Doctoral Dissertations 2014-current

Off-campus UMass Amherst users: To download campus access dissertations, please use the following link to log into our proxy server with your UMass Amherst user name and password.

Non-UMass Amherst users: Please talk to your librarian about requesting this dissertation through interlibrary loan.

Dissertations that have an embargo placed on them will not be available to anyone until the embargo expires.

ADAPTIVE STEP-SIZES FOR REINFORCEMENT LEARNING

William C. Dabney, *Computer Sciences*

Document Type
Open Access Dissertation

Degree Name
Doctor of Philosophy (PhD)

Degree Program
Computer Science

Year Degree Awarded
Fall 2014

First Advisor
Andrew G. Barto

Second Advisor
Sridhar Mahadevan

Third Advisor
Benjamin Marlin

Fourth Advisor
Hari Jagannathan Balasubramanian

Keywords
reinforcement learning, machine learning, step-size, learning rate, evaluation, adaptive

Subject Categories
Artificial Intelligence and Robotics

Abstract
The central theme motivating this dissertation is the desire to develop reinforcement learning algorithms that “just work” regardless of the domain in which they are applied. The largest impediment to this goal is

the sensitivity of reinforcement learning algorithms to the step-size parameter used to rescale incremental updates. Adaptive step-size algorithms attempt to reduce this sensitivity or eliminate the step-size parameter entirely by automatically adjusting the step size throughout the learning process. Such algorithms provide an alternative to the standard "guess-and-check" methods used to find parameters known as parameter tuning.

However, the problems with parameter tuning are currently masked by the way experiments are conducted and presented. In this dissertation we seek algorithms that perform well over a broad subset of reinforcement learning problems with minimal parameter tuning. To accomplish this we begin by addressing the limitations of current empirical methods in reinforcement learning and propose improvements with benefits far outside the area of adaptive step-sizes.

In order to study adaptive step-sizes in reinforcement learning we show that the general form of the adaptive step-size problem is a combination of two dissociable problems (adaptive scalar step-size and update whitening). We then derive new parameter-free adaptive scalar step-size algorithms for the reinforcement learning algorithm Sarsa(λ) and use our improved empirical methods to conduct a thorough experimental study of step-size algorithms in reinforcement learning. Our adaptive algorithms (VES and PARL2) both eliminate the need for a tunable step-size parameter and perform at least as well as Sarsa(λ) with an optimized step-size value. We conclude by developing natural temporal difference algorithms that provide an approximate solution to the update whitening problem and improve performance over their non-natural counterparts.