches to robot control model domain dynamics and derive mathematically-based policies. Though theoretically well-founded, these approaches depend heavily upon the accuracy of the world model. Not only does this model require considerable expertise to develop, but approximations such as linearization are often introduced for computational tractability, thereby degrading performance. Other approaches, such as Reinforcement Learning, guide policy learning by providing reward feedback about the desirability of visiting particular states. To define a function to provide the reward, however, is known to be difficult and requires considerable expertise to address. Furthermore, building the policy requires gathering information by visiting states to receive rewards, which is non-trivial for a robot learner executing actual actions in the real world.

Considering these challenges, LfD has many attractive points for both learner and teacher. LfD formulations typically do not require expert knowledge of the domain dynamics, which removes performance brittleness resulting from model simplifications. The absence of this expert domain knowledge requirement also opens policy development to non-robotics-experts, satisfying a need that increases as robots become more commonplace. Furthermore, demonstration has the attractive feature of being an intuitive medium for communication from humans, who already use demonstration to teach other humans. Demonstration also has the practical feature of focusing the dataset to areas of the state-space actually encountered during task execution.

### 1.2. Problem statement

LfD can be seen as a subset of *Supervised Learning*. In Supervised Learning the agent is presented with labeled training data and learns an approximation to the function which produced the data. Within LfD, this training dataset is composed of example executions of the task by a demonstration teacher (Fig. 1, top).

We formally construct the LfD problem as follows. The world consists of states  $S$  and actions  $A$ , with the mapping between states by way of actions being defined by a probabilistic transition function  $T(s^i|s, a) : S \times A \times S \rightarrow [0, 1]$ . We assume that the state is not fully observable. The learner instead has access to observed state  $Z$ , through the mapping  $M : S \rightarrow Z$ . A policy  $\pi : Z \rightarrow A$  selects actions based on observations of the world state. A single cycle of policy execution at time  $t$  is shown in Fig. 1 (bottom).

The set  $A$  ranges from containing low-level motions to high-level behaviors. For some simulated world applications, state may be fully transparent, in which case  $M = I$ , the identity mapping. For all other applications state is not fully transparent and must be observed, for example through sensors in the real world. For succinctness, throughout the text we will use “state” interchangeably with “observed state.” It should be assumed, however, that state is always the observed state, unless explicitly noted otherwise. This assumption will be reinforced by use of the  $Z$  notation throughout the text.

Throughout the teacher execution, states and selected actions are recorded. We represent a demonstration  $d_j \in D$  formally as  $k_j$  pairs of observations and actions:  $d_j = \{(z_j^i, a_j^i)\}$ ,  $z_j^i \in Z$ ,  $a_j^i \in A$ ,  $i = 0 \dots k_j$ . These demonstrations set LfD apart from other learning approaches. The set  $D$  of the demonstrations is made available to the learner. The policy derived from this dataset enables the learner to select an action based on the current state.

### 1.3. Terminology and context

Before continuing, we pause to place the intents of this survey within the context of previous LfD literature. The aim of this survey is to review the broad topic of LfD, to provide a categorization that highlights differences between approaches, and to identify research areas within LfD that have not yet been explored.

We begin with a comment on terminology. Demonstration-based learning techniques are described by a variety of terms within the published literature, including Learning by Demonstration (LbD), Learning from Demonstration (LfD), Programming by Demonstration (PbD), Learning by Experienced Demonstrations, Assembly Plan from Observation, Learning by Showing, Learning by Watching, Learning from Observation, behavioral cloning, imitation and mimicry. While the definitions for some of these terms, such as imitation, have been loosely borrowed from other sciences, the overall use of these terms is often inconsistent or contradictory across articles.

Within this article, we refer to the general category of algorithms in which a policy is derived based on demonstrated data as *Learning from Demonstration (LfD)*. Within this category, we further distinguish between approaches by their various characteristics, as outlined in Section 2, such as the source of the demonstrations and the learning techniques applied. Subsequent sections introduce terms used to characterize algorithmic differences. Due to the already contradictory use of terms in the existing literature, our definitions will not always agree with those of other publications. Our intent, however, is not for others in the field to adopt the terminology presented here, but rather to provide a consistent set of definitions that highlight distinctions between techniques.

Regarding a categorization for approaches, we note that many legitimate criteria could be used to subdivide LfD research. For example, one proposed categorization considers the broad spectrum of *who, what, when and how to imitate*, or subsets thereof [2,3]. Our review aims to focus on the specifics of implementation. We therefore categorize approaches according to the computational formulations and techniques required to implement an LfD system.

To conclude, readers may also find useful other related surveys of the LfD research area. In particular, the book *Imitation in*




Fig. 1. Control policy derivation and execution.

*Animals and Artifacts* [4] provides an interdisciplinary overview of research in imitation learning, presenting leading work from neuroscience, psychology and linguistics as well as computer science. A narrower focus is presented in the chapter “Robot Programming by Demonstration” [2] within the book *Handbook of Robotics*. This work particularly highlights techniques which may augment or combine with traditional LfD, such as giving the teacher an active role during learning. By contrast, our focus is to provide a categorical structure for LfD approaches, in addition to presenting the specifics of implementation. We do refer the reader to this chapter for a more comprehensive historical overview of LfD, as the scope of our survey is restricted to recently published literature. Additional reviews that cover specific sub-areas of LfD research in detail are highlighted throughout the article.

## 2. Design choices

There are certain aspects of LfD which are common among all applications to date. One is the fact that a teacher demonstrates execution of a desired behavior. Another is that the learner is provided with a set of these demonstrations, and from them derives a policy able to reproduce the demonstrated behavior.

However, the developer still faces many design choices when developing a new LfD system. Some of these decisions, such as the choice of a discrete or continuous action representation, may be determined by the domain. Other design choices may be up to the preference of the developer. As we discuss in the later sections, these design decisions strongly influence how the learning problem is structured and solved. In this section we highlight those decisions that are the most significant to make.

To illustrate these choices, we pair the presentation with a running *pick and place* example in which a robot must move a box from a table to a chair. To do so, the object must be (1) picked up, (2) relocated and (3) put down. We present alternate representations and/or learning methods for this task, to illustrate how these particular choices influence task formalization and learning.

### 2.1. Demonstration approach

Within the context of gathering teacher demonstrations, two key decisions must be made: the choice of *demonstrator*, and the choice of *demonstration technique*. We discuss various choices for each of these decisions below. Note that these decisions are at times affected by factors such as the complexity of the robot and task. For example, teleoperation is rarely used with high degree of freedom humanoids, since their complex motions are typically difficult to control via joystick.

#### 2.1.1. The choice of demonstrator

Most LfD work to date has made use of human demonstrators, although some techniques have also examined the use of robotic teachers, hand-written control policies and simulated planners. The choice of demonstrator further breaks down into the subcategories of (i) who *controls* the demonstration and (ii) who *executes* the demonstration.

For example, consider a robot learning to move a box, as described above. One demonstration approach could have a robotic teacher pick up and relocate the box using its own body. In this case a *robot* teacher controls the demonstration, and its *teacher* body executes the demonstration. An alternate approach could have a human teacher teleoperate the robot learner through the task of picking up and relocating the box. In this case a *human* teacher controls the demonstration, and the *learner* body executes the demonstration. The choice of demonstrator has a significant impact on the type of learning algorithms that can be applied. As we discuss in Section 3, the similarity between the state and action spaces of the teacher and learner determines the kinds of algorithms that may be required to process the data.

#### 2.1.2. Demonstration technique

The choice of demonstration technique refers to the *strategy* for providing data to the learner. One option is to perform *batch learning*, in which case the policy is learned only once all data has been gathered. Alternatively, *interactive approaches* allow the policy to be updated incrementally as training data becomes available, possibly provided in response to current policy performance. Examples of both approaches are highlighted throughout the article.

### 2.2. Problem space continuity

The question of continuity plays a prominent role within the context of state and action representation, and many valid representations frequently exist for the same domain. Within our robot box moving example, one option could be to discretize state, such that the environment is represented by Boolean features such as *box on table* and *box held by robot*. Alternatively, a continuous state representation could be used in which the state is represented by the 3D positions of the robot’s end effector and the box. Similar discrete or continuous representations could be chosen for the robot’s actions.

In designing a domain, the continuity of the problem space may be determined by many factors, such as the desired learned behavior, the set of available actions and whether the world is simulated or real. As discussed in Section 4, the question of continuity heavily influences how suitable the various policy derivation techniques are for addressing a given problem.

Additionally, we note that LfD can be applied at a variety of action control levels, depending on the problem formulation. We roughly group actions into three control levels: low-level actions for motion control, basic high-level actions (often called action primitives) and complex behavioral actions for high-level control. Note that this is a somewhat different consideration to action–space continuity. For example, a low-level motion could be formulated as discrete or continuous, and so this action level can map to either space. As a general technique, LfD can be applied at any of these action levels. Most important in the context of policy derivation, however, is whether actions are continuous or discrete, and not their control level. For the remainder of the article, we distinguish between representations based on continuity only.

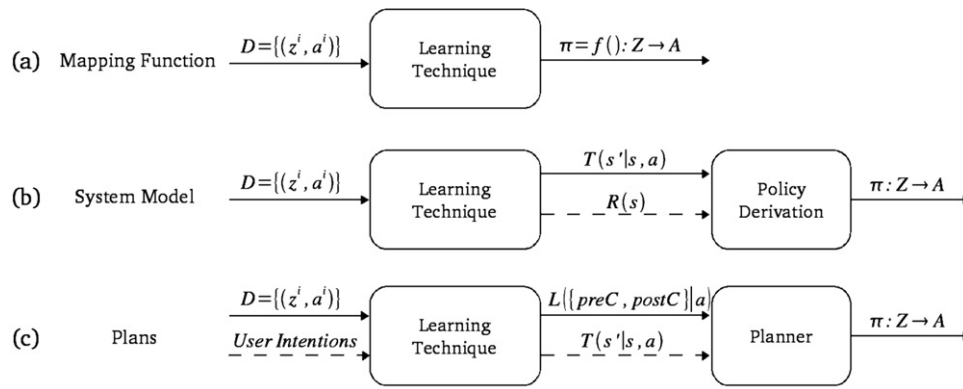
### 2.3. Policy derivation and performance

In selecting an algorithm for generating a policy, we consider two key decisions: the general technique used to *derive the policy*, and whether performance can *improve beyond the teacher’s demonstrations*. These decisions are influenced by action-continuity, as described in the previous section, which is in turn determined both by task and robot capabilities.

#### 2.3.1. Policy derivation technique

As summarized in Fig. 2, research within LfD has seen the development of three core approaches to policy derivation from demonstration data, which we define as *mapping function*, *system model*, and *plans*:

- *Mapping function* (Section 4.1): Demonstration data is used to directly approximate the underlying function mapping from the robot’s state observations to actions ( $f() : Z \rightarrow A$ ).
- *System model* (Section 4.2): Demonstration data is used to determine a model of the world dynamics ( $T(s'|s, a)$ ), and possibly a reward function ( $R(s)$ ). A policy is then derived using this information.



**Fig. 2.** Policy derivation using the generalization approach of determining (a) an approximation to the state  $\rightarrow$  action mapping function, (b) a dynamics model of the system and (c) a plan of sequenced actions.

- *Plans* (Section 4.3): Demonstration data, and often additional user intention information, is used to learn rules that associate a set of pre- and post-conditions with each action ( $L(\{preC, postC\}|a)$ ), and possibly a sparsified state dynamics model ( $T(s'|s, a)$ ). A sequence of actions is then planned using this information.

Returning again to our example, suppose a mapping function approach is used to derive the policy. A function  $f() : Z \rightarrow A$  is learned that maps the observed state of the world, for example the 3D location of the robot's end effector, to an action which guides the learner towards the goal state, for example the desired end effector motor speed. Consider instead using a system model approach. Here a state transition model  $T(s'|s, a)$  is learned, for example that taking the *pick up* action when in state *box on table* results in state *box held by robot*. Using this model, the derived policy indicates the best action to take when in a given state, to guide the robot towards the goal state. Finally, consider using a planning approach. The pre- and post-conditions of executing an action  $L(\{preC, postC\}|a)$  are learned from the demonstrations. For example, the *pick up* action requires the *box on table* pre-condition, and results in the *box held by robot* post-condition. A planner uses this learned information to produce a sequence of actions that ends with the robot in the goal state. Each of the three approaches are discussed in detail within Section 4.

### 2.3.2. Dataset limitations

Training examples obtained from demonstration are inherently limited by the performance of the teacher. In many domains, it is possible that this teacher performance is suboptimal when compared with the abilities of the learner. For example, a human teacher may not be physically able to execute actions as quickly or accurately as a robot. Since the learner derives its policy from these examples, the performance of this policy is therefore also limited by the teacher's abilities. Many LfD learning systems, however, have been augmented to enable learner performance to improve beyond what was provided in the demonstration dataset. Examples include the incorporation of teacher advice or Reinforcement Learning techniques. These approaches are discussed in depth within Section 5.

## 3. Gathering examples: How the dataset is built

In this section, we discuss various techniques for executing and recording demonstrations. The LfD dataset is composed of state-action pairs recorded during teacher executions of the desired behavior. Exactly *how* they are recorded, and *what* the teacher uses as a platform for the execution, varies greatly across approaches. Examples range from sensors on the robot learner

recording its own actions as it is passively teleoperated by the teacher, to a camera recording a human teacher as she executes the behavior with her own body.

For LfD to be successful, the states and actions in the learning dataset must be usable by the student. In the most straightforward setup, the states and actions of the teacher executions map directly to the learner. In reality, however, a direct mapping will often not be possible, as the learner and teacher will likely differ in sensing or mechanics. For example, a robot learner's camera will not detect state changes in the same manner as a human teacher's eyes, nor will its gripper apply force in the same manner as a human hand. The challenges which arise from these differences are referred to broadly as *Correspondence Issues* [5].

### 3.1. Correspondence

The issue of correspondence deals with the identification of a mapping between the teacher and the learner that allows the transfer of information from one to the other. In this survey, we define correspondence with respect to two mappings, shown in Fig. 3: the *record mapping*, and the *embodiment mapping*.

- The *Record Mapping* (Teacher Execution  $\rightarrow$  Recorded Execution) refers to whether the exact states/actions experienced by the teacher during the demonstration execution are recorded within the dataset.
- The *Embodiment Mapping* (Recorded Execution  $\rightarrow$  Learner) refers to whether the states/actions recorded within the dataset are exactly those that the learner would observe/execute.

When the *record mapping* is the identity  $I(z, a)$ , the states/actions experienced by the teacher during execution are directly recorded in the dataset. Otherwise this teacher information is encoded according to some record mapping function  $g_R(z, a) \neq I(z, a)$ , and this encoded information is recorded within the dataset. Similarly, when the *embodiment mapping* is the identity  $I(z, a)$ , the states/actions in the dataset map directly to the learner. Otherwise the embodiment mapping consists of some function  $g_E(z, a) \neq I(z, a)$ . For any given learning system, it is possible to have neither, either or both of the record and embodiment mappings be the identity. Note that the mappings do not change the content of the demonstration data, but only the reference frame within which it is represented. Fig. 4 shows the intersection of these configurations, which we discuss further within subsequent sections.

The embodiment mapping is particularly important when considering real robots, compared with simulated agents. Since actual robots execute real actions within a physical environment, providing them with a demonstration involves a physical execution by the teacher. Learning within this setting depends heavily upon an

Fig. 3. Mapping a teacher execution to the learner.

Fig. 4. Intersection of the record and embodiment mappings. The left and right columns represent an identity (Demonstration) and non-identity (Imitation) embodiment mapping, respectively. Each column is then subdivided by an identity (top) or non-identity (bottom) record mapping. Typical approaches to providing data are listed within the quadrants.

accurate mapping between the recorded dataset and the learner's abilities.

Recalling again our box relocation example, consider a human teacher using her own body to demonstrate moving the box, and that a camera records the demonstration. Let the teacher actions,  $A_T$ , be represented as human joint angles, and the learner actions,  $A_L$ , be represented as robot joint angles. In this context, the robot observes the teacher's demonstration of the task through the camera images. The teacher's exact actions are unknown to the robot; instead, this information must be extracted from the image data. This is an example of a  $g_R(z, a) \neq I(z, a)$  record mapping  $A_T \rightarrow D$ . Furthermore, the physical embodiment of the teacher is different from that of the robot and his actions ( $A_T$ ) are therefore *not* the same as those of the robot ( $A_L$ ). Therefore in order to make the demonstration data meaningful for the robot, a mapping  $D \rightarrow A_L$  must be applied to convert the demonstration into the robot's frame of reference. This is one example of a  $g_E(z, a) \neq I(z, a)$  embodiment mapping.

The categorization of LfD data sources that we present in this article groups approaches according to the absence or presence of the record and embodiment mappings. We select this categorization to highlight the levels at which correspondence plays a role in demonstration learning. Within a given learning approach, the inclusion of each additional mapping introduces a potential injection point for correspondence difficulties; in short, the more mappings, the more difficult it is to recognize and reproduce the teacher's behavior. However, mappings also reduce constraints on the teacher and increase the generality of the demonstration technique.

In our categorization, we first split LfD data acquisition approaches into two categories based on the embodiment mapping, and thus by execution platform:

- *Demonstration*: There is no embodiment mapping, because demonstration is performed on the actual robot learner (or a physically identical platform). Thus  $g_E(z, a) \equiv I(z, a)$ .

- *Imitation*: There exists an embodiment mapping, because demonstration is performed on a platform which is *not* the robot learner (or a not physically identical platform). Thus  $g_E(z, a) \neq I(z, a)$ .

We then further distinguish approaches within each of these categories according to record mapping, relating to how the demonstration is recorded. Fig. 5 introduces our full categorization of the various approaches for building the demonstration dataset. We structure our discussion of data acquisition in subsequent sections according to this categorization.

### 3.2. Demonstration

When teacher executions are *demonstrated*, by our definition there exists no embodiment mapping issue between the teacher and learner. This situation is presented in the left column of Fig. 4. There may exist a non-direct record mapping, however, for state and/or actions, if the states experienced (actions taken) by the demonstrator are not recorded directly, and must instead be inferred from the data. Based on this distinction, we identify two common approaches for providing demonstration data to the robot learner as:

- *Teleoperation* (Section 3.2.1): A demonstration technique in which the teacher *operates* the robot learner platform and the robot's sensors record the execution. The record mapping is direct; thus  $g_R(z, a) \equiv I(z, a)$ .
- *Shadowing* (Section 3.2.2): A demonstration technique in which the robot learner records the execution using its own sensors while attempting to match or *mimic* the teacher motion as the teacher executes the task. There exists a non-direct record mapping; thus  $g_R(z, a) \neq I(z, a)$ .

Again, for both teleoperation and shadowing the robot records from its own sensors as its body executes the behavior, and so the embodiment mapping is direct,  $g_E(z, a) \equiv I(z, a)$ .

The record mapping distinction plays an important role in the application and development of demonstration algorithms. As described below, teleoperation is not suitable for all learning platforms, while shadowing techniques require an additional processing component to enable the learner to mimic the teacher. In the following subsections we discuss various works that utilize these demonstration techniques.

#### 3.2.1. Teleoperation

During teleoperation, a robot is operated by the teacher while recording from its own sensors. Since the robot directly records the states/actions experienced during the execution, the record mapping is direct and  $g_R(z, a) \equiv I(z, a)$ . Teleoperation provides the most direct method for information transfer within demonstration learning. However, teleoperation requires that operating the robot be manageable, and as a result not all systems are suitable for this technique. For example low-level motion demonstrations are difficult on systems with complex motor control, such as high degree of freedom humanoids.

Demonstrations recorded through human teleoperation via a joystick are used in a variety of applications, including flying a robotic helicopter [6], robot kicking motions [7], object grasping [8, 9], robotic arm assembly tasks [10] and obstacle avoidance and navigation [11,12]. Teleoperation is also applied to a wide variety



















