
Exploring compact reinforcement-learning representations with linear regression ^{*}

Thomas J. Walsh[†] **István Szita**[‡]
[†]Dept. of Computer Science
Rutgers University
Piscataway, NJ 08854

Carlos Diuk[†] **Michael L. Littman**[†]
[‡]Dept. of Computing Science
University of Alberta
Edmonton, AB Canada T6G 2E8

Abstract

This paper presents a new algorithm for on-line linear regression whose efficiency guarantees satisfy the requirements of the KWIK (Knows What It Knows) framework. The algorithm improves on the computational and storage complexity bounds of the current state-of-the-art procedure in this setting. We explore several applications of this algorithm for learning compact reinforcement-learning representations. We show that KWIK linear regression can be used to learn the reward function of a factored MDP and the probabilities of action outcomes in Stochastic STRIPS and Object Oriented MDPs, none of which have been proven to be efficiently learnable in the RL setting before. We also combine KWIK linear regression with other KWIK learners to learn larger portions of these models, including experiments on learning factored MDP transition and reward functions together.

^{*}This Technical Report corrects an earlier version of this work (UAI-09) that erroneously reported an $O(n^2)$ improvement in the KWIK sample complexity of the described linear regression algorithm based on an assumption about the inputs that is not universally valid in the KWIK setting. The reported computational and space improvements do, however, hold. Detailed proofs and a further explanation of the discrepancy are provided in this document.