Dissertations

Autonomous Robot Skill Acquisition

George D. Konidaris, *University of Massachusetts - Amherst*

Date of Award
5-13-2011

Document Type
Open Access Dissertation

Degree Name
Doctor of Philosophy (PhD)

Degree Program
Computer Science

First Advisor
Andrew G. Barto

Second Advisor
Neil E. Berthier

Third Advisor
Roderic A. Grupen

Keywords
Reinforcement Learning, Robotics, Skill Acquisition

Subject Categories
Artificial Intelligence and Robotics | Computer Sciences

Abstract
Among the most impressive of aspects of human intelligence is skill acquisition—the ability to identify important behavioral components, retain them as skills, refine them through practice, and apply them in new task contexts. Skill acquisition underlies both our ability to choose to spend time and effort to specialize at particular tasks, and our ability to collect and exploit previous experience to become able to solve harder and harder problems over time with less and less cognitive effort.

Hierarchical reinforcement learning provides a theoretical basis for skill acquisition, including principled methods for learning new skills and deploying them during problem solving. However, existing work focuses largely on small, discrete problems. This dissertation addresses the question of how we scale such methods up to high-dimensional, continuous domains, in order to design robots that are able to acquire skills autonomously. This presents three major challenges; we introduce novel methods addressing each of these challenges.

First, how does an agent operating in a continuous environment discover skills? Although the literature contains several methods for skill discovery in discrete environments, it offers none for the general continuous case. We introduce skill chaining, a general skill discovery method for continuous domains. Skill chaining incrementally builds a skill tree that allows an agent to reach a solution state from any of its start states by executing a sequence (or chain) of acquired skills. We empirically demonstrate that skill chaining can improve performance over monolithic policy learning in the Pinball domain, a challenging dynamic and continuous reinforcement learning problem.

Second, how do we scale up to high-dimensional state spaces? While learning in relatively small domains is generally feasible, it becomes exponentially harder as the number of state variables grows. We introduce abstraction selection, an efficient algorithm for selecting skill-specific, compact representations from a library of available representations when creating a new skill. Abstraction selection can be combined with skill chaining to solve hard tasks by breaking them up into chains of skills, each defined using an appropriate abstraction. We show that abstraction selection selects an appropriate representation for a new skill using very little sample data, and that this leads to significant performance improvements in the Continuous Playroom, a relatively high-dimensional reinforcement learning problem.

Finally, how do we obtain good initial policies? The amount of experience required to learn a reasonable policy from scratch in most interesting domains is unrealistic for robots operating in the real world. We introduce CST, an algorithm for rapidly constructing skill trees (with appropriate abstractions) from sample trajectories obtained via human demonstration, a feedback controller, or a planner. We use CST to construct skill trees from human demonstration in the Pinball domain, and to extract a sequence of low-dimensional skills from demonstration trajectories on a mobile robot. The resulting skills can be reliably reproduced using a small number of example trajectories.

Finally, these techniques are applied to build a mobile robot control system for the uBot-5, resulting in a mobile robot that is able to acquire skills autonomously. We demonstrate that this system is able to use skills acquired in one problem to more quickly solve a new problem.