

Robots that imitate humans

Cynthia Breazeal and Brian Scassellati

The study of social learning in robotics has been motivated by both scientific interest in the learning process and practical desires to produce machines that are useful, flexible, and easy to use. In this review, we introduce the social and task-oriented aspects of robot imitation. We focus on methodologies for addressing two fundamental problems. First, how does the robot know what to imitate? And second, how does the robot map that perception onto its own action repertoire to replicate it? In the future, programming humanoid robots to perform new tasks might be as simple as showing them.

Cynthia Breazeal

The Media Lab, Massachusetts Institute of Technology, 77 Massachusetts Ave NE18-5FL, Cambridge MA 02139, USA.

Brian Scassellati

Dept of Computer Science, Yale University, 51 Prospect Street, New Haven, CT 06520, USA.

The study of the mechanisms that enable an individual to acquire information or skills from another individual has been a seminal topic in many areas of cognitive science. For example, ethologists attempt to understand how bees communicate the location of food sources, to describe how successive generations of blue-tits learn to open milk cans, and to categorize the spread of tool use in chimpanzee troops. Developmental psychologists study the emergence of social learning mechanisms in human infants from the very early (but simple) imitative responses of the newborn [1] to the complex replication of task goals that toddlers demonstrate [2].

Research in robotics has focused on social learning for many reasons. Commercial interest in building robots that can be used by ordinary people in their homes, their workplaces, and in public spaces such as hospitals and museums, invoke social learning as a mechanism for allowing users to customize systems to particular environments or user preferences. Research in artificial intelligence has focused on social learning as a possible means for building machines that can acquire new knowledge autonomously, and become increasingly more complex and capable without requiring additional effort from human designers. Other researchers implement models of social behavior in machines to gain a deeper understanding of social learning in animals (including humans).

Differences between the study of social learning in animals and machines

The methods for studying social learning in artificial systems differ significantly from methods used to study social learning in biological systems. When studying animals, researchers attempt to determine the minimal set of capabilities required to produce an observed behavior. Precise taxonomies of the types of required skill have been developed; however, none of these is universally accepted (see Box 1). Although these descriptions often focus on cognitive skills, they do not

completely capture the ways in which these skills can be constructed or combined to produce the observed behavior.

Whereas biological studies tend to be descriptive, studies of social learning in artificial systems are primarily generative; researchers attempt to construct a desired behavior from a minimal set of capabilities. These studies often use imprecise definitions of the external behavior (often using the word *imitation* to mean any type of social learning), but can precisely specify the underlying mechanisms of the system (see Box 2). Although these methodological differences do produce terminology problems between these related disciplines, on the whole, the literature on social learning in animals is a very accessible source of inspiration for robots, both physical and simulated (see Box 3).

Many different underlying mechanisms can produce the same observable behavior

There are many ways in which a robot can be made to replicate the movement of a human. Animatronic devices (such as those used in amusement parks) continuously replay movements that have been recorded either by manually putting the machine into a sequence of postures or by using devices that record the joint angles of a human actor. Although these machines can perform very high fidelity playback, they are non-interactive; they neither respond to changes in their environment nor do they adapt to new situations.

Other research has focused on the development of robots that can learn to perform tasks by observing a person perform that action. This technique, often called 'learning from demonstration', has been reviewed in detail by Schaal [3]. Early explorations did not focus on perceiving the movement of the human demonstrator, but rather focused on observing the effects of those movements on objects in the environment (such as stacking blocks [4] or peg insertion [5]). In other work, the robot observes the human's performance as well, using both object and human movement information to estimate a control policy for the desired task. Providing the robot with knowledge of the goal (in the form of an evaluation function) allows the robot to further improve its performance through trial and error, for instance, for a 'ball-in-cup' task [6] or the task of playing air hockey (Fig. 1). Atkeson and Schaal [7] demonstrated that far fewer real-world practice trials were needed if the robot could simulate its experience using a predictive forward model for a pendulum-swing-up task. Although systems that learn from demonstration have been programmed to perform impressive feats, the systems are limited by the fact that information flows only in one direction: from human to machine.

Imitation and social interaction in robots

Studies of social learning in robotic systems have looked at a wide range of learning situations and techniques. Initial studies of social learning in robotics focused on allowing one robot to learn to navigate through mazes [8] or an unknown landscape [9] by using simple perception (proximity and infrared sensors) to follow another robot that was adept at maneuvering in the environment. Other work in social learning for autonomous robots addressed learning inter-pers. commun. protocols between similar robots, between robots with similar morphology but which differ in scale [10], and with a human instructor [11]. Other approaches have looked at expressive imitation involving facial displays and head gestures [12–14].

Although the individual tasks in each of these studies varied considerably, each of these studies looked at social interaction as a means to address two fundamental issues. First, how does the robot decide what to imitate? Second, how does the robot act upon that decision to perform a similar action? For simplicity, in the following discussion we look only at systems that involve social learning between a human and a robot that has a similar physical body structure to a human (see [15] for a discussion of the difficulties that arise when body structures are radically different).

How does a robot know what to imitate?

When attempting to imitate a human, how does the robot determine what perceptual aspects are relevant to the task? The robot needs to detect the demonstrator, observe his or her actions, and determine which are relevant to the task, which are part of the instructional process, and which are circumstantial [16]. This is a challenging problem for perceptual systems and involves not only the ability to perceive human movement, but also the abilities to determine saliency (i.e. what is important) and to direct attention.

Perception of movement

The visual perception of 3-D movement of humans or objects continues to be a difficult problem for robot vision systems. This problem can be avoided by using motion capture technologies, such as an externally worn exoskeleton that measures joint angle (e.g. a Sarcos *SenSuit*), or placing magnetic markers on certain joints and tracking them (e.g. the *FastTrak* system) [17]. Other simplifications, such as marking relevant objects with magnetic tags or distinctive colors, are often used [4,5,7, 18,19].

More general solutions to the problem of perceiving human movement through vision have yet to be realized [20,21], but many researchers are turning to techniques such as hidden Markov models [22], or perceptual–motor primitives (see Box 4) [23,24] to provide basic information on how a human is moving in a visual scene. These techniques combine task-based knowledge with predictive models in an attempt to link expectations of what the scene should look like with sensory data. Although these techniques can provide information on how a person is moving, subsequent extensive tuning to

the particular robot and environment are often necessary to produce usable data.

Attention

The problems of perception are closely tied to models of attention. Some attention models selectively direct computational resources to areas containing task-related information. They do this either by using fixed criteria [23,25] (such as ‘always look at red objects when trying to pick apples’) or by using adaptive models that modify the attentional process based on the robot’s social context and internal state. For example, the humanoid robot *Cog* (see Fig. 2) was biased to attend to objects with colors that matched skin tones when it was ‘lonely’, and to attend to objects that were brightly colored when ‘bored’ [26]. Another strategy is to use imitative behavior as an implicit attentional mechanism that allows the imitator to share a similar perceptual state with the demonstrator [27,9]. This approach is used in the ‘learning-by-imitation’ paradigm, in which the ability to imitate is given *a priori* and acts as a mechanism for reinforcing further learning and understanding. Hence, as Demiris and Hayes put it, ‘the learner isn’t imitating because it understands what the demonstrator is showing, but instead learns to understand because it is imitating’ [24]. For instance, those authors used this technique to teach a robot a control policy for how to traverse a series of corridors by following another robot [8].

Shared attention, the ability to attend to the demonstrator’s object of attention, has also been explored as a means for a robot to determine critical task elements [13]. Many machine vision systems have looked at the problems of identifying cues that indicate attention, such as pointing [28], head pose [29], or gaze direction [30]. However, only in the past few years has it become practical to use these systems in real time on robotic systems [31,32].

How does a robot know how to imitate?

Once a relevant action has been perceived, the robot must convert that perception into a sequence of its own motor responses to achieve the same result. Nehaniv and Dautenhahn have termed this the correspondence problem [15]. Although it is possible to specify the solution to the correspondence problem *a priori*, this is practical only in simple systems that use the learning-by-imitation paradigm described above. When the solution to the correspondence problem is acquired through experience, more complex perceptions and actions can be accommodated, and this is then referred to as ‘learning to imitate’.

Representing perceived movement in motor-based terms

One strategy to attempt to solve the correspondence problem is to represent the demonstrator’s movement trajectory in the coordinate frame of the imitator’s motor coordinates. This approach was explored by Billard and Schaal [33], who recorded human arm movement data using a Sarcos *SenSuit*, and then projected that data into an intrinsic frame of reference for a 41-degree-of-freedom humanoid simulation [34]. Another approach, the use of perceptual–motor primitives [35,36], is inspired by the

discovery of 'mirror neurons' in primates, which are active both when a goal-oriented action is observed and when the same action is performed [37–40]. Mataric adapted this idea to allow a simulated upper-torso humanoid robot to learn to imitate a sequence of arm trajectories [23] (see Fig. 3 and Box 4).

Representing motor movements in task-based terms

An alternative to converting perceptions into motor responses is to represent the imitator's motor acts in task space, where they can be compared directly with the observed trajectory. Predictive forward models have been proposed as a way to relate observed movement to those motor acts that the robot can perform [19,24,41,42]. Their power has been demonstrated in model-based imitation learning: Atkeson and Schaal have shown how a forward model and *a priori* knowledge of the task goal can be used to acquire a task-level policy from reinforcement learning in very few trials [18]. They demonstrated an anthropomorphic robot learning how to perform a pole-balancing task in a single trial, and a pendulum swing up task in three to four trials [18,19].

Demiris and Hayes [24] present a related technique that emphasizes the bi-directional interaction between perception and action, whereby movement recognition is *directly* accomplished by the movement-generating mechanisms. They call this 'active imitation' to distinguish it from passive imitation (which follows a one-way perceive–recognize–act sequence). To accomplish this, a forward model for a behavior is built directly into the behavior module responsible to producing that movement.

Conclusion

Imitation-inspired mechanisms have played three dominant (and related) roles in robotics research to date. First, imitation can be an easy way to program a robot to perform novel actions simply by observing a demonstration (see Fig. 1). Second, imitation can be a mechanism for communicating (between a robot and a human or between two robots). Shared meaning for gestures (Fig. 3) or a lexicon (Fig. 4) have been accomplished by learning to map shared sensory–motor experiences between two different bodies (robot to human, or robot to robot). 'Learning to imitate' frames the motor learning problem as one of acquiring a mapping between a perceived behavior and the underlying movement primitives. By representing perceptual–motor primitives as predictive forward models, both the observation and the output of the primitive share the same coordinate representation, so measuring similarity is computationally efficient. A solution to the correspondence problem is not given to the robot in 'learning by imitation'. Instead, the learner acquires a state–action policy by following the model and thereby sharing a similar perceptual and motor state [8,9,27]. This mapping often represents a shared interpersonal communication protocol, where the model announces the labels for particular sensory–motor states as they occur and the follower learns their association.

Third, imitation has been an effective tool for efficient motor learning in high-dimensional spaces. For a humanoid robot with many articulated joints, the state–action space becomes prohibitively large to search for a viable solution in reasonable time. The issue of learning efficiency has been addressed both by building more compact state–action spaces using movement primitives (Box 4) (inspired by their biological counterpart [40]), and by constraining the search through state–action space by using a human demonstration of the skill as an example [3]. Alternatively, a predictive forward model can be learned from the human demonstration, and used as simulated experience to accelerate trial and error learning [7].

Imitation and other forms of social learning hold tremendous promise as a powerful means for robots (humanoid and otherwise) to acquire new tasks and skills. Unfortunately, the most advanced robots we have currently are less adept than 2-year-old children at imitating the actions and goals of people. This review focused on two fundamental issues (what to imitate and how to imitate) that are far from solved, but there are many other important research areas that need to be addressed (see Questions for future research). It is our belief that research on these issues in artificial systems will both benefit from, and inform, research on imitation in biological systems. The synthetic approach of building systems that imitate requires attention to details that are often not part of the analytic study of social behavior in animals. For example, the process of selecting which object to imitate is not often addressed in literature on animal social learning but is a critical part of any robotic implementation. Further, we believe that imitating robots offer unique tools to evaluate and explore models of animal (and human) behavior. Just as simulations of neural networks have been useful in evaluating the applicability of models of neural function, these robots can serve as a test-bed for evaluating models of human and animal social learning.

Imitation is a sophisticated form of socially mediated learning. To date, however, robots that learn by imitation-inspired mechanisms are not particularly social themselves. In the examples above, the interaction is in one direction, from demonstrator (or model) to learner, rather than there being a bi-directional exchange of information. In human infants, imitation is hypothesized to play an important early role in the development of social cognition, serving as a discovery procedure for understanding persons, and providing the earliest 'like me' experiences of the self in relation to others [2]. Beyond ease of programming and skill transfer from human to robot, imitation could one day play a role in understanding the social cognition of robots as they begin to co-exist with people.

References

- 1 Meltzoff, A.N. and Moore, M.K. (1977) Imitation of facial and manual gestures by human neonates. *Science* 198, 74–78
- 2 Meltzoff, A.N. (1995) Understanding the intentions of others: Re-enactment of intended acts by 18 month-old children. *Dev. Psychol.* 31, 838–850

- 3 Schaal, S. (1999) Is imitation learning the route to humanoid robots? *Trends Cogn. Sci.* 3, 233–242
- 4 Kuniyoshi, Y. *et al.* (1994) Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *IEEE Trans. Robot. Autom.* 10, 799–822
- 5 Hovland, G.E. *et al.* (1996) Skill acquisition from human demonstration using a hidden Markov Model. In *IEEE International Conference on Robotics and Automation (ICRA 96)*, pp. 2706–2711, IEEE
- 6 Miyamoto, H. *et al.* (1996) A Kendama learning robot based on bi-directional theory. *Neural Netw.* 9, 1181–1302
- 7 Atkeson, C.G. and Schaal, S. (1997) Learning tasks from single demonstration. In *IEEE International Conference on Robotics and Automation (ICRA 97)*, pp. 1706–1712, IEEE
- 8 Hayes, G.M. and Demiris, J. (1994) A robot controller using learning by imitation. In *Proc. Second Int. Symp. Intell. Robots Syst.* (Borkowski, A. and Crowleg, J.L., eds), pp. 198–204, LIFTA-IMAG
- 9 Dautenhahn, K. (1995) Getting to know each other: Artificial social intelligence for autonomous robots. *Robot. Auton. Syst.* 16, 333–356
- 10 Billard, A. and Dautenhahn, K. (1998) Grounding communication in autonomous robots: An experimental study. *Robot. Auton. Syst.* 24, 71–81
- 11 Billard, A. (2002) Play, dreams and imitation in Robota. In *Socially Intelligent Agents: Creating Relationships with Computers and Robots* (Dautenhahn, K. *et al.*, eds), pp. 165–172, Kluwer
- 12 Demiris, J. *et al.* (1997) Deferred imitation of human head movements by an active stereo vision head. In *IEEE 1997 International Workshop on Robot Human Communication*, pp. 45–51, IEEE
- 13 Scassellati, B. (1998) Imitation and mechanisms of joint attention: A developmental structure for building social skills in a humanoid robot. In *Computation for Metaphors, Analogy and Agents* (Nehaniv, C., ed.), Vol. 1562 of *Springer Lecture Notes in Artificial Intelligence*, Springer-Verlag
- 14 Hara, F. and Kobayashi, H. (1996) A face robot able to recognize and produce facial expression. In *Proc. Int. Conf. Intell. Robots Syst.*, pp. 1600–1607, XXXXXXX
- 15 Nehaniv, C.L. and Dautenhahn, K. (2002) The correspondence problem. In *Imitation in Animals and Artifacts* (Dautenhahn, K. and Nehaniv, C.L., eds), pp. 41–61, MIT Press
- 16 Breazeal, C. and Scassellati, B. (2002) Challenges in building robots that imitate people. In *Imitation in Animals and Artifacts* (Dautenhahn, K. and Nehaniv, C.L., eds), pp. 363–389, MIT Press
- 17 Ude, A. *et al.* (2000) Automatic generation of kinematic models for the conversion of human motion capture data into humanoid robot motion. In *Proc. First IEEE-RAS Int. Conf. Humanoid Robots*, pp. xxx–xxx, XXXXXXX
- 18 Atkeson, C.G. and Schaal, S. (1997) Robot learning from demonstration. In *Int. Conf. Machine Learn.*, pp. 12–20, XXXXXXX
- 19 Schaal, S. (1997) Learning from demonstration. In *Advances in Neural Information Processing Systems* (Vol. 9) (Mozer, M.C. *et al.*, eds), pp. 1040–1046, MIT Press
- 20 Essa (1999) Computers seeing people. *AI Magazine* 20, 69–82
- 21 Ude, A. (1999) Robust estimation of human body kinematics from video. In *Proc. IEEE-RAS Conf. Intell. Robots Syst.*, pp. 1489–1494, IEEE/RSJ
- 22 Yang, J. *et al.* (1997) Human action learning via hidden Markov model. In *IEEE Trans. On Systems, Man and Cybernetics A: Systems and Humans* 27, 34–44
- 23 Mataric, M.J. (2002) Sensory-motor primitives as a basis for imitation: Linking perception to action and biology to robotics, in *Imitation in Animals and Artifacts* (Dautenhahn, K. and Nehaniv, C.L., eds), pp. 391–422, MIT Press
- 24 Demiris, J. and Hayes, G.M. (2002) Imitation as a dual-route process featuring predictive and learning components: A biologically plausible computational model. In *Imitation in Animals and Artifacts* (Dautenhahn K. and Nehaniv, C.L., eds), pp. 321–361, MIT Press
- 25 Mataric, M.J. and Pomplun, M. (1998) Fixation behavior in observation and imitation of human movement. *Cogn. Brain. Res.* 7, 191–202
- 26 Breazeal, C. and Scassellati, B. (1999) A context-dependent attention system for a social robot. In *Proc. Sixteenth Int. Joint Conf. Artif. Intell. (IJCAI 99)*, pp. 1146–1151, XXXXXXX
- 27 Billard, A. (2002) Imitation: a means to enhance learning of a synthetic proto-language in an autonomous robot. In *Imitation in Animals and Artifacts* (Dautenhahn K. and Nehaniv, C.L., eds), pp. 281–310, MIT Press
- 28 Darrell, T. and Pentland, A. (1996) Active gesture recognition using learned visual attention. In *Advances in Neural Information Processing Systems (NIPS)* (Touretzky, D.S. *et al.*, eds), p. 8, MIT Press
- 29 Morency, L.P. *et al.* (2002) Fast stereo-based head tracking for interactive environment. In *Proc. Int. Conf. on Automatic Face and Gesture Recognition*, pp. xxx
- 30 Matsumoto, Y. and Zelinsky, A. (2000) An algorithm for real-time stereo vision implementation of head pose and gaze direction measurement. In *Proc. IEEE Fourth Int. Conf. Face and Gesture Recognition (FG'2000)*, pp. 499–505, IEEE
- 31 Scassellati, B. (2002) Theory of mind for a humanoid robot. *Auton. Robots* 12, 13–24
- 32 Kozima, H. (1998) Attention-sharing and behavior-sharing in human-robot communication. In *IEEE Int. Workshop Robot Hum. Commun. (RoMan-98)*, pp. 9–14, IEEE
- 33 Billard, A. and Schaal, S. (2001) A connectionist model for on-line learning by imitation. In *Proc. 2001 IEEE-RSJ Int. Conf. Intell. Robots Syst.*, pp. xxx–xxx, IEEE/RSJ
- 34 Billard, A. (2001) Learning motor skills by imitation: A biologically inspired robotic model. *Cybern. Syst. J.* 32, 155–193
- 35 Weber, S. *et al.* (2000) Experiments in imitation using perceptuo-motor primitives. In *Autonomous Agents*, pp. 136–137, ACM Press
- 36 Jenkins, O.C. and Mataric, M.J. (2000) Primitive-based movement classification for humanoid imitation. *Technical Report IRIS-00-385*, University of Southern California, Institute for Robotics and Intelligent Systems
- 37 Rizzolatti, G. *et al.* (1988) Functional organization of inferior area 6 in the macaque monkey: II. Area F5 and the control of distal movements. *Exp. Brain Res.* 71, 491–507
- 38 Jeannerod, M. *et al.* (1995) Grasping objects: the cortical mechanisms of visuomotor transformation. *Trends Neurosci.* 18, 314–320
- 39 Murata, A. *et al.* (1997) Object representation in the ventral premotor cortex (area F5) of the monkey. *J. Neurophysiol.* 78, 2226–2230
- 40 Arbib, M.A. (1981) Perceptual structures and distributed motor control. In *Handbook of Physiology, Section 2: The Nervous System (Vol. II, Motor Control, Part 1)* (Brooks, V.B., ed.), pp. 1449–1480, American Physiological Society
- 41 Wolpert, D.M. and Kawato, M. (1998) Multiple paired forward and inverse models for motor control. *Neural Netw.* 11, 1317–1329
- 42 Wolpert, D.M. *et al.* (2001) Perspectives and problems in motor learning. *Trends Cogn. Sci.* 5, 487–494

Questions for future research

- Just as children develop the ability to imitate the goal of an action rather than a specific act, can we construct robots that are capable of making this inference? Today's robots respond only to the observable behavior without any understanding of the intent of an action.

- Who should the robot learn from, and when is imitative learning appropriate? Robots that imitate humans today are programmed to imitate any human within view.
- Can robots capitalize on the two-way communication of social interactions to enhance learning? What capabilities would be gained if the robot could interrupt an instructional session to ask questions, or when the instructor notices that the robot is performing an action incorrectly?

Box 1. Taxonomies of social learning

There has been little consensus on operational definitions for many of the behavioral terms used to describe social learning, although many taxonomies have been developed [a–c]. The following incomplete set of simplified definitions (adapted from [d]) is provided as an example of the range of behaviors considered under social learning.

Let A and B represent two individuals or sub-populations of individuals:

Imitation: A learns a behavior performed by B that is novel to A's behavioral repertoire. A is capable of performing the behavior in the absence of B.

Goal emulation: after observing B's actions, A produces the same end product as B. The form of A's behavior differs from B's.

Stimulus enhancement: A's attention is drawn to an object or location as a result of B's behavior.

Social support: A is more likely to learn B's behavior because B's performance produces a similar motivational state in A.

Exposure: as a result of A's association with B, both are exposed to comparable environments and thus acquire comparable behaviors.

Social facilitation: an innate behavior is released in A as a result of B's performance.

Other attempts at categorizing types of social behavior have focused on the distinction between the observable behavior and the underlying behavioral goal [e]. For example, suppose a robot were to observe a person picking up a paintbrush and applying paint to a wall. The robot could imitate the surface form of this event by moving its arm through a similar trajectory, perhaps even encountering a wall or a brush along the way. However, the underlying organizational structure of applying paint to a wall involves recognizing the intent of the action as well as the usefulness of the tool in accomplishing the goal. Meltzoff [f] has noted that by 18 months of age human children are capable of responding to both the surface form and the intended action.

References

- Galef, B.G., Jr (1988) Imitation in animals: History, definitions, and interpretation of data from the psychology laboratory. In *Social Learning: Psychological and Biological Perspectives* (Zentall, T. and Galef, B.G., eds), pp. 3–28, Lawrence Erlbaum
- Whiten, A. and Ham, R. (1992) On the nature and evolution of imitation in the animal kingdom: Reappraisal of a century of research. In *Advances in the Study of Behavior* 21, 239–283
- Caro, T.M. and Hauser, M.D. (1992) Is there teaching in nonhuman animals? *Q. Rev. Biol.* 67, 151–174
- Hauser, M.D. (1996) *The Evolution of Communication*. MIT Press
- Byrne, W. (1999) Imitation without intentionality: Using string parsing to copy the organization of behavior. *Anim. Cogn.* 2, 63–72
- Meltzoff, A.N. (1995) Understanding the intentions of others: Re-enactment of intended acts by 18 month-old children. *Dev. Psychol.* 31, 838–850

Box 2. Terms used to describe social learning in robotics

Imitative behavior refers to a robot's ability to replicate the movement of a demonstrator [a]. This ability can either be learned or specified *a priori*. For instance, in **learning by imitation** [b–d], the robot is given the ability to engage in imitative behavior, which serves as a mechanism that reinforces further learning and understanding. When the ability to imitate is learned, called **learning to imitate** [e–g], the robot learns how to solve the 'correspondence problem' through experience. In **learning by demonstration** [h–j], a new task is acquired by the robot, but this may or may not involve imitative behavior. In the case where it does not, called **task-level imitation**, the robot learns how to perform the physical task of the demonstrator (such as an assembly task [k,l]) without imitating the behaviors of the demonstrator. When given knowledge of the task goal, robots have learned to perform a physical task (e.g. learning the game of 'ball in cup' [m], or a tennis forehand [n]) by making use of both the demonstrator's movement and that of the object. Finally, the ability of a robot to learn a novel task, where it acquires both the goal and the manner of achieving it from demonstration, is referred to as **true imitation**.

References

- Schaal, S. (1999) Is imitation learning the route to humanoid robots? *Trends Cogn. Sci.* 3, 233–242
- Billard, A. (2002) Imitation: a means to enhance learning of a synthetic proto-language in an autonomous robot. In *Imitation in Animals and Artifacts* (Dautenhahn K. and Nehaniv, C.L., eds), pp. 281–310, MIT Press
- Hayes, G.M. and Demiris, J. (1994) A robot controller using learning by imitation. In *Proc. Second Int. Symp. Intell. Robots Syst.* (Borkowski, A. and Crowleg, J.L., eds), pp. 198–204, LIFTA-IMAG
- Dautenhahn, K. (1995) Getting to know each other: Artificial social intelligence for autonomous robots. *Robot. Auton. Syst.* 16, 333–356
- Billard, A. (2001) Learning motor skills by imitation: A biologically inspired robotic model. *Cybern. Syst. J.* 32, 155–193
- Mataric, M.J. (2002) Sensory-motor primitives as a basis for imitation: linking perception to action and biology to robotics, in *Imitation in Animals and Artifacts* (Dautenhahn, K. and Nehaniv, C.L., eds), pp. 391–422, MIT Press
- Demiris, J. and Hayes, G.M. (2002) Imitation as a dual-route process featuring predictive and learning components: A biologically plausible computational model. In *Imitation in Animals and Artifacts* (Dautenhahn K. and Nehaniv, C.L., eds), pp. 321–361, MIT Press
- Atkeson, C.G. and Schaal, S. (1997a) Learning tasks from single demonstration. In *IEEE Int. Conf. Robotics and Automation (ICRA 97)*, pp.1706–1712 IEEE
- Atkeson, C.G. and Schaal, S. (1997b) Robot learning from demonstration. In *Int. Conf. Machine Learning*, pp. 12–20.

- j Ude, A. (1999) Robust estimation of human body kinematics from video. In *Proc. IEEE-RAS Conf. Intell. Robots Syst.*, pp. 1489–1494, IEEE/RSJ
- k Kuniyoshi, Y. *et al.* (1994) Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *IEEE Trans. Robot. Autom.* 10, 799–822
- l Hovland, G.E. *et al.* (1996) Skill acquisition from human demonstration using a hidden Markov Model. In *IEEE Int. Conf. Robotics and Automation*, pp. 2706–2711, IEEE
- m Miyamoto, H. *et al.* (1996) A Kendama learning robot based on bi-directional theory. *Neural Netw.* 9, 1181–1302
- n Miyamoto, H. and Kawato, M. (1998) A tennis serve and upswing learning robot based on bi-directional theory. *Neural Netw.* 11, 1131–1344

Box 3. Robotic platforms: physical and simulated

The robotic community has explored the topic of imitation on a wide assortment of platforms, including physical robots and sophisticated physics-based simulations.

Humanoid robots can engage in physical and social imitation tasks and serve as extremely compelling demonstrations. They are also expensive, challenging to build, and require continual maintenance. Some systems are primarily upper torsos [a–d], some are full-body systems [e], some are only a head with a vision system [f], and some have an expressive face [g]. Although many other full-body humanoid robots have been constructed (e.g. Honda's child sized *Asimo* and Sony's knee-height *SDR-4X*) they have not yet been used in social learning systems. Simpler robots, such as small mobile robots [h,i] or robot dolls [j], have also been used to explore the social dimension of imitation. Robotic arms are popular for exploring learning how to perform physical tasks by demonstration [k–o].

Physics-based 3-D rigid-body simulations of humanoid robots are a popular alternative, allowing researchers to implement and evaluate systems quickly. Simulations produce results that are more easily replicated, as the software can often be shared among researchers. The primary difficulty with simulations is in transferring results from simulation to physical robots. Solutions that tend to work even in complex simulations often fail in the real world because of the inherent lower fidelity of simulations. A few collaborations exist allowing researchers who work mostly with simulated humanoids to test their theories and implementations on actual robots [p,q].

References

- a Brooks, R. *et al.* (1999) The Cog project: building a humanoid robot. In *Computation for Metaphors, Analogy and Agents* (Nehaniv, C.L. ed.), Vol. 1562 Springer Lecture Notes in Artificial Intelligence, Springer-Verlag
- b Kozima, H. and Zlatev, J. (2000) An epigenetic approach to human-robot communication. In *IEEE Int. Workshop Robot Hum. Commun. (RoMan-2000)*, pp. 346–351, IEEE
- c Kuniyoshi, Y. and Nagakubo, A. (1997) Humanoid as a research vehicle into flexible complex interaction. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS97)*, pp. xxx–xxx, IEEE
- d http://vesuvius.jsc.nasa.gov/er_er/html/robonaut/robonaut.html
- e Kotosaka, S. *et al.* (2000) Humanoid robot DB. In *Proc. Int. Conf. Machine Automation (ICMA2000)*, pp. 21–26, Xxxxxxx
- f Nehaniv, C.L. and Dautenhahn, K. (2002) The correspondence problem. In *Imitation in Animals and Artifacts* (Dautenhahn, K. and Nehaniv, C.L., eds), pp. 41–61, MIT Press
- g Hara, F. and Kobayashi, H. (1996) A face robot able to recognize and produce facial expression. In *Proc. Int. Conf. Intell. Robots Syst.*, pp. 1600–1607, Xxxxxxx
- h Hayes, G.M. and Demiris, J. (1994) A robot controller using learning by imitation. In *Proc. Second Int. Symp. Intell. Robots Syst.* (Borkowski, A. and Crowleg, J.L., eds), pp. 198–204, LIFTA-IMAG
- i Dautenhahn, K. (1995) Getting to know each other: Artificial social intelligence for autonomous robots. *Robot. Auton. Syst.* 16, 333–356
- j Billard, A. (2002) Play, dreams and imitation in *Robota*. In *Socially Intelligent Agents: Creating Relationships with Computers and Robots* (Dautenhahn, K. *et al.*, eds), pp. 165–172, Kluwer
- k Kuniyoshi, Y. *et al.* (1994) Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *IEEE Trans. Robot. Autom.* 10, 799–822
- l Miyamoto, H. *et al.* (1996) A Kendama learning robot based on bi-directional theory. *Neural Netw.* 9, 1181–1302
- m Atkeson, C.G. and Schaal, S. (1997) Learning tasks from single demonstration. In *IEEE International Conference on Robotics and Automation (ICRA 97)*, pp. 1706–1712, IEEE
- n Atkeson, C.G. and Schaal, S. (1997) Robot learning from demonstration. In *Int. Conf. Machine Learning*, pp. 12–20, Xxxxxx
- o Schaal, S. (1997) Learning from demonstration. In *Advances in Neural Information Processing Systems* (Vol. 9) (Mozer, M.C. *et al.*, eds), pp. 1040–1046, MIT Press
- p Mataric, M.J. (2000) Getting humanoids to move and imitate. *IEEE Intell. Syst.* 15, 18–23
- q Atkeson, C.G. *et al.* (2000) Using humanoid robots to study human behavior. *IEEE Intell. Syst.* 15, 46–56

Box 4. Movement primitives

'Movement primitives' (also referred to as perceptual–motor primitives, basis behaviors, motor schemas, macro actions, or motor programs [a,b]) are a compact representation of action sequences for generalized movements that accomplish a goal. From a computational perspective, a movement primitive can be formalized as a 'control policy', encoded using a few parameters in the form of a parameterized motor controller for achieving a particular task [c,d]. Examples of movement primitives include behaviors such as 'walking', 'grasping', or 'reaching', and they are often characterized as discrete straight-line movements, continuous oscillatory movements, or postures [e]. The primitives of a system serve as the basis set of motor programs (a movement 'vocabulary'), which are sufficient, through combination operators, for generating the robot's entire movement repertoire. The primitives allow positions and trajectories to be represented with fewer parameters, although with a corresponding loss of granularity and/or generality. As a result, more recent work has focused on using imitation as a way of acquiring new primitives (as new sequences or combinations of existing primitives) that can be added to the repertoire [f,g].

References

- a Arbib, M.A. (1981) Perceptual structures and distributed motor control. In *Handbook of Physiology, Section 2: The Nervous System (Vol. II, Motor Control, Part 1)* (Brooks, V.B., ed.), pp. 1449–1480, American Physiological Society
- b Bizzi, E. *et al.* (1991) Computations underlying the execution of movement: a biological perspective. *Science* 253, 287–291
- c Sternad, D. and Schaal, D. (1999) Segmentation of endpoint trajectories does not imply segmented control. *Exp. Brain Res.* 124, 118–136
- d Schaal, S. (1999) Is imitation learning the route to humanoid robots? *Trends Cogn. Sci.* 3, 233–242
- e Mataric, M.J. *et al.* (1998) Movement control methods for complex, dynamically simulated agents: *Adonis* dances the Macarena. In *Proc. Second Int. Conf. Autonomous Agents* (Johnson, W.L., ed.), pp. 317–324, Xxxxxx
- f Mataric, M.J. (2002) Sensory-motor primitives as a basis for imitation: linking perception to action and biology to robotics, in *Imitation in Animals and Artifacts* (Dautenhahn, K. and Nehaniv, C.L., eds), pp. 391–422, MIT Press
- g Demiris, J. and Hayes, G.M. (2002) Imitation as a dual-route process featuring predictive and learning components: A biologically plausible computational model. In *Imitation in Animals and Artifacts* (Dautenhahn K. and Nehaniv, C.L., eds), pp. 321–361, MIT Press



Fig. 1. *DB*, a full-torso humanoid robot offered commercially by Sarcos, which can learn to play air hockey by observing the movements that a human player makes. The robot's visual system attends to the green puck and the positions of the human player's red paddle. By playing against experienced opponents, the robot learns to position its own paddle to defend its goal successfully and to shoot at the opponent's goal.

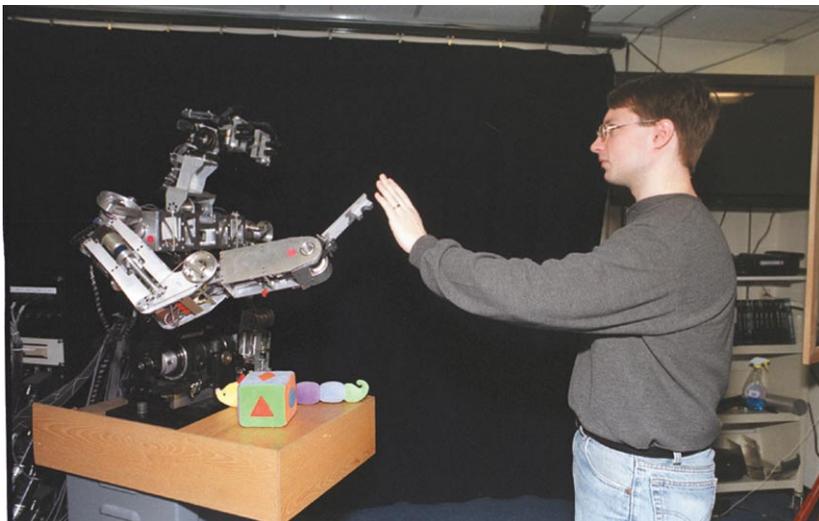


Fig. 2. *Cog*, an upper-torso robot capable of mimicking arm gestures. *Cog* uses an attention system based on models of human visual attention to locate multiple objects of interest in the environment (such as the author's hand here), selects object trajectories that display animate characteristics (i.e. trajectories that display self-propelled motion) and that the human instructor is attending to (based on the instructor's head orientation), and attempts to map these trajectories onto the movement of its own arm.



Fig. 3. *Adonis*, a rich physics-based simulation of a humanoid upper torso, which has learned to dance the Macarena based on motion capture data from a human dancer. *Adonis* uses motion primitives to map the recorded movements of the human dancer to the range of possible motions that it is capable of performing.



Fig. 4. *Robota* is a robot doll currently under development at USC. It is able to mimic a few simple gestures of a person wearing infrared markers, such as raising an arm or turning one's head. The demonstrator presses a sequence of keys on a keyboard (each key represents a label such as 'move', 'arm', 'left', etc.), at the same time as performing the corresponding gesture. Using a recurrent, associative neural network, the doll learns the association between the sequence of keystrokes and how they map onto its actions and perceptions of different parts of its body. After training, the demonstrator can press a new sequence of keys without performing the corresponding gesture, and the robot performs it.