

A survey of robot learning from demonstration

Brenna D. Argall^{a,*}, Sonia Chernova^b, Manuela Veloso^b, Brett Browning^a

^a Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA

^b 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 (B.D. Argall), soniac@cs.cmu.edu (S. Chernova), mveloso@cs.cmu.edu (M. Veloso), 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*

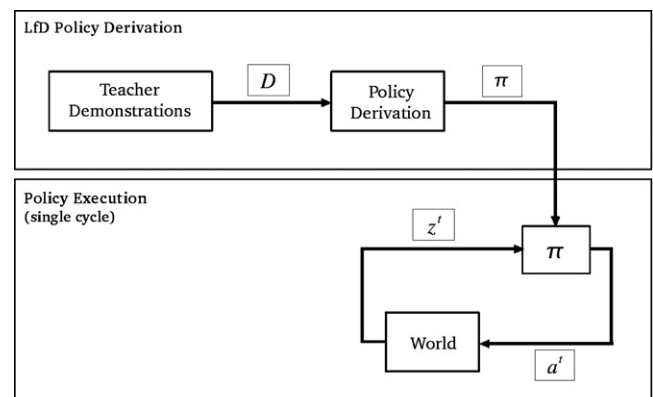


Fig. 1. Control policy derivation and execution.

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.

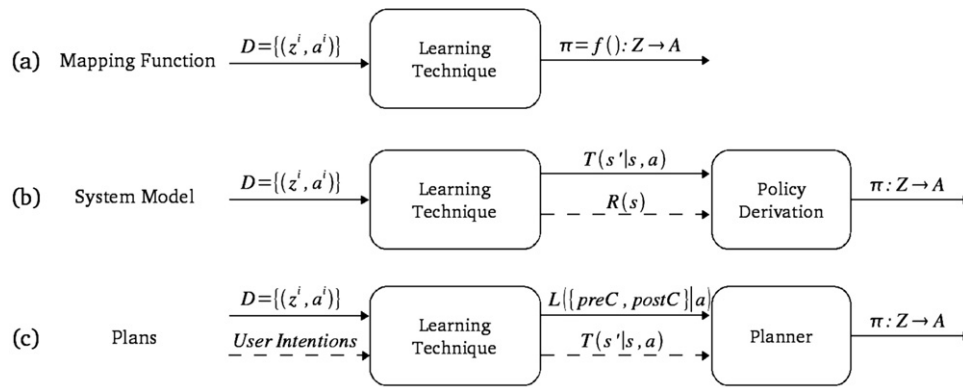


Fig. 2. Policy derivation using the generalization approach of determining (a) an approximation to the state \rightarrow action mapping function, (b) a dynamics model of the system and (c) a plan of sequenced actions.

- *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.

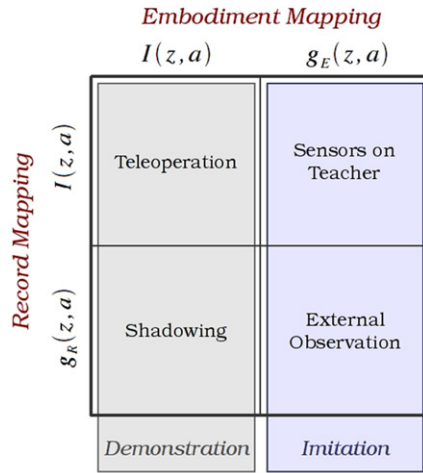


Fig. 4. Intersection of the record and embodiment mappings. The left and right columns represent an identity (Demonstration) and non-identity (Imitation) embodiment mapping, respectively. Each column is then subdivided by an identity (top) or non-identity (bottom) record mapping. Typical approaches to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

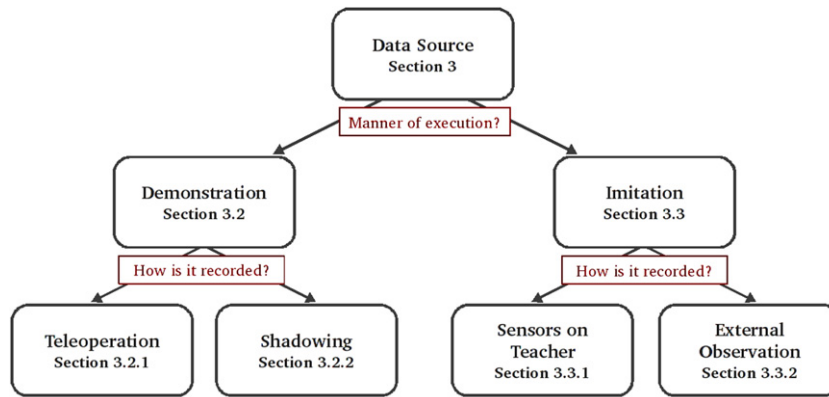


Fig. 5. Categorization of approaches to building the demonstration dataset.

of simulated domains, ranging from static mazes [13,14] to dynamic driving [15,16] and soccer domains [17], and many other applications.

Human teachers also employ techniques other than direct joysticking. In kinesthetic teaching, a humanoid robot is not actively controlled but rather its passive joints are moved through desired motions [18]. Demonstration may also be performed through speech dialog, in which the teacher specifically tells the robot what actions to execute in various states [19–21]. Both of these techniques might be viewed as variants on traditional teleoperation, or alternatively as a sort of high-level teleoperation. In place of a human teacher, hand-written controllers are also used to teleoperate robots [22–24,12].

We note that demonstration data recorded using real robots frequently does not represent the *full* observation state of the teacher. This occurs if, while executing, the teacher employs extra sensors that are not recorded. For example, if the teacher observes parts of the world that are inaccessible from the robot’s cameras (e.g., behind the robot, if its cameras are forward-facing), then state, as observed by the teacher, differs from what is actually recorded as data. A small number of works do address this problem. For example, in Grudic and Lawrence [25] a vision-based robot is teleoperated while the teacher looks exclusively at a screen displaying the robot’s camera output.

3.2.2. Shadowing

During shadowing, the robot platform mimics the teacher’s demonstrated motions while recording from its own sensors. The record mapping is not direct, $g_R(z, a) \neq I(z, a)$, because the states/actions of the true demonstration execution are not recorded. Rather, the learner records its own mimicking execution, and so the teacher’s states/actions are *indirectly* encoded within the dataset. In comparison to teleoperation, shadowing requires an extra algorithmic component which enables the robot to track and actively shadow (rather than passively be teleoperated by) the teacher.

Navigational tasks learned from shadowing have a robot follow an identical-platform robot teacher through a maze [26], follow a human teacher past sequences of colored markers [27] and mimic routes determined from observations of human teacher executions [28]. A humanoid learns arm gestures, as well as the rules of a turn-taking gesture game, by mimicking the motions of a human demonstrator [29].

3.3. Imitation

As previously defined, embodiment issues do exist between the teacher and learner for *imitation* approaches. This situation is

presented in the right column of Fig. 4, indicating the presence of a non-identity embodiment mapping $g_E(z, a) \neq I(z, a)$. Within this setting, we further divide approaches for providing imitation data by whether the record mapping is the identity or not:

- *Sensors on teacher* (Section 3.3.1): An imitation technique in which sensors located on the executing body are used to record the teacher execution. The record mapping is direct; thus $g_R(z, a) \equiv I(z, a)$.
- *External observation* (Section 3.3.2): An imitation technique in which sensors external to the executing body are used to record the execution. These sensors may or may not be located on the robot learner. There exists a non-direct record mapping; thus $g_R(z, a) \neq I(z, a)$.

The record mapping distinction again plays an important role in the application and development of imitation algorithms, just as it did with demonstration approaches. The sensors on teacher approach provides precise measurements, but requires a lot of overhead in the way of specialized sensors and customized surroundings. By contrast, external observation is a more general method, but also provides less reliable measurements. In the following subsections, we discuss these imitation techniques and the various works that utilize them. Additionally, we point the reader to a review that examines the problems of correspondence and knowing what to imitate [30].

3.3.1. Sensors on teacher

The sensors-on-teacher approach utilizes recording sensors located *directly* on the executing platform. This means no record mapping, and so $g_R(z, a) \equiv I(z, a)$, which alleviates one potential source for correspondence difficulties. The strength of this technique is that the teacher provides precise measurements of the example execution. However, the overhead attached to the specialized sensors, such as human-wearable sensor-suits, or customized surroundings, such as rooms outfitted with cameras, is non-trivial and limits the applicability settings of this technique.

Human teachers commonly use their own bodies to perform example executions by wearing sensors able to record the person’s state and actions. This is especially true when working with humanoid or anthropomorphic robots, since the body of the robot resembles that of a human. The recorded joint angles of a human teach drumming patterns to a 30-DoF humanoid [31], and in later work walking patterns as well [32]. A humanoid learns referee signals by pairing kinesthetic teachings (3.2.1) with human teacher executions recorded via wearable motion sensors [33]. Another approach has a human wearing sensors control a simulated human, which maps to a simulated robot and then to a real robot arm [34].

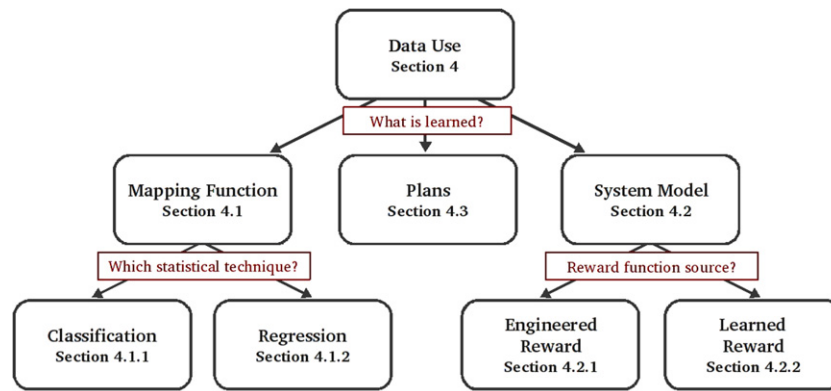


Fig. 6. Categorization of approaches to learning a policy from demonstration data.

3.3.2. External observation

Imitation performed through external observation relies on data recorded by sensors located *externally* to the executing platform, meaning that a record mapping exists, and $g_R(z, a) \neq I(z, a)$. Since the robot does not directly record the actual states/actions experienced by the teacher during the execution, they must be inferred, introducing a new source of uncertainty for the learner. Some LfD implementations first extract the teacher states/actions from this recorded data, and then map the extracted states/actions to the learner. Others map the recorded data directly to the learner, without ever explicitly extracting the states/actions of the teacher. Compared to the sensors-on-teacher approach, the data recorded under this technique is less precise and less reliable. The method, however, is more general and is not limited by the overhead of specialized sensors and settings.

An early seminal work has a robotic arm learn pole balancing via stereo-vision and human demonstration [35]. Unlike most other approaches in this section, human demonstration trains a non-anthropomorphic robot and data is obtained by tracking the movement of the pole, not the teacher's arm itself. Similarly, the position of the human player's (teacher's) puck is tracked when teaching a 37-DoF humanoid to play air hockey [36]. The robot learner itself directly observes the teacher executions, as in other humanoid systems, where cameras are often located directly on the robot body as a parallel to human eyes.

Typically, the external sensors used to record human teacher executions are vision-based. Motion capture systems utilizing visual markers are applied to teaching human motion [37–39] and manipulation tasks [40]. Visual features are tracked in biologically-inspired frameworks that link perception to abstract knowledge representation [41] and teach an anthropomorphic hand to play the game Rock–Paper–Scissors [42]. Background subtraction is used to extract teacher motion from images [29]. This work takes a unique approach to building the embodiment mapping $g_E(z, a)$; the correspondence is *learned*, rather than being hand-engineered, by having the human teacher shadow the robot's motions.

A number of systems combine external sensors with other information sources; in particular, with sensors located directly on the teacher. A force-sensing glove is combined with vision-based motion tracking to teach grasping movements [43,44]. Movement, end effector position, stereo vision, and tactile information teaches generalized dexterous motor skills to a humanoid robot [45]. The combination of 3D marker data with torso movement and joint angle data is used for applications on a variety of simulated and robot humanoid platforms [46]. Other approaches combine speech with visual observation of human gestures [47] and object tracking [48], within the larger goal of speech-supported LfD learning.

Several works also explore learning through external observation of non-human teachers. A robot learns a manipulation task

through visual observation of an identical robotic teacher [49], and a simulated agent learns about its own capabilities and unvisited parts of the state space through observation of other simulated agents [50]. A generic framework for solving the correspondence problem between differently embodied robots allows a robotic agent to learn new behaviors through imitation of another, possibly physically different, agent [51].

3.4. Other approaches

Within LfD there do exist exceptions to the data source categorization we have so far presented. These exceptions record only states during demonstration, without recording actions. For example, by drawing a path through a 2D representation of the physical world, high-level path-planning demonstrations are provided to a rugged outdoor robot [52] and a small quadruped robot [53,54]. Since the dataset does not provide actions, no state–action mapping is learned for action selection. Instead, actions are selected at runtime by employing low-level motion planners and controllers [52,53], or by providing transition models $T(s'|s, a)$ [54].

4. Deriving a policy: The source of the state to action mapping

Given a dataset of state–action examples that have been acquired using one of the methods described in the previous section, we now discuss methods for deriving a policy using this data. LfD has seen the development of three core approaches to deriving policies from demonstration data, as summarized in Fig. 2. Learning a policy can involve simply learning an approximation to the state–action mapping (*mapping function*), or learning a model of the world dynamics and deriving a policy from this information (*system model*). Alternately, a sequence of actions can be produced by a planner after learning a model of action pre- and post-conditions (*plans*). Across all of these learning techniques, minimal parameter tuning and fast learning times requiring few training examples are desirable.

We introduce a full categorization of the various approaches to deriving a policy from the demonstration dataset in Fig. 6. Approaches are initially split between the three core derivation techniques described above. Further splits, if present, are approach-specific. We discuss each of these categories in depth in the following sections. Additionally, we direct the reader to a literature review that focuses on statistical and mathematical approaches to deriving motor control policies [3].

4.1. Mapping function

The *mapping function* approach to policy learning calculates a function that approximates the state to action mapping, $f() : Z \rightarrow A$, for the demonstrated behavior (Fig. 2a). The goal of this type of algorithm is to reproduce the underlying teacher policy, which is unknown, and to generalize over the set of available training examples such that valid solutions are also acquired for similar states that may not have been encountered during demonstration.

The details of function approximation are influenced by many factors. These include whether the state input and action output are continuous or discrete, whether the generalization technique uses the data to approximate a function prior to execution time or directly at execution time, whether it is feasible or desirable to keep the entire demonstration dataset around throughout learning, and whether the algorithm updates online.

In general, mapping approximation techniques fall into two categories depending on whether the prediction output of the algorithm is discrete or continuous. *Classification* techniques produce discrete output (Section 4.1.1), and *regression* techniques produce continuous output (Section 4.1.2). Many techniques for performing classification and regression have been developed outside of LfD and we refer the reader to [55] for a full discussion.

4.1.1. Classification

Classification approaches categorize their input into discrete classes, thereby grouping similar input values together. In the context of policy learning, the input to the classifier is robot states and the discrete output classes are robot actions. Below, we present a summary of classification methods applied at three action control levels (basic motion control, action primitives, complex behaviors), as defined in Section 2.2. Although a single classification algorithm can be applied at any level, we distinguish between them to highlight the generality of this learning technique.

Low-level robot actions include basic commands such as moving forward or turning. Example applications that learn a mapping from states to low-level actions include controlling a car within a simulated driving domain using *Gaussian Mixture Models (GMMs)* [56], flying a simulated airplane using *decision trees* [57] and learning obstacle avoidance and navigation behaviors using *Bayesian network* [11] and *k-Nearest Neighbors (kNN)* [58] classifiers.

When states are mapped to motion primitives, the primitives typically are then composed or sequenced together. For example, Pook and Ballard [8] classify primitive membership using kNN, and then recognize each primitive within the larger demonstrated task via *Hidden Markov Models (HMMs)*, to teach a robotic hand and arm an egg flipping manipulation task. HMMs and teach a basic assembly task [59], and motor-skill tasks by identifying and generalizing upon the intentions of the user [60]. A biologically-inspired framework automatically extracts primitives from demonstration data, classifies data via *vector-quantization* and then composes and/or sequences primitives within a hierarchical Neural Network [46]. The work is applied to a variety of test beds, including a 20 DoF simulated humanoid torso, a 37 DoF avatar (a simulated humanoid which responds to external sensors), Sony AIBO dogs and small differential drive Pioneer robots. Under this framework, the simulated humanoid torso is taught a variety of dance, aerobics and athletics motions [61], the avatar reaching patterns [38] and the Pioneer sequenced location visiting tasks [62].

Similar approaches have been used for the classification of high-level behaviors. The behaviors themselves are generally developed (by hand or learned) prior to task learning. A similarity measure on an eigenvector representation recognizes gesture behaviors for an anthropomorphic hand [44], and within this framework HMMs classify demonstrations into gestures for a box sorting task with a

Pioneer robot [63]. The Bayesian likelihood method selects actions for a humanoid robot in a button pressing task [64], and *Support Vector Machines (SVMs)* classify behaviors for a robotic ball sorting task [65].

4.1.2. Regression

Regression approaches map demonstration states to continuous action spaces. Similar to classification, the input to the regressor are robot states, and the *continuous* output are robot actions. Since the continuous-valued output often results from combining multiple demonstration set actions, typically regression approaches apply to low-level motions and not high-level behaviors. A key distinction between methods is whether the mapping function approximation occurs at run time, or prior to run time. Below we present a summary of regression techniques for LfD along a continuum between these two extremes.

At one extreme lies Lazy Learning [66], where function approximation does not occur until a current observation point in need of mapping is present. The simplest Lazy Learning technique is kNN, applied to action selection within robotic marble-maze [67] and simulated ball interception [22] domains. More complex approaches include *Locally Weighted Regression (LWR)* [68]. One LWR technique further anchors local functions to the phase of nonlinear oscillators [69] to produce rhythmic movements, specifically drumming [70] and walking [32] patterns with a humanoid robot. While Lazy Learning approaches are fast and expend no effort approximating the function in areas unvisited during execution time, they do require keeping around all of the training data.

In the middle of the data processing continuum lie techniques in which the original data is converted to another, possibly sparsified, representation prior to run time. This *converted* data is then used by Lazy Learning techniques at run time. For example, *Receptive Field Weighted Regression (RFWR)* [71] first converts demonstration data to a Gaussian and local linear model representation. *Locally Weighted Projection Regression (LWPR)* [72] extends this approach to scale with input data dimensionality and redundancy. Both RFWR and LWPR are able to incrementally update the number of representative Gaussians as well as regression parameters online. Successful robotics applications using LWPR include an AIBO robot performing basic soccer skills [23] and a humanoid playing the game of air hockey [67]. This latter work actually employs multiple techniques, first using kNN to select a behavior class and then LWPR to generate low-level actions using the demonstration data within this class. The approaches at this position on the continuum benefit from not needing to evaluate all of the training data at run time, but at the cost of extra computation and generalization prior to execution.

At the opposite extreme lie approaches which form a complete function approximation prior to execution time. At run time, they no longer depend on the presence of the underlying data (or any modified representations thereof). A seminal work uses a *Neural Network (NN)* to enable autonomous driving of a van at speed on a variety of road types [73]. NN techniques also enable a robot arm to perform a peg in hole task [49], and a set of b-spline wavelets approximate time-sequenced data for full body humanoid motion [39]. Also possible are statistical approaches that represent the demonstration data in a single or mixture distribution sampled at run time. Demonstration data encoded as a joint distribution over objects, hand position and hand orientation is used by a humanoid for grasping novel and known household objects [47]. *Gaussian Mixture Regression (GMR)* teaches a humanoid robot basketball referee signals [33], and *Sparse On-Line Gaussian Processes (SOGP)* teaches an AIBO robot to perform basic soccer skills [74]. The regression techniques in this area all have the advantage of *no* training data evaluations at run time,

but are the most computationally expensive prior to execution. Additionally, some of the techniques, for example NN, can suffer from “forgetting” mappings learned earlier in the learning run, a problem of particular importance when updating the policy online.

4.2. System models

The *system model* approach to LfD policy learning uses a state transition model of the world, $T(s'|s, a)$, and from this derives a policy $\pi : Z \rightarrow A$ (Fig. 2b). This approach is typically formulated within the structure of *Reinforcement Learning (RL)*. Demonstration data, and any additional autonomous exploration the robot may do, generally determines the transition function $T(s'|s, a)$. To derive a policy from this transition model, a reward function $R(s)$, which associates reward value r with world state s , is either learned from demonstrations or defined by the user.

The goal of RL is to maximize cumulative reward over time. The expected cumulative future reward of the state under the current policy is represented by further associating each state s with a value according to the function $V(s)$ (or associating state–action pair s, a with a Q -value according to the function $Q(s, a)$). State values may be represented by the *Bellman Equation*:

$$V^\pi(s) = \int_a \pi(s, a) \int_{s'} T(s'|s, a) [r(s) + \gamma V^\pi(s')] ds'$$

where $V^\pi(s)$ is the value of state s under policy π , and γ represents a discounting factor on future rewards. How to use these values to update a policy, and how to update them effectively online, is a subject of much research outside of LfD. Unlike function approximation techniques, RL approaches typically do *not* generalize state and demonstrations must be provided for every discrete state. For a full review of RL, we refer the reader to [1].

Reward is the motivation behind state and action selection, and consequently also guides task execution. Defining a reward function to accurately capture the desired task behavior, however, is not always obvious. We therefore categorize LfD implementations of the system model method by the source of the reward function. Some take the approach of an *engineered* reward function, as seen in classical RL implementations (Section 4.2.1). In other approaches, the reward function is *learned* from the demonstration data (Section 4.2.2).

4.2.1. Engineered reward functions

Within most applications of the LfD system model approach, the user manually defines the reward function. User-defined rewards tend to be sparse, meaning that the reward value is zero except for a few states, such as around obstacles or near the goal. Various demonstration-based techniques therefore have been defined to aid the robot in locating the rewards and to prevent extensive periods of blind exploration.

Demonstration may be used to highlight interesting areas of the state space in domains with sparse rewards, for example using teleoperation to show the robot the reward states and thus eliminating long periods of initial exploration that acquire no reward feedback [75]. Both reward and action demonstrations influence the performance of the learner in the supervised actor-critic reinforcement learning algorithm [1]. Applied to a basic assembly task using a robotic manipulator [24], reward feedback penalizes negative and reinforces positive behavior, while action demonstrations highlight recommended actions and suggest promising directions for state exploration.

A number of algorithms that use hand-engineered rewards also use interactions with other autonomous agents to enhance learning. However, these approaches have mostly been studied in simulation. In the Ask For Help framework [13], RL agents request advice from other similar agents when all possible actions

in a given state have relatively equal quality estimates. In a similarly motivated system [76], agents select, exchange and incorporate advice from other agents, and to combine this advice with information acquired locally through exploration, to speed up learning. Novice agents learn in a multiagent system [50] by passively observing expert agents in the environment, without explicit teaching.

Finally, real robot applications tend to employ RL in ways not typically seen in the classical RL policy derivation algorithms, since even in discrete state–action spaces the cost of visiting every state, and taking every action from each state, becomes prohibitively high and further could be physically dangerous. One approach uses simulation to seed an initial world model, for example using a planner operating on a simulated version of a marble maze robot [77]. The opposite approach derives a model of robot dynamics from demonstration but then uses the model in simulation, for example to simulate robot state when employing RL to optimize a NN controller for autonomous helicopter flight [78] and inverted helicopter hovering [6]. The former implementation additionally pairs the hand-engineered reward with a high penalty for visiting any state not encountered during demonstration, and the latter additionally makes the explicit assumption of a limited class of feasible controllers.

4.2.2. Learned reward functions

Defining an effective reward function in real world systems can be a non-trivial issue. *Inverse Reinforcement Learning* [79] is one subfield within RL that addresses this issue by *learning*, rather than hand-defining, the reward function. Within LfD, several algorithms examine learning a reward function from demonstrations, under a variety of system model frameworks.

A reward function learned by associating greater reward with states similar to those encountered during demonstration combines with RL techniques to teach a pendulum swing-up task to a robotic arm [35]. A modified RL algorithm learns task reward entirely based on individual rewards interactively provided by a human user [80]. An additional guidance component also enables the user to recommend actions for the robot to perform.

Rewarding similarity between teacher demonstrated and robot-executed trajectories is another approach to learning a reward function. Example approaches include a Bayesian model formulation [81], and another that further provides greater reward to positions close to the goal and thus enables adaptation to task changes like a new goal location or path obstacles [82]. A trajectory comparison based on state features first pioneered in [15] guarantees learning a correct *reward* function if the feature counts are properly estimated. This guarantee extends to learning the correct *behavior*, and on a real robot system implementation, in [52]. This work learns a mapping from observation input to state costs using margin-maximization techniques, and uses the resulting terrain cost map by a navigational planner to generate goal-achieving paths as sequences of states. This work extends to consider the decomposition of task demonstration into hierarchies for a small legged robot [54,53], and further to resolve ambiguities in feature counts and reward functions by employing the principle of maximum entropy [83]. All of these works learn the reward function $R(s)$ exclusively; however, [84] learns both a reward function *and* transition function $T(s'|s, a)$ for acrobatic helicopter maneuvers.

4.3. Plans

An alternative to mapping states directly to actions is to represent the desired robot behavior as a plan (Fig. 2c). The planning framework, represents the policy as a sequence of actions that lead from the initial state to the final goal state.

Actions are often defined in terms of *pre-conditions*, the state that must be established before the action can be performed, and *post-conditions*, the state resulting from the action's execution. Unlike other LfD approaches, planning techniques frequently rely not only on state–action demonstrations, but also on additional information in the form of *annotations* or *intentions* from the teacher. Demonstration-based algorithms differ in how the rules associating pre- and post-conditions with actions are learned, and whether additional information is provided by the teacher.

One of the first papers to learn execution plans based on demonstration learns a plan for object manipulation based on observations of the teacher's hand movements [85]. A humanoid learns a generalized plan, able to represent sequential tasks with repetitions, from only two demonstrations of a repetitive ball collection task [86]. Other approaches use spoken dialog both as a technique to demonstrate task plans, and also to enable the robot to verify unspecified or unsound parts of a plan through dialog. Under such a framework, multiple teachers provide verbal navigation instructions, with the intention of building a vocabulary that maps to sensorimotor primitives [20], and a single teacher presents a robot with a series of conditional statements that are then processed into a plan [21].

Other planning-based methods require teacher annotations. Overall task intentions provided in response to learner queries are used to encode plan post-conditions [87], and human feedback draws attention to particular elements of the domain when learning a high-level plan from shadowing [88]. Goal annotations include indicating the demonstrator's goal and time points at which it is achieved or abandoned when learning non-deterministic plans [89], and providing binary feedback that grows or prunes the goal set used to infer the demonstrator's goal [90]. Annotations of the demonstrated action sequence combine with high-level instructions, for example that some actions can occur in any order, when learning a domain-specific hierarchical task model [91].

5. Limitations of the demonstration dataset

LfD systems are inherently linked to the information provided in the demonstration dataset. As a result, learner performance is heavily limited by the quality of this information. In this article we identify two distinct causes for poor learner performance within LfD frameworks and survey the techniques that have been developed to address each limitation. The first cause, discussed in Section 5.1, is due to dataset sparsity, or the existence of areas of the state space that have not been demonstrated. The second, discussed in Section 5.2, is due to poor quality of the dataset examples, which can result from a teacher's inability to perform the task optimally.

5.1. Undemonstrated state

The availability of training data limits all LfD policies because, in all but the most simple domains, the teacher is unable to demonstrate the correct action from every possible state. This raises the question of how the robot should act when it encounters a state without a demonstration. We categorize existing methods for dealing with undemonstrated state into two approaches: *generalization from existing demonstrations*, and *acquisition of new demonstrations*.

5.1.1. Generalization from existing demonstrations

Here we discuss methods for using the existing data to deal with undemonstrated state. We present techniques as used within each of the core approaches for policy derivation.

Within the function mapping approach, state generalization deals with undemonstrated state. Generalizing across inputs is a feature inherent to robust function approximation. The exact nature of this generalization depends upon the specific classification or regression technique used. Approaches range from strict grouping with the nearest dataset point state, as in kNN, to soft averaging across multiple states, as in kernelized regression.

Within the system model approach, state exploration addresses undemonstrated state. State exploration is often accomplished by providing the learner with an exploration policy that guides action selection within undemonstrated states. Rewards provided by the world, automatically evaluate the performances of these actions. This motivates the issue of *Exploration vs. Exploitation*; that is, of determining how frequently to take exploratory actions versus following the set policy. We note that taking exploratory steps on a real robot system is often unwise, for safety and stability reasons. A safe exploration policy is employed for the purposes of gathering robot demonstration data, both to seed and then update a transition model and reward function [12]. To guide exploration within this work, initially a suboptimal controller is provided, and once a policy is learned small amounts of Gaussian noise are randomly added to the greedy actions of this policy.

Generalization to unseen states is not a common feature amongst traditional planning algorithms. This is in large part due to the common assumption that every action has a set of known, deterministic effects on the environment that lead to a particular known state. However, this assumption is typically dropped in real-world applications, and among LfD approaches several algorithms do explore the generalization of demonstrated sequences. For example, the teacher specifies the level of generality that was implied by a demonstration by answering the learner's questions (e.g. should the box be picked up only from the chair, or from any surface besides the table?) [87]. Other approaches focus on the development of generalized plans with flexible action order [88,89].

5.1.2. Acquisition of new demonstrations

A fundamentally different approach to dealing with undemonstrated state is to re-engage the teacher and acquire additional demonstrations when novel states are encountered by the robot. Under such a paradigm, the responsibility of selecting states for demonstration now is shared between the robot and teacher.

One set of approaches acquire new demonstrations by enabling the learner to evaluate its confidence in selecting a particular action, based on the confidence of the underlying classification algorithm. A robot requests additional demonstration in states very different from previously demonstrated states or in which a single action can not be selected with certainty [16,56], and a teacher chooses to provide additional demonstrations based on the robot indicating its certainty in performing various elements of the task [23].

5.2. Poor quality data

The quality of a learned LfD policy depends heavily on the quality of the provided demonstrations. In general, approaches assume the dataset to contain high quality demonstrations performed by an expert. In reality, however, teacher demonstrations may be ambiguous, unsuccessful or suboptimal in certain areas of the state space. A naively learned policy will likely perform poorly in such areas, [35]. In this section we discuss proposed approaches for improving demonstration data and dealing with poor learner performance.

5.2.1. Suboptimal or ambiguous demonstrations

One approach to dealing with inefficient demonstrations is to eliminate parts of the teacher's execution that are suboptimal. Other approaches address the ambiguities that can result from single or multiple demonstrations.

Suboptimality is most common in demonstrations of low-level tasks, such as movement trajectories for robotic arms. In these domains, the demonstration serves as guidance instead of offering a complete solution. Removing actions that do not contribute to the solution of the task allows for the active identification and removal of elements of the demonstration that are unnecessary or inefficient [92]. Other approaches smooth or generalize over suboptimal demonstrations in such a way as to improve upon the teacher's performance. For example, when learning motion control, data from multiple repeated demonstrations is used to obtain a smoothed optimal path [34,93,94], and data from multiple teachers encourages more robust generalization [8].

Uncaptured elements of the record mapping may cause dataset ambiguity. Differences in teacher and learner perspectives during externally recorded demonstrations are accounted for by having the robot maintain separate belief estimates for itself and the human demonstrator [19]. Teacher demonstrations that are inconsistent between multiple executions also may cause dataset ambiguity. This type of ambiguity arises due to the presence of multiple equally applicable actions (e.g. both turning left and moving forward are equally valid actions) among which the teacher selects arbitrarily, or due to state corruption through sensor noise. The result in both cases is that identical, or nearly identical, states are mapped to different actions in the demonstration. Approaches designed to address these inconsistencies include requests for additional clarification demonstrations from the teacher [56] and modeling the choice of actions explicitly in the robot policy [95].

5.2.2. Learning from experience

An altogether different approach to dealing with poor quality data has the student to learn from experience. If the learner is provided with feedback that evaluates its performance, this may be used to update its policy. This evaluation is generally provided via teacher feedback, or through a reward as in RL.

Automatic policy evaluations enable a policy *derived* under the mapping function approximation approach to be *updated* using RL techniques [67]. The function approximation of this work considers Q -values associated with paired query-demonstration points, and these Q -values are updated based on learner executions and a developer-defined reward function. Binary policy evaluations from a human teacher prune a set of inferred teacher goals [90], and human-provided binary tags flag portions of the learner execution for poor performance enable updating a continuous-action mapping function approximation [22].

Most use of RL-type rewards, however, occurs under the system model approach. We emphasize that these are rewards seen during *learner* execution with its policy, and not during teacher demonstration executions. [12] seeds the robot's reward and transition functions with a small number of suboptimal demonstrations, and then optimizes the policy through exploration and RL. A similar approach is taken learning a motor primitives for a 7-DoF robotic arm with the Natural Actor Critic algorithm [96]. [77] seeds the world model is seeded with trajectories produced by a planner operating on a simulated version of their robot; learner execution results in state values being updated, and consequently also the policy. In addition to updating state values, student execution information may be used to update the actual learned dynamics model $T(s'|s, a)$, as in work that seeds the world model with teacher demonstrations [97]. Policy updates in this work consist of rederiving the world dynamics after adding learner executions with this policy (and later updated policies) to the demonstration set.

Beyond providing just an evaluation of performance, teacher feedback may also provide a *correction* on the executed behavior. The correct discrete action provided by a human teacher updates the structure of a hierarchical Neural Network of robot behaviors [88] and an action classifier [95]. Further challenging is for a human to provide corrections in continuous action spaces. Kinesthetic teaching, where passive humanoid joints are moved by the teacher through the desired motions is one approach that provides continuous-valued corrections [33]. Another approach has a set of mathematical operations transform high-level advice from a human teacher into continuous-valued corrections on students executions [98].

Before concluding, we underline the advantage of providing demonstrations and employing LfD *before* using traditional policy learning techniques, such as exploration-based methods like RL. As mentioned in Section 4.2.1, many traditional learning approaches are not immediately applicable to robots. This is largely due to the constraints of gathering experience on a real system. To physically execute *all* actions from *every* state is likely infeasible, may be dangerous, and will not scale with continuous state spaces. However, the use of these techniques to evaluate and improve upon the experiences gained through an LfD system has proven very successful. We consider this to be an area with much potential for future research, which we discuss further within Section 6.4.

6. Future directions

As highlighted by the discussion in the previous sections, current approaches to LfD address a wide variety of problems under many different conditions and assumptions. In this section, we aim to highlight several promising areas of LfD research that have received limited attention, ranging from data representation to issues of system robustness and evaluation metrics.

6.1. Learned state features

Within all policy learning approaches, state information is typically represented by a set of *state features* that describe various attributes of the environment. Before learning begins, a developer or user determines the set of features available for learning. A robot learning a task knows nothing about the world beyond the provided features. In selecting features, developers face a decision between a larger number of features, which provide a richer representation and allow for learning a wider range of tasks, and a smaller number of features, which typically results in faster learning times with less training data.

Most policy learning algorithms are typically provided with exactly the data needed to learn a task, nothing more and nothing less. Some approaches, however, do deal with extra features by generalizing over irrelevant features [23] or filtering them out through feature selection [99].

In addition to facing these challenges, we propose that future LfD algorithms are likely to face a new problem—*insufficient features*. Unlike standard robotics approaches, LfD does not require extensive programming experience but rather only the ability to demonstrate the task in the chosen manner. Due to the intuitive nature of demonstration, LfD algorithms have the potential to make robot programming accessible for everyday users who have no programming experience and want to customize robot behaviors. The use of current approaches by the general public, however, is likely to lead to situations in which users attempt to teach the robot tasks that cannot be learned simply due to the lack of features describing some aspect of the task. An open area of future work is to address this shortcoming by identifying intuitive methods for introducing new features into the learning system.

6.2. Temporal data

Actions within almost all robot behaviors have a temporal ordering, a natural progression in which actions of the task are performed. Many policy learning algorithms, especially within the area of planning, naturally encode the temporal ordering. However, the vast majority of LfD algorithms developed to date result in a purely reactive policy that maps directly from a state to an action without considering past history or events. Reactive approaches have repeatedly been shown to achieve good performance, however they suffer from a number of drawbacks. By discarding all temporal information from the training data, such algorithms have a difficulty representing repetitive tasks, such as ones in which an action sequence must be repeated some fixed number of times. Additionally, learning becomes difficult in the presence of actions which have no perceivable effect on the environment and that therefore do not change the state of the robot. Both of these challenges can be solved by including special state features that represent the memory of recent events of this nature; for example, anchoring demonstrations of repetitive rhythmic tasks to the phase of a nonlinear oscillator [70]. However, such features are always task specific and do not present a general solution to this problem. Encoding temporal information into the demonstration sequence is one approach that may achieve a more universal solution. The development of such generalized LfD policy learning techniques is an open area for future research.

6.3. Execution failures

With any LfD approach, it is possible that the policy will be unable to execute as intended by the demonstration teacher. This could be, for example, because the state representation does not model crucial aspects of the world, or because the policy failed to generalize properly from the demonstration data. All LfD systems run the risk of encountering failures. The ability of a system to identify and address failures and their causes is both extremely valuable and extremely non-trivial.

A small number of works do directly address policy failures. Policy derivation failure is considered by identifying when a HMM fails to classify a demonstration and requesting that the human teacher re-demonstrate the task [63]. Other work considers policy execution failures: a robot learner recognizes when it is unable to complete its task (for example, due to a physical obstruction) and solicits the help of a human [62]. Though in this work the robot identifies task failure, it does not identify the failure cause or modify its policy; the intent is that human aid will enable task completion (for example, remove the obstruction).

6.4. New techniques for learning from experience

The ability to learn from experience is a desirable trait in any learning system. The most popular approaches implemented within LfD systems to date offer reward-type evaluations of policy performance (Section 5.2.2). Reward functions are often sparse, however, and only offer an indication of how desirable being in a given state is; reward gives no indication of which action would have been more appropriate to select instead, for example. Richer forms of performance evaluation, however, could benefit approaches able to incorporate policy performance feedback.

There is no reason to expect that formalizing these richer evaluations is any easier than the admittedly challenging task of formally defining reward functions (Section 4.2.2); if anything, it is possible that a richer evaluation is even more difficult to capture. One promising solution has the demonstration teacher provide feedback on learner performance, in addition to providing the example executions. Demonstration teachers presumably possess

a measure of task expertise, and so at the very least have existing measures by which to evaluate their own performance on the task. Furthermore, metrics difficult to formalize or parameterize but easy for a human to evaluate are available if the teacher is human. Steps in this direction utilize human teacher evaluations to provide corrections on discrete action selection [16,88] to correct state sequencing [54]. A human teacher also corrects continuous action selection [98] and penalizes poorly performing demonstration data [22]. Exactly how to provide effective feedback, and then how to use this information, is for the most part still an area open for future research.

6.5. Multi-robot demonstration learning

LfD approaches to date focus almost exclusively on the problem of a single robot being taught by a single teacher. However, solutions to complex tasks often require the cooperation of multiple robots. Multi-robot coordination has been extensively studied in robotics research and presents many challenges, such as issues of action coordination, communication, noise in shared information, and physical interaction. Current LfD research is beginning to take steps in this direction. Multiple autonomous agents request advice and provide demonstrations for each other [76], and agents with both similar and dis-similar embodiments to learn movements from each other through imitation [51,100]. A confidence-based algorithm enables a single person to teach multiple robots at the same time through the demonstration of physical and communication actions required to perform the task [65]. These studies only begin to explore the various applications of LfD in the multi-robot setting, and many directions remain open for future work.

6.6. Evaluation metrics

LfD is a relatively young but rapidly growing field. As highlighted by this survey, a wide variety of approaches address the challenges presented by this learning method. However, to date there exists little direct comparison between algorithms.

One reason for this absence of direct comparisons is that most approaches are tested using only a single domain and robotic platform. As a result, such techniques are often customized to that particular domain and do not present a general solution to a wide class of problems. Additionally, the field of LfD research currently lacks a standard set of evaluation metrics, further complicating performance comparisons across algorithms and domains.

Improved methods for evaluation and comparison are a fundamentally important area for future work. Existing evaluation methods include several proposed for specific areas of LfD [101–104], and potentially some from the Human–Robot Interaction (HRI) community [105]. The formalization of evaluation criteria would help to drive research and the development of widely applicable general-purpose learning techniques.

7. Conclusion

In this article we have presented a comprehensive survey of *Learning from Demonstration (LfD)* techniques employed to address the robotics control problem. LfD has the attractive characteristics of being an intuitive communication medium for human teachers and of opening control algorithm development to non-robotics-experts. Additionally, LfD complements many traditional policy learning techniques, offering a solution to some of the weaknesses in traditional approaches. Consequently, LfD has been successfully applied to many robotics applications.

The LfD robotics community, however, suffers from the lack of an established structure in which to organize approaches. In

this survey, we have contributed such a structure, through a categorization of LfD techniques. We first segment the LfD learning problem into two fundamental phases. The first phase consists of *gathering* the demonstration examples. Within this phase we categorize according to how the demonstration is *executed* and how the demonstration is *recorded*. The second phase consists of *deriving* a policy from the demonstration examples. Within this phase we categorize according to what is learned: an approximation to the state–action mapping (*mapping function*), a model of the world dynamics and/or reward (*system model*), or a model of action pre- and post-conditions (*plans*). We have further contributed a categorization of current approaches into this structure.

Though LfD has proven a successful tool for robot policy development, there still exist many open areas for research, several of which we have identified. In particular, learning state features and considering the temporal nature of demonstration data are interesting directions to explore within the context of building a dataset and policy. System robustness could increase by explicitly dealing with policy derivation and execution failures. New approaches to learning from experience, that are specific to LfD systems, hold much promise for improving policy performance, and LfD extends naturally into multi-robot domains, which is an intriguing area that remains largely open. Finally, a standardized set of evaluation metrics would facilitate comparisons between approaches and implementations, to the benefit of the LfD community as a whole.

Acknowledgements

This research is partly sponsored by the Boeing Corporation under Grant No. CMU-BA-GTA-1, BBNT Solutions under subcontract No. 950008572, via prime Air Force contract No. SA-8650-06-C-7606, and the Qatar Foundation for Education, Science and Community Development. The views and conclusions contained in this document are solely those of the authors.

The authors would like to thank J. Andrew Bagnell and Darrin Bentivegna for feedback on the content and scope of this article.

References

- [1] R.S. Sutton, A.G. Barto, Reinforcement learning: An introduction, The MIT Press, Cambridge, MA, London, England, 1998.
- [2] A. Billard, S. Callinon, R. Dillmann, S. Schaal, Robot programming by demonstration, in: B. Siciliano, O. Khatib (Eds.), Handbook of Robotics, Springer, New York, NY, USA, 2008 (Chapter 59).
- [3] S. Schaal, A. Ijspeert, A. Billard, Computational approaches to motor learning by imitation, Philosophical Transactions: Biological Sciences 358 (1431) (2003) 537–547.
- [4] K. Dautenhahn, C.L. Nehaniv (Eds.), Imitation in Animals and Artifacts, MIT Press, Cambridge, MA, USA, 2002.
- [5] C.L. Nehaniv, K. Dautenhahn, in: K. Dautenhahn, C.L. Nehaniv (Eds.), The Correspondence Problem, MIT Press, Cambridge, MA, USA, 2002 (Chapter 2).
- [6] A. Ng, A. Coates, M. Diel, V. Ganapathi, J. Schulte, B. Tse, E. Berger, E. Liang, Inverted autonomous helicopter flight via reinforcement learning, in: International Symposium on Experimental Robotics, 2004.
- [7] B. Browning, L. Xu, M. Veloso, Skill acquisition and use for a dynamically-balancing soccer robot, in: Proceeding of the 19th National Conference on Artificial Intelligence, AAAI'04, 2004.
- [8] P.K. Pook, D.H. Ballard, Recognizing teleoperated manipulations, in: Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'93, 1993.
- [9] J.D. Sweeney, R.A. Grupen, A model of shared grasp affordances from demonstration, in: Proceedings of the IEEE-RAS International Conference on Humanoids Robots, 2007.
- [10] J. Chen, A. Zelinsky, Programming by demonstration: Coping with suboptimal teaching actions, The International Journal of Robotics Research 22 (5) (2003) 299–319.
- [11] T. Inamura, M. Inaba, H. Inoue, Acquisition of probabilistic behavior decision model based on the interactive teaching method, in: Proceedings of the Ninth International Conference on Advanced Robotics, ICAR'99, 1999.
- [12] W.D. Smart, Making Reinforcement Learning Work on Real Robots. Ph.D. Thesis, Department of Computer Science, Brown University, Providence, RI, 2002.
- [13] J.A. Clouse, On integrating apprentice learning and reinforcement learning. Ph.D. Thesis, University of Massachusetts, Department of Computer Science, 1996.
- [14] R.P.N. Rao, A.P. Shon, A.N. Meltzoff, A bayesian model of imitation in infants and robots, in: K. Dautenhahn, C. Nehaniv (Eds.), Imitation and Social Learning in Robots, Humans, and Animals: Behavioural, Social and Communicative Dimensions, Cambridge University Press, Cambridge, UK, 2004 (Chapter 11).
- [15] P. Abbeel, A.Y. Ng, Apprenticeship learning via inverse reinforcement learning, in: Proceedings of the 21st International Conference on Machine Learning, ICML'04, 2004.
- [16] S. Chernova, M. Veloso, Multi-thresholded approach to demonstration selection for interactive robot learning, in: Proceedings of the 3rd ACM/IEEE International Conference on Human–Robot Interaction, HRI'08, 2008.
- [17] R. Aler, O. Garcia, J.M. Valls, Correcting and improving imitation models of humans for robosoccer agents, Evolutionary Computation 3 (2–5) (2005) 2402–2409.
- [18] A.G. Billard, S. Calinon, F. Guenter, Discriminative and adaptive imitation in uni-manual and bi-manual tasks, in: The Social Mechanisms of Robot Programming by Demonstration, Robotics and Autonomous Systems 54 (5) (2006) 370–384 (special issue).
- [19] C. Breazeal, M. Berlin, A. Brooks, J. Gray, A.L. Thomaz, Using perspective taking to learn from ambiguous demonstrations, in: The Social Mechanisms of Robot Programming by Demonstration, Robotics and Autonomous Systems 54 (5) (2006) 385–393 (special issue).
- [20] S. Lauria, G. Bugmann, T. Kyriacou, E. Klein, Mobile robot programming using natural language, Robotics and Autonomous Systems 38 (2002) 171–181.
- [21] P.E. Rybski, K. Yoon, J. Stolarz, M.M. Veloso, Interactive robot task training through dialog and demonstration, in: Proceedings of the 2nd ACM/IEEE International Conference on Human–Robot Interactions, HRI'07, 2007.
- [22] B. Argall, B. Browning, M. Veloso, Learning from demonstration with the critique of a human teacher, in: Proceedings of the 2nd ACM/IEEE International Conference on Human–Robot Interactions, HRI'07, 2007.
- [23] D.H. Grollman, O.C. Jenkins, Dogged learning for robots, in: Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'07, 2007.
- [24] M.T. Rosenstein, A.G. Barto, Supervised actor-critic reinforcement learning, in: J. Si, A. Barto, W. Powell, D. Wunsch (Eds.), Learning and Approximate Dynamic Programming: Scaling Up to the Real World, John Wiley & Sons, Inc., New York, NY, USA, 2004 (Chapter 14).
- [25] G.Z. Grudic, P.D. Lawrence, Human-to-robot skill transfer using the spore approximation, in: Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'96, 1996.
- [26] Y. Demiris, G. Hayes, in: K. Dautenhahn, C.L. Nehaniv (Eds.), Imitation as a dual-route process featuring predictive and learning components: A biologically-plausible computational model, MIT Press, Cambridge, MA, USA, 2002 (Chapter 13).
- [27] M.N. Nicolescu, M.J. Mataric, Experience-based representation construction: Learning from human and robot teachers, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'01, 2001.
- [28] U. Nehmzow, O. Akanyeti, C. Weinrich, T. Kyriacou, S. Billings, Robot programming by demonstration through system identification, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'07, 2007.
- [29] M. Ogino, H. Toichi, Y. Yoshikawa, M. Asada, Interaction rule learning with a human partner based on an imitation faculty with a simple visuomotor mapping, in: The Social Mechanisms of Robot Programming by Demonstration, Robotics and Autonomous Systems 54 (5) (2006) 414–418 (special issue).
- [30] C. Breazeal, B. Scassellati, Robots that imitate humans, Trends in Cognitive Sciences 6 (11) (2002) 481–487.
- [31] A.J. Ijspeert, J. Nakanishi, S. Schaal, Movement imitation with nonlinear dynamical systems in humanoid robots, in: Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'02, 2002.
- [32] J. Nakanishi, J. Morimoto, G. Endo, G. Cheng, S. Schaal, M. Kawato, Learning from demonstration and adaptation of biped locomotion, Robotics and Autonomous Systems 47 (2004) 79–91.
- [33] S. Calinon, A. Billard, Incremental learning of gestures by imitation in a humanoid robot, in: Proceedings of the 2nd ACM/IEEE International Conference on Human–Robot Interactions, HRI'07, 2007.
- [34] J. Aleotti, S. Caselli, Robust trajectory learning and approximation for robot programming by demonstration, in: The Social Mechanisms of Robot Programming by Demonstration, Robotics and Autonomous Systems 54 (5) (2006) 409–413 (special issue).
- [35] C.G. Atkeson, S. Schaal, Robot learning from demonstration, in: Proceedings of the Fourteenth International Conference on Machine Learning, ICML'97, 1997.
- [36] D.C. Bentivegna, A. Ude, C.G. Atkeson, G. Cheng, Humanoid robot learning and game playing using PC-based vision, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'02, 2002.
- [37] R. Amit, M. Mataric, Learning movement sequences from demonstration, in: Proceedings of the 2nd International Conference on Development and Learning, ICDL'02, 2002.
- [38] A. Billard, M. Mataric, Learning human arm movements by imitation: Evaluation of biologically inspired connectionist architecture, Robotics and Autonomous Systems 37 (2–3) (2001) 145–160.

- [39] A. Ude, C.G. Atkeson, M. Riley, Programming full-body movements for humanoid robots by observation, *Robotics and Autonomous Systems* 47 (2004) 93–108.
- [40] N. Pollard, J.K. Hodgins, Generalizing demonstrated manipulation tasks, in: *Workshop on the Algorithmic Foundations of Robotics, WAFR'02*, 2002.
- [41] A. Chella, H. Dindo, I. Infantino, A cognitive framework for imitation learning, in: *The Social Mechanisms of Robot Programming by Demonstration*, *Robotics and Autonomous Systems* 54 (5) (2006) 403–408 (special issue).
- [42] I. Infantino, A. Chella, H. Dindo, I. Macaluso, A posture sequence learning system for an anthropomorphic robotic hand, *Robotics and Autonomous Systems* 42 (2004) 143–152.
- [43] M. Lopes, J. Santos-Victor, Visual learning by imitation with motor representations, *IEEE Transactions on Systems, Man, and Cybernetics, Part B* 35 (3) (2005) 438–449.
- [44] R.M. Voyles, P.K. Khosla, A multi-agent system for programming robots by human demonstration, *Integrated Computer-Aided Engineering* 8 (1) (2001) 59–67.
- [45] J. Lieberman, C. Breazeal, Improvements on action parsing and action interpolation for learning through demonstration, in: *4th IEEE/RAS International Conference on Humanoid Robots*, 2004.
- [46] M.J. Matarić, in: K. Dautenhahn, C.L. Nehaniv (Eds.), *Sensory-motor primitives as a basis for learning by imitation: Linking perception to action and biology to robotics*, MIT Press, Cambridge, MA, USA, 2002 (Chapter 15).
- [47] J. Steil, F. Rothling, R. Haschke, H. Ritter, Situated robot learning for multi-modal instruction and imitation of grasping, in: *Robot Learning by Demonstration*, *Robotics and Autonomous Systems* 2–3 (47) (2004) 129–141 (special issue).
- [48] Y. Demiris, B. Khadhour, Hierarchical attentive multiple models for execution and recognition of actions, in: *The Social Mechanisms of Robot Programming by Demonstration*, *Robotics and Autonomous Systems* 54 (5) (2006) 361–369 (special issue).
- [49] R. Dillmann, M. Kaiser, A. Ude, Acquisition of elementary robot skills from human demonstration, *International Symposium on Intelligent Robotic Systems, SIRS'95*, 1995.
- [50] B. Price, C. Boutilier, Accelerating reinforcement learning through implicit imitation, *Journal of Artificial Intelligence Research* 19 (2003) 569–629.
- [51] A. Alissandrakis, C.L. Nehaniv, K. Dautenhahn, Towards robot cultures?, *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems* 5 (1) (2004) 3–44.
- [52] N. Ratliff, J.A. Bagnell, M.A. Zinkevich, Maximum margin planning, in: *Proceedings of the 23rd International Conference on Machine Learning, ICML'06*, 2006.
- [53] N. Ratliff, D. Bradley, J.A. Bagnell, J. Chestnutt, Boosting structured prediction for imitation learning, in: *Proceedings of the Advances in Neural Information Processing Systems, NIPS'07*, 2007.
- [54] J.Z. Kolter, P. Abbeel, A.Y. Ng, Hierarchical apprenticeship learning with application to quadruped locomotion, in: *Proceedings of the Advances in Neural Information Processing, NIPS'08*, 2008.
- [55] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer, New York, NY, USA, 2001.
- [56] S. Chernova, M. Veloso, Confidence-based learning from demonstration using Gaussian Mixture Models, in: *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems, AAMAS'07*, 2007.
- [57] C. Sammut, S. Hurst, D. Kedzier, D. Michie, Learning to fly, in: *Proceedings of the Ninth International Workshop on Machine Learning*, 1992.
- [58] J. Saunders, C.L. Nehaniv, K. Dautenhahn, Teaching robots by moulding behavior and scaffolding the environment, in: *Proceedings of the 1st ACM/IEEE International Conference on Human–Robot Interactions, HRI'06*, 2006.
- [59] G. Hovland, P. Sikka, B. McCarragher, Skill acquisition from human demonstration using a hidden markov model, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'96*, 1996.
- [60] K.R. Dixon, P.K. Khosla, Learning by observation with mobile robots: A computational approach, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'04*, 2004.
- [61] O.C. Jenkins, M.J. Matarić, S. Weber, Primitive-based movement classification for humanoid imitation, in: *Proceedings of the 1st IEEE-RAS International Conference on Humanoid Robots*, 2000.
- [62] M. Nicollescu, M.J. Matarić, Learning and interacting in human-robot domains, *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans* 31 (5) (2001) 419–430 (special issue).
- [63] P.E. Rybski, R.M. Voyles, Interactive task training of a mobile robot through human gesture recognition, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'99*, 1999.
- [64] A. Lockerd, C. Breazeal, Tutelage and socially guided robot learning, in: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'04*, 2004.
- [65] S. Chernova, M. Veloso, Teaching multi-robot coordination using demonstration of communication and state sharing (short paper), in: *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems, AAMAS'08*, 2008.
- [66] C.G. Atkeson, A.W. Moore, S. Schaal, Locally weighted learning, *Artificial Intelligence Review* 11 (1997) 11–73.
- [67] D.C. Bentivegna, Learning from Observation Using Primitives. Ph.D. Thesis, College of Computing, Georgia Institute of Technology, Atlanta, GA, July 2004.
- [68] W.S. Cleveland, C. Loader, *Smoothing by local regression: Principles and methods*. Technical Report, AT & T Bell Laboratories, Murray Hill, NJ, 1995.
- [69] S. Schaal, D. Sternad, Programmable pattern generators, in: *3rd International Conference on Computational Intelligence in Neuroscience*, 1998.
- [70] A.J. Ijspeert, J. Nakanishi, S. Schaal, Learning rhythmic movements by demonstration using nonlinear oscillators, in: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'02*, 2002.
- [71] S. Schaal, C.G. Atkeson, Constructive incremental learning from only local information, *Neural Computation* 10 (8) (1998) 2047–2084.
- [72] S. Vijayakumar, S. Schaal, Locally weighted projection regression: An $o(n)$ algorithm for incremental real time learning in high dimensional space, in: *Proceedings of the 17th International Conference on Machine Learning, ICML'00*, 2000.
- [73] D. Pomerleau, Efficient training of artificial neural networks for autonomous navigation, *Neural Computation* 3 (1) (1991) 88–97.
- [74] D.H. Grollman, O.C. Jenkins, Sparse incremental learning for interactive robot control policy estimation, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'08*, 2008.
- [75] W.D. Smart, L.P. Kaelbling, Effective reinforcement learning for mobile robots, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'02*, 2002.
- [76] E. Oliveira, L. Nunes, Learning by exchanging advice, in: R. Khosla, N. Ichalkaranje, L. Jain (Eds.), *Design of Intelligent Multi-Agent Systems*, Springer, New York, NY, USA, 2004 (Chapter 9).
- [77] M. Stolle, C.G. Atkeson, Knowledge transfer using local features, in: *Proceedings of the IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning, ADPRL'07*, 2007.
- [78] J.A. Bagnell, J.G. Schneider, Autonomous helicopter control using reinforcement learning policy search methods, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'01*, 2001.
- [79] S. Russell, Learning agents for uncertain environments (extended abstract), in: *Proceedings of the eleventh annual conference on Computational learning theory, COLT'98*, 1998.
- [80] A.L. Thomaz, C. Breazeal, Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance, in: *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI'06)*, 2006.
- [81] M. Ollis, W.H. Huang, M. Happold, A bayesian approach to imitation learning for robot navigation, in: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'07*, 2007.
- [82] F. Guenter, A. Billard, Using reinforcement learning to adapt an imitation task, in: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'07*, 2007.
- [83] B.D. Ziebart, A. Maas, J.D. Bagnell, A.K. Dey, Maximum entropy inverse reinforcement learning, in: *Proceedings of the 23rd National Conference on Artificial Intelligence, AAAI'08*, 2008.
- [84] P. Abbeel, A. Coates, M. Quigley, A.Y. Ng, An application of reinforcement learning to aerobatic helicopter flight, in: *Proceedings of the Advances in Neural Information Processing, NIPS'07*, 2007.
- [85] Y. Kuniyoshi, M. Inaba, H. Inoue, Learning by watching: Extracting reusable task knowledge from visual observation of human performance, *IEEE Transactions on Robotics and Automation* 10 (1994) 799–822.
- [86] H. Veeraraghavan, M. Veloso, Teaching sequential tasks with repetition through demonstration (short paper), in: *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems, AAMAS'08*, 2008.
- [87] H. Friedrich, R. Dillmann, Robot programming based on a single demonstration and user intentions, in: *3rd European Workshop on Learning Robots at ECML'95*, 1995.
- [88] M.N. Nicollescu, M.J. Matarić, Methods for robot task learning: Demonstrations, generalization and practice, in: *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems, AAMAS'03*, 2003.
- [89] M. van Lent, J.E. Laird, Learning procedural knowledge through observation, in: *Proceedings of the 1st international conference on Knowledge capture, K-CAP'01*, 2001.
- [90] B. Jansen, H. Belpaeme, A computational model of intention reading in imitation, in: *The Social Mechanisms of Robot Programming by Demonstration*, *Robotics and Autonomous Systems* 54 (5) (2006) 394–402 (special issue).
- [91] A. Garland, N. Lesh, *Learning Hierarchical TaskModels by Demonstration*. Technical Report TR2003-01, Mitsubishi Electric Research Laboratories, 2003.
- [92] M. Kaiser, H. Friedrich, R. Dillmann, Obtaining good performance from a bad teacher, in: *Programming by Demonstration vs. Learning from Examples Workshop at ML'95*, 1995.
- [93] N. Delson, H. West, Robot programming by human demonstration: Adaptation and inconsistency in constrained motion, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA'96*, 1996.
- [94] M. Yeasin, S. Chaudhuri, Toward automatic robot programming: Learning human skill from visual data, *IEEE Transactions on Systems, Man and Cybernetics – Part B: Cybernetics* 30 (2000).
- [95] S. Chernova, M. Veloso, Learning equivalent action choices from demonstration, in: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'08*, 2008.
- [96] J. Peters, S. Schaal, Natural actor-critic, *Neurocomputing* 71 (7–9) (2008) 1180–1190.

- [97] P. Abbeel, A.Y. Ng, Exploration and apprenticeship learning in reinforcement learning, in: Proceedings of the 22nd International Conference on Machine Learning, ICML'05, 2005.
- [98] B. Argall, B. Browning, M. Veloso, Learning robot motion control with demonstration and advice-operators, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS'08, 2008.
- [99] I. Guyon, A. Elisseeff, An introduction to variable and feature selection, *Journal of Machine Learning Research* 3 (2003) 1157–1182.
- [100] A. Alissandrakis, C.L. Nehaniv, K. Dautenhahn, Correspondence mapping induced state and action metrics for robotic imitation, *IEEE Transactions on Systems, Man and Cybernetics – Part B: Cybernetics* 37 (2) (2007) 299–307.
- [101] A. Alissandrakis, C.L. Nehaniv, K. Dautenhahn, J. Saunders, Evaluation of robot imitation attempts: Comparison of the system's and the human's perspectives, in: Proceedings of the 1st ACM/IEEE International Conference on Human–Robot Interactions, HRI'06, 2006.
- [102] S. Calinon, A.G. Billard, What is the teacher's role in robot programming by demonstration? Toward benchmarks for improved learning, *Interaction Studies* 8 (24) (2007) 441–464.
- [103] C. Nehaniv, K. Dautenhahn, Like me?—Measures of correspondence and imitation, *Cybernetics and Systems* 32 (41) (2001) 11–51.
- [104] M. Pomplun, M. Mataric, Evaluation metrics and results of human arm movement imitation, in: Proceedings of the 1st IEEE-RAS International Conference on Humanoid Robotics, 2000.
- [105] A. Steinfeld, T. Fong, D. Kaber, M. Lewis, J. Scholtz, A. Schultz, M. Goodrich, Common metrics for human-robot interaction, in: Proceedings of the 1st ACM/IEEE International Conference on Human–Robot Interactions, HRI'06, 2006.



Sonia Chernova is a Ph.D. student in the Computer Science Department at Carnegie Mellon University. She received her undergraduate degree in Computer Science and robotics from Carnegie Mellon University in 2003. Her research interests include learning and interaction in robotic systems.



Manuela Veloso is Herbert A. Simon Professor of Computer Science at Carnegie Mellon University. She received a *licenciatura* in Electrical Engineering in 1980, and an M.Sc. in Electrical and Computer Engineering in 1984 from the Instituto Superior Tecnico in Lisbon. She earned her Ph.D. in Computer Science from Carnegie Mellon in 1992. Veloso researches in planning, control learning, and execution algorithms, in particular for multi-robot teams. With her students, Veloso has developed teams of robot soccer agents, which have been RoboCup world champions several times. She is a Fellow of AAAI, the Association for the Advancement of Artificial Intelligence, an IEEE Senior member, and the President Elect (2008) of the International RoboCup Federation.



Brenna D. Argall is a Ph.D. candidate in the Robotics Institute at Carnegie Mellon University. Prior to graduate school, she held a Computational Biology position in the Laboratory of Brain and Cognition, at the National Institutes of Health, while investigating visualization techniques for neural fMRI data. Argall received an M.S. in Robotics in 2006, and in 2002 a B.S. in Mathematics, both from Carnegie Mellon. Her research interests focus upon machine learning techniques to develop and improve robot control systems, under the guidance of a human teacher.



Brett Browning is a Senior Systems Scientist in Carnegie Mellon University's School of Computer Science, where he has been a faculty member of the Robotics Institute since 2002. Prior to that, he was a postdoctoral fellow at Carnegie Mellon working with Manuela Veloso. Browning received his Ph.D. from the University of Queensland in 2000, and a B.Electrical Engineer and B.Sc (Math) from the same institution in 1996. His research interests are on robot autonomy and in particular real-time robot perception, applied machine learning, and teamwork.