

Aude Billard, Daniel Grollman

EPFL

Lausanne

Switzerland

{aude.billard,daniel.grollman}@epfl.ch

Synonyms

Learning/Programming from/by Demonstration, Apprenticeship Learning

Definition

Imitation is the ability to recognize and reproduce others' actions – By extension, imitation learning is a means of learning and developing new skills from observing these skills performed by another agent. Imitation learning as applied to robots is a technique to reduce the complexity of search spaces for learning. When observing either good or bad examples, one can reduce the search for a possible solution, by either starting the search from the observed good solution (local optima), or conversely, by eliminating from the search space what is known as a bad solution. Imitation learning offers an implicit means of training a machine, such that explicit and tedious programming of a task by a human user can be minimized or eliminated. Imitation learning is thus a “natural” means of training a machine, meant to be accessible to lay people.




Figure 1: Examples of robots learning via imitation. In (a), a robot observes a human demonstrator and combines known motor primitives to play air hockey. In (b), a robot is taught kinesthetically and learns the arm trajectory necessary to manipulate the die.

Images copyright ATR and LASA

Theoretical Background

Like other robot learning paradigms, imitation learning has grown from observations and models of how animals such as rats, monkeys, birds, dolphins and of course humans acquire and adapt skills via interactions with others. The term itself is drawn from ethology and developmental psychology, where it describes a wide variety of imitative behaviors growing in complexity across species (Breazeal and Scassellati, 2002). One distinguishes between simple imitative behavior, also referred to as copying, mimicry or even social facilitation, which one finds in rats and monkeys, and “true imitation,” where the animal can learn behaviors not part of its innate repertoire, and is in the purview of high primates, particularly humans. When used in the context of robot learning, IL also allows for some gradation in the complexity of the behavior that can be learned and the amount of apriori information given to the system. While one body of work is devoted to learning very low-level skills (e.g. encoding skills as joint trajectories, Figure 1b), other work focus on learning complex behaviors by sequencing and aggregating built-in behaviors (Figure 1a) (Billard et al., 2008).

Robot imitation learning was developed to address several difficulties with the standard methods of instantiating robot controllers. Primary is the inability of humans to pre-program a robot for every possible scenario, but secondary concerns include making the robot’s behavior appear more ‘natural,’ and enabling non-programmers to modify robot behavior. The approach itself proceeds with the robot observing performances of a task by the expert (demonstrations), and subsequently developing a control law with the goal of successfully completing the task in novel situations. As intended, this approach to controller creation stands in contrast to explicit programming, where a user must first analyze the task and then explicitly tell (program) the robot how to perform it, but also differs from reinforcement learning, where the user only gives out rewards or punishments and does not indicate what appropriate behavior is. Instead, in IL the analysis of what is important to reproduce (*what to imitate*) and how to reproduce it (*how to imitate*) to achieve the task is done primarily by the robot. This analysis can sometimes be done under close guidance from the teacher, a process to which one refers as incremental learning or tutelage.

At its root, imitation learning seeks to establish a robot control policy that maps a robot’s state to actions, or changes to that state. Both state and actions can be external to the robot (perceived by sensors and carried out by actuators) or internal (residing within the robot’s ‘mind’). Demonstrations can arise from many methods, ranging from free-space demonstrations by users that have different embodiments and perceptions (Figure 1a) to direct actuator teleoperation or kinesthetic manipulation (Figure 1b). If they are generated externally, then some method of transforming the demonstrations into the robot’s frame of reference must be used (Argall et al., 2009). Either way, demonstrations give rise to state-actions pairs, indicating that for this particular task, a particular action can be associated with a particular state.

If the demonstrations cover the entire state space (or conversely, the robot were limited to experience only the states that occurred in the demonstration), then robot performance of the task could be as simple as ‘replaying’ the appropriate demonstration, assuming that the robot’s memory

and computational processes were fast enough to do so. However, in the real world state and action spaces are continuous and possibly infinite, so the necessary data cannot be gathered. Thus, as the goal of imitation learning is to enable the robot to accomplish the task in novel situations, the robot must *generalize* the demonstrations.

Important Scientific Research and Open Questions

Imitation Learning for robots is still a young field, and as such there is much debate as to how it should be gone about (Schaal et al., 2003). A few surveys exist that review the work to date (Argall et al., 2009; Billing and Hellstrm, 2010; Billard et al., 2008). However, it has emerged that the task of learning by imitation can be further broken down, with each subproblem having its own varied and ongoing research community (Nehaniv et al., 1999):

- Learning *what* to imitate - some portions of the demonstration may be irrelevant for the task.
- Learning *how* to imitate - there may be multiple possible ways to perform or encode the task.
- Learning *when* to imitate - the task can only be attempted in certain circumstances, and the demonstrator may not always be performing the task.

For considering what to imitate, one approach is to take multiple demonstrations (from the same or different users) and compare them. Similarities across the demonstrations are more likely to be associated with the task itself, while differences may be user or example dependent. An alternative is to give a more active role to the users, and enable them to indicate which portions of the demonstration are more important.

Learning how to imitate relates both to the physical actions carried out by the robot, and the underlying encoding (and learning) of the controller. As the expert and learner may have different embodiments, it may be impossible for the robot to perform the task exactly as it was shown. Even if it were possible, it may not be necessary or desirable, as there are often multiple, equally valid methods of attaining a particular goal. The method of representing and encoding the controller affects what the robot's behavior as well. For instance, one representation would be to directly encode the mapping from states to actions, so that controlling the robot becomes a function evaluation. Here the robot learns that precise trajectories through state and action space are important. Another approach would be to only determine which portions of the state-action space are desirable, in which case control becomes a path-finding problem. In this view, the robot learns that goals are important, and paths perhaps less so.

Even within these different representations, there are many possible learning methods, and often the same method can be used in different representations. It is in this area that imitation learning draws heavily from the field of machine learning for techniques. Two popular families of approaches are *Neural Networks*, which attempt to simulate how living brains work, and *Statistical Learning*,

which uses mathematical models to describe and infer the policy. Each family has a plethora of different flavors, each with their own advantages and disadvantages.

The issue of when to imitate arises when you consider robots that exist over longer timescales, or are capable of performing multiple tasks. While now demonstrations are often collected during dedicated training sessions and applied in explicit test sessions, future always-on robots will have to appropriately identify opportunities to learn, and correct situations to apply the tasks they know.

Research in all of these areas is following multiple avenues simultaneously as the field of imitation learning continues to grow and reshape itself. Some current research examines ways in which imitative learning can be combined with other learning methods, such as reinforcement learning. In that vein, robots first begin learning from imitation, but then can improve their skills by practice and feedback from humans.

Cross-References

- Learning Algorithms
- Robot Learning
- Robot Learning from Demonstration
- Model-based imitation learning

References

- B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469 – 483, May 2009.
- A. Billard, S. Calinon, R. Dillmann, and S. Schaal. Survey: Robot Programming by Demonstration. In *Handbook of Robotics*, volume chapter 59. MIT Press, 2008.
- E. A. Billing and T. Hellstrm. A formalism for learning from demonstration. *Paladyn*, 1(1):1–13, 2010.
- C. Breazeal and B. Scassellati. Robots that imitate humans. *Trends in Cognitive Science*, 6: 481–487, 2002.
- C. L. Nehaniv, H. A. Ab, and K. Dautenhahn. Of hummingbirds and helicopters: An algebraic framework for interdisciplinary studies of imitation and its applications. In *Interdisciplinary Approaches to Robot Learning*, pages 136–161. World Scientific Press, 1999.
- S. Schaal, A. Ijspeert, and A. Billard. Computational approaches to motor learning by imitation. *Philosophical Transaction of the Royal Society of London*, 358(1431):537–547, Mar. 2003.

Glossary