

Is Robot Learning a New Subfield?

The Robolearn-96 Workshop

Henry Hexmoor, Lisa Meeden, and Robin R. Murphy

Learning takes place when the system makes changes to its internal structure so as to improve some metric on its long-term future performance, as measured by a fixed standard (Russell 1991, p. 141).

Let us take "learning" to mean, roughly, the improvement of a system's behavior by making it more appropriate for the environment in which it is embedded (Kaelbling 1993, p. 3).

Although both definitions given here, offered by Stuart Russell and Leslie Kaelbling, describe learning as an improvement in external performance, they differ in how this improvement is measured. Rather than employing a fixed standard, for Kaelbling, learning is evaluated by how appropriate the resulting behavior is for a particular environment. This focus on the situatedness of the learning system being embedded in its environment reflects the recent experience gained by much direct experimentation with physical robots.

Learning to control a physical robot has proved to be a challenge because of the unpredictableness of the real world coupled with the inaccuracies inherent in physical sensors and actuators. The recent Robolearn-96 Workshop¹ (Hexmoor and Meeden 1996) focused on learning as it is applied to real robots to try to examine why successful learning algorithms from the machine-learning tradition do not seem to translate well to robot-learning domains.

This article is a position paper on the role of robot learning relative to other disciplines. Our discussion reflects the sentiments expressed at the Robolearn-96 Workshop on this topic. Robot learning is most closely related to the fields of machine learning and robotics but also encompasses aspects of AI and various social sciences such as cognitive psychology. We believe that robot learning is not a proper subset of any of these current fields. A

This article posits the idea of robot learning as a new subfield. The results of the Robolearn-96 Workshop provide evidence that learning in modern robotics is distinct from traditional machine learning. The article examines the role of robotics in the social and natural sciences and the potential impact of learning on robotics, generating both a continuum of research issues and a description of the divergent terminology, target domains, and standards of proof associated with robot learning. The article argues that although robot learning is a new subfield, there is significant potential for synergy with traditional machine learning if the differences in research cultures can be overcome.

fundamental question we address is, Why is robot learning a distinct field of research?

The first section describes the role that robotics can play in the social and natural sciences. The second section clarifies the distinctions between robotics and robot learning, suggesting that they represent extremes on a

continuum of domain knowledge. The third section argues that the machine-learning and robotics communities reflect different cultures, target domains, terminology, and acceptable proofs that result in a de facto separation. The unique constraints placed on representation by robot learning are characterized in the fourth section. Finally, we close with some concluding remarks.

Robotics in the Social and Natural Sciences

Robotics is the study and practice of building robots, but robots can be useful beyond their specific applications within robotics. Robots are general tools that can profitably be used by many other disciplines, including psychology and biology. Building robots helps us think about the workings of biological systems. For example, from the mechanics of robots, we can learn about the mechanics of limbs, power transfer, and locomotion; from the organization of robot software and hardware components, we can learn about the biological circuitry of sensing and acting; from path planning and navigation, we can learn about search strategies; from the endowing of agents with mental attributes such as goals, desires, and feelings, we can better understand aspects of cognition, awareness, and emotions; and from group behavior among robots, we can learn about dynamics of cooperation, communication, coordination, and other social interactions. Just as the development of pumps ultimately made it possible to understand the working of the heart, robots are technologically anchored systems that might one day help us understand the working of biological systems (Michie, Ambler, and Ross 1974).

Although robots can be useful tools for gaining an understanding of biological systems, they should not be seen as merely tools. Robots are interesting models in their own right. Just as airplanes are their own best models of mechanical flight and not of biological flight, robots are important models for exploring control behaviors and learning in machines.

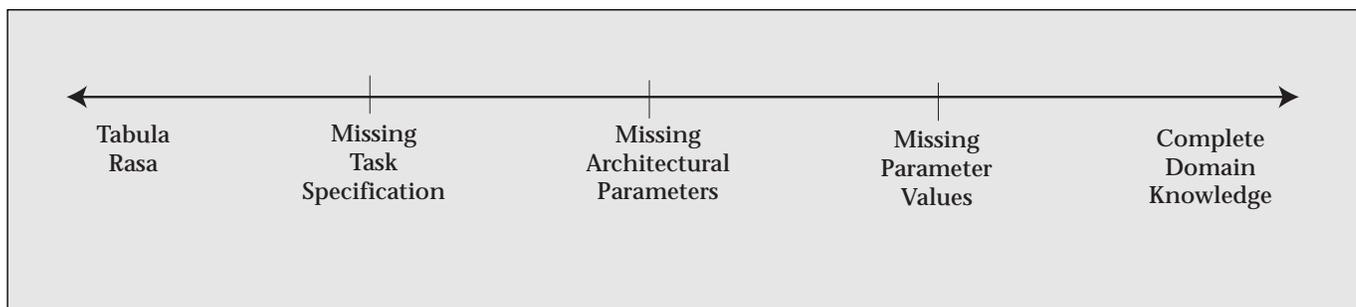


Figure 1. *The Continuum of Domain Knowledge for Robot Learning.*

Robotics Versus Robot Learning

As our mobile robots become more and more complex and as we attempt to make them more sophisticated, perhaps by incorporating more sensors for greater perceptual acuity or by adding more actuators for finer dexterity, the software can become severely strained in trying to deal with so many inputs and outputs. One area of research is to investigate how, when, and where learning algorithms can be incorporated into a robot's intelligence system to alleviate the programmer's burden (Jones and Flynn 1993, p. 272).

As the quote suggests, for the typical roboticist, learning is incorporated into the control system only when engineering a solution becomes too complex. Roboticists tend to strive for advances in hardware and software to solve robotics problems because this route to success is still more direct than the use of learning. In contrast, advocates of robot learning predict that eventually, the human designer will be the bottleneck in the design process of robotics systems. Thus, even though learning is currently a less efficient mechanism for building a robot control system, it is crucial to begin examining ways to improve its effectiveness now.

Robots are being developed for a wide variety of purposes: from industrial robots that are designed for repetitive tasks in highly structured environments to mobile office robots that are designed to interact with humans

in a more natural, open-ended setting. The engineering approach is well suited to the clearly defined tasks of a factory robot working on an assembly line, but as we push toward more autonomous robots, the tasks become increasingly less defined, and the ability of a robot to adapt becomes essential. Robot learning can mean varying degrees of change to internal structure and external behavior. It is difficult to catalog all the different types of learning, but briefly they might include (1) changing values of parameters in feedback loops, (2) building topological models, (3) developing fine-grained perceptual or motor skills, (4) encoding successful routine associations between sensing and acting, (5) acquiring new concepts, and (6) learning architectural features such as coordinations and concurrency.

Figure 1 depicts a continuum of knowledge. In one extreme, the robots have complete domain knowledge. Such a robot is an engineered artifact, and no learning is necessary. Near this extreme are examples of robots with missing parameter values, which are engineered systems containing control modules that lack optimal values for their parameters. Most of the research presented at the Robolearn-96 Workshop fell near this end of the continuum. Many of the presenters argued that to construct a successful system, learning must be limited in use to portions of the system where the designer's knowledge is too incomplete to engineer a solution. Learning from scratch is too inefficient for any problem of reasonable complexity. Therefore, starting with no domain knowledge, the other extreme on the continuum, was viewed as impractical.

Although beginning with a significant amount of domain knowledge is certainly more practical, there can be a disadvantage. In a sense, these systems are given an encoded version of the possible sensory input. Chris Atkeson (Georgia Institute of Technology), the keynote speaker at the workshop, pointed out that the more interesting systems are those that discover the appropriate input and output for their tasks. These systems determine the relevance of sensory information to tasks, synthesize this sensory information, and compose behaviors without ever being explicitly programmed for such interactions.

Machine Learning and Robot Learning: A Comparison and Search for Synergy

An objective of the robot-learning community is to exploit the natural synergy between machine learning and robotics. Machine learning supplies roboticists with valuable, robust tools, but roboticists provide feedback from situated-agent applications. Initially, it seemed feasible to simply do machine learning using robots. However, the biggest impediment to synergy is the different research cultures. The differences between machine learning and robotics reflect the implicit assumptions associated with their historical traditions. As a result, machine learning and robot learning have become nearly as separate as classical and reactive planning; both are subfields of AI but have divergent terminology, target domains, and standards of proof.

Machine learning grew out of a need to build robust computer programs that learned a search path, a game strategy, or a category. These programs were useful in recording partial paths and successful strategies in domains that were fully specified. Time was never a concern. Practically infinite amounts of training examples were provided. The training data were usually unambiguous, ignoring the realities of sensor-based observations: noise, uncertainty, ambiguity, quality, and contributions of nonhomogeneous sensors. The only requirement was a before-and-after run of the program that showed an improvement in search time or strategy. In retrospect, the machine-learning community appears to have tended toward methods that act as *disembodied software agents*: independent subsystems that do not have to interact directly with other agents, can learn offline, and could be assumed to function in a closed world.

The autonomous robot culture, however, was always first and foremost in the *situated-agent tradition*. In this tradition, the robot is viewed as a physical agent expected to function robustly in a world too large to explicitly model. It must operate in real time to be successful, and because of limited resources (time limits, energy, computation, and so on), it might not be able to collect or process large training sets. Learning almost always must be done online to be useful; for example, consider learning to walk (Mark Pendrith and Malcolm Ryan [both of University of New South Wales]), Andrija Kun and W. Thomas Miller [both of University of New Hampshire]), refining behaviors (Tae-Hoon Choi, Eunbin Yim, and Keith Doty [all of University of Florida]), and navigating (Minoru Asada and Takayuki Nakamura [both of Osaka University]). The situated-agent tradition offers some advantages for learning, including being able to test the sufficiency of learning directly and immediately (Robin Murphy [Colorado School of Mines] and Marcel Schoppers [Robotics Research Harvesting]) and continuing learning if needed.

Another attribute that distinguishes robot learning from traditional ma-

chine learning is the role of the human supervisor. In robotics, semiautonomous control is common; the human can assume control of the robot to perform a task beyond the robot's capabilities, such as obstacle avoidance (Noel Sharkey, Jan Heemkerk, and John Neary [all of University of Sheffield]). This provides the robot with the opportunity to learn by direct observation of the process as well as by the result. Humans can also provide assistance, as seen by the dialog-based map acquisition of Hideki Asoh, Yoichi Motomura, Hara, Akaho, Hayamizu, and Toshi Matsui (all of Electrotechnical Laboratory, Japan).

In some regards, the machine-learning and robotics traditions represent two different paradigms: bottom up and top down. By and large, the focus of the machine-learning community has been on taking specific instances and generalizing them into rules, models, and so on, for one application. For example, an artificial neural network for learning to recognize English spoken by a U.S. native can be implemented and trained separately from one learning to recognize English spoken with a British accent. A recent research trend in machine learning is *transfer*, which attempts to transfer what has been learned in one domain to another. Note that this is a bottom-up approach: First, learn a specific domain, then try to find the common elements with a new application. The needs of robotic learning in practice have tended to concentrate on using general strategies to learn specific instances. This is more of a top-down approach, where a general strategy exists but is tailored to the context of each new application. For example, the objective of Sven Koenig and Reid Simmons (both of Carnegie Mellon University) is to put a brand-new robot with general office-navigational capabilities in a new building and have it learn the particulars of its office building. The map-and-place learning systems of Murphy and Schoppers and Ryohei Nakano, Naonori Ueda, Kazumi Satto (all of NTT, Japan), and Mutsumi Takahashi (CSK Corporation, Japan) can also be viewed as specifying instances rather than generalizing concepts. There are notable excep-

tions, of course, including the concept-learning methods of Jun Hakura, Hiroshi Yokoi, and Yukinori Kakazu (all of Hokkaido University) and Klingspor (University of Dortmund).

Also, the acceptable standard of proof for the robotics community favors direct demonstrations over simulations for at least two reasons: First, it is difficult to satisfactorily simulate the open world, especially the response of the sensors and emergent properties of multiple behaviors. Second, as the price of commercial research robots has dropped into the \$25,000 range, many researchers now have access to platforms. Therefore, the burden of proof for a robot-learning algorithm favors convincing demonstrations that either reproduce or closely mimic the target robot application. This is not to say that simulations do not have an important place in robot learning; consider the work by Alan

Robots are being developed for a wide variety of purposes: from industrial robots that are designed for repetitive tasks in highly structured environments to mobile office robots that are designed to interact with humans in a more natural, open-ended setting.

Schultz, John Grefenstette, and William Adams (all of the U.S. Naval Research Lab) that used a simulator to learn a complex shepherding behavior, then transferred it to an operational robot. The point is that simulations are treated as weaker forms of proof of viability. The theoretical properties of algorithms, for example, proof of convergence, are secondary to the central question: Will it work in the real world?

Representation and Learning

Expressiveness versus tractability is recognized as a fundamental trade-off in knowledge representation (Levesque and Brachman 1985). We posit that for learning systems, an additional fundamental feature of the knowledge representation scheme is its manipulability. Learning systems must be able to modify their initial representations to encode new knowledge. Three primary forms of change are (1) cumulative, (2) parametric, and (3) structural. *Cumulative change* is simply the gradual accretion of new knowledge. *Parametric change* involves the tuning of control variables to improve performance in a particular environment. *Structural change* is the most drastic and indicates alterations to the very architecture of the system.

The representation choice for a robot-learning system is highly dependent on the task to be accomplished. For example, given the task of developing a controller for a gripper mechanism, one might choose a system that associates sensory input with actuator output. The parameters of this system—the strength of associations between sensing and acting—are easier to manipulate in a connectionist framework than in first-order logic. In contrast, consider the task of developing controllers for interacting robots that must cooperate to move heavy equipment. In this case, a higher-level, rule-based system that gradually deduces new rules would probably be more successful. The appropriate manipulability for a robot-learning system's representation scheme is best judged in the context of the task.

Conclusions

The breadth of papers presented at the RoboLearn-96 Workshop indicates that learning is relevant to all aspects of robotics: improving execution and control, navigating and map learning, and general concept learning. Also, the workshop papers reported on a wide variety of methods being used, including reinforcement learning and genetic algorithms.

Although we argue that the unique characteristics of robotics has led to robot learning being distinct from machine learning, we want to emphasize that robot learning cannot and should not proceed in ignorance of machine learning or robotics. Research in robot learning should not reinvent the wheels developed by both fields. Many of the wheels might not be useful at this nascent time in robot learning, but this does not preclude their eventual utility. One potential problem for robot learning is that many robotics researchers come from non-AI backgrounds; they might not be familiar with the breadth of the machine-learning field. The popular push in robotics for empirical demonstrations must be balanced with the well-developed metrics from machine learning, especially the theoretical analysis of the learning properties. The consensus of the RoboLearn-96 participants was that continuing with yearly workshops dedicated to robot learning, providing a forum for the machine learning and robotics communities to learn from each other, is highly desirable. In accordance with this consensus, RoboLearn-97 continued some of these discussions (see www.cs.buffalo.edu/~hexmoor/robolearn97-cfp.html).

Note

1. Information about the workshop and related discussions can be obtained at www.cs.buffalo.edu/~hexmoor/robolearn96.html.

References

Hexmoor, H., and Meeden, L., eds. 1996. *RoboLearn-96: An International Workshop on Learning for Autonomous Agents*, Technical Report, 96-11, Department of Computer Science, State University of New York at Buffalo.

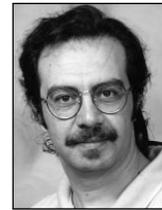
Jones, J. L., and Flynn, A. M. 1993. *Mobile Robots: From Inspiration to Implementation*. Wellesley, Mass.: A. K. Peters.

Kaelbling, L. P. 1993. *Learning in Embedded Systems*. Cambridge, Mass.: MIT Press.

Levesque, H., and Brachman, R. 1985. A Fundamental Trade-Off in Knowledge Representation and Reasoning. In *Readings in Knowledge Representation*, eds. R. Brachman and H. Levesque, pp. 41–71. San Francisco, Calif.: Morgan Kaufmann.

Michie, D.; Ambler, P.; and Ross, R. 1974. Memory Mechanisms and Machine Learning. In *On Machine Intelligence*, ed. D. Michie. New York: Wiley.

Russell, S. 1991. Prior Knowledge and Autonomous Learning. In *Toward Learning Robots*, ed. Walter Van de Velde. Cambridge, Mass.: MIT Press.



Henry Hexmoor is an assistant professor in the Department of Computer Science at the University of North Dakota. His research includes intelligent robotic systems and fusion. His e-mail address is hexmoor@cs.und.edu.



Lisa Meeden is an assistant professor of computer science at Swarthmore College in Swarthmore, Pennsylvania. Her main research interests are in connectionist and evolutionary models of cognition. Her e-mail address is meeden@cs.swarthmore.edu.



Robin R. Murphy is an assistant professor in the Department of Mathematical and Computer Sciences and is the associate director of the Center for Robotics and Intelligent Systems at the Colorado School of Mines. She writes a guest column on robot competitions for *Robotics and Autonomous Systems* and coedited *Artificial Intelligence and Mobile Robots* (MIT Press, 1997). Her e-mail address is rmurphy@mines.edu.