Learning Autonomous Robot Control from Demonstration

Assistant Professor **Chad Jenkins**' research leverages the idea that robot control and computational perception are better learned from human demonstration rather than explicit computer programming.

In recent years, technological advancements have resulted in more capable robot hardware with increasingly broad applicability through society, such as Sony's AIBO and iRobot's Roomba. Despite the great strides in sensing, actuation, and computation, developing autonomous control policies (programs to control the robot) has remained a difficult problem in robotics. A robot is defined as a machine (embodied in some physical form) that uses a control policy to act (or actuate its motors) based on information about its environment (provided by its sensors). Even when restricted to highly constrained domains, the manual crafting of a robot control policy can be a tedious and time-consuming endeavor requiring a considerable amount of technical skill. Furthermore, such controllers often lack scalability for modifying the robot's behavior as new or unforeseen functionality becomes desired.

ABOVE Brown's two RoboCup teams will use the Sony Aibo. Teams of four dogs will compete against each other in a soccer match. **BELOW** Robo-Dog makes a save! (and the fans go wild...)

A viable approach to development of autonomous robot controllers is learning them from human demonstration. Human beings are apt at devising control policies that yield functional behavior in a dynamic environment, which is especially true for video games. Unfortunately, most humans are unable to transfer their expertise onto the robot. The barrier exists because explicitly programming a computer language is often required for robot programming and most people are not computer programmers. Our aim is to implicitly transfer human expertise into robot control policies through human-robot collaboration or instruction.

Our research in this area aims to extend methods for learning from demonstration towards applicability for general sets of tasks and multi-robot domains. That is, given a set of sensory inputs and corresponding motor outputs, learning their functional relationship. By observing a human demonstrator, our systems attempt to learn the function exhibited by the human that produces motor commands from the data provided by the robot's sensors. Once learned, the robot uses this sensory-motor function to autonomously assess situational context, assign tasks to robots, and perform the learned tasks at run-time. In addition, we seek to develop methods that allow a robot to generalize and refine a learned policy for new situations and better performance in familiar scenarios.

RoboCup

One interesting application of this work is in the domain of robot soccer, specifically, RoboCup. RoboCup is a worldwide research and education initiative focused on advancing the state-of-the-art through annual robotic soccer competitions. Unlike typical competitions (e.g., FIRST, BattleBots), the robots operate in full autonomy without external control or resources. The high-level goal of RoboCup is to foster artificial intelligence and robotics research with respect to a standard problem. The dream of this initiative is to develop a team of humanoid robots capable of winning against the official human WorldCup Champion before 2050. During the fall of 2005, the Brown RoboCup team was established to field competitive robot soccer teams and explore methods for learning soccer skills and strategies from demonstrations. While several RoboCup leagues exist for different aspects of the greater challenge, we focus on the Sony Four-Legged Robot League, where Sony Aibo robots play in teams of four.

With the support of a Brown Salomon grant, Brown RoboCup is sending two teams, Brown # (read: Pound) and DemonstraDogs, to compete in the RoboCup US Open, our first competition, in Atlanta this April. The Brown # uses manually programmed control routines created by our development team of Sc.M. students including captain, Ethan Leland, and programmers Brendan Dickinson and Mark Moseley. DemonstraDogs leverages the research of doctoral student Dan Grollman to learn soccer control policies from human teleoperation of robot teams.

"RoboCup is a worldwide research and education initiative focused on advancing the state-of-the-art through annual robotic soccer competitions."

Research Directions

Though learning from demonstration is a promising avenue for advancing robotics, there are several general questions that remain. Among these questions, the issue of generalization figures to be the most prominent. Generalization deals with the ability to extrapolate beyond the scenarios explicitly demonstrated by a human. Machine learning offers many methods for addressing the generalization problem through regression and reinforcement learning. We are currently exploring which methods in this space, if any, are suited to generalization in limited and broader domains.

In parallel, we are undertaking an effort to collect large datasets of sensory-motor data from human control to minimize the magnitude of the generalization problem. Given human capabilities and interest in video games, we believe there is an opportunity to learn superior soccer policies while broadening interest in computer science.

Over the long term, we expect a convergence between learning from demonstration with time-critical decision making, as in our collaborations with Brown Assistant Professor Meinolf Sellmann. In this regard, learning is used to perform symbol grounding, that is, the discretization of a continuous control policy into a set of symbols. Once symbolized, we will explore the use of combinatorial optimization for decision making about the group behavior of a soccer team. If optimization can be performed in a time-critical fashion, it has the potential to find optimal play selections and task allocations (i.e., assignments of robots to roles) for competitive soccer. **C!**

C! Comments?

Send your views to: Conduit,
Department of Computer Science,
Brown University, Box 1910,
Providence, RI 02912
or email conduit@cs.brown.edu

Ph.D. Profiles

Aris Anagnostopoulos

Aris did his undergraduate studies in the computer engineering and informatics department at the University of Patras, Greece. He joined Brown University in 2000 and started working with his advisor, Professor Eli Upfal, in the area of algorithms, and in particular on stochastic analysis of protocols for networks. During the summers of 2004 and 2005 he interned at the IBM T..J. Watson Research Center, under the supervision of Andrei Broder, on problems related to web searching. Their paper "Sampling Search-Engine Results,"(with David Carmel at the IBM Haifa Research Lab) was accepted in the WWW 2005 conference and was runner up for best paper. Aris' other research interests include information security, stream processing and combinatorial optimization.

Ionut Aron

Ionut started his Ph.D. at Brown in 2001, working in the area of combinatorial optimization under the direction of Professor Pascal Van Hentenryck. His Ph.D. work has focused on the integration of constraint programming, developed and used mainly by the artificial intelligence community, with mathematical programming, used primarily in operations research. Due to the interdisciplinary nature of this work, Ionut spent three semesters at Carnegie Mellon University. While at CMU, he worked with researchers from the operations research group at the Tepper School of Business. Ionut returned to Brown in January, 2004, and is now completing his dissertation work. At the end of the 2006 spring semester, he will join the optimization group in the mathematical sciences department at the IBM T.J. Watson Research Center.

Nesime Tatbul

Nesime received her B.S. and M.S. degrees in computer engineering from the Middle East Technical University in Turkey. In 1999, she started her Ph.D. studies at Brown, working with Professor Stan Zdonik in the area of database systems. Her thesis research focuses on load management techniques for large-scale data stream processing systems. In particular, she has designed and implemented load shedding techniques for the Aurora/Borealis prototype systems, which support continuous real-time processing and monitoring over high-speed data streams. Nesime spent the summer of 2002 at the IBM Almaden Research Center, where she worked on replication-based update handling mechanisms for database caching. From 2003–2005, she worked as a consultant for the U.S. Army Research Institute of Environmental Medicine on developing data management solutions for wireless sensor networks. Nesime expects to complete her Ph.D. by summer 2006. **C!**