Dissertations

Basis Construction and Utilization for Markov Decision Processes Using Graphs

Jeffrey Thomas Johns, *University of Massachusetts - Amherst*

Date of Award

2-2010

Document Type

Open Access Dissertation

Degree Name

Doctor of Philosophy (PhD)

Degree Program

Computer Science

First Advisor

Sridhar Mahadevan

Second Advisor

Andrew G. Barto

Third Advisor

Shlomo Zilberstein

Keywords

Markov decision process, Reinforcement learning, Representation discovery

Subject Categories

Computer Sciences

Abstract

The ease or difficulty in solving a problemstrongly depends on the way it is represented. For example, consider the task of multiplying the numbers 12 and 24. Now imagine multiplying XII and XXIV. Both tasks can be solved, but it is clearly more difficult to use the Roman numeral representations of twelve and twenty-four. Humans excel at finding appropriate representations for solving complex problems. This is not true for artificial systems, which have largely relied on humans to provide appropriate representations. The ability to autonomously construct useful representations and to efficiently exploit them is an important challenge for artificial intelligence. This dissertation builds on a recently introduced graph-based approach to learning representations for sequential decision-making problems modeled as Markov decision processes (MDPs). Representations, or basis functions, forMDPs are abstractions of the problem's state space and are used to approximate value functions, which quantify the expected long-term utility obtained by following a policy. The