Research Statement

Thomas J. Walsh

The long standing goal of Artificial Intelligence (AI) research has been to develop agents that can perform tasks in general settings. To this end, many traditional AI systems were developed using high-powered representations, such as first order logic. However, many of these systems, despite their rather ingenious representations, were not designed to adapt based on feedback from an environment. In contrast, research in the AI subfield of Machine Learning (ML) has focused on extracting models from data with guarantees about the amount of data needed to reach the target model. But often the more general models from traditional AI have been eschewed in favor of specific models catered to the problem at hand. My research has focused on data driven machine-learning algorithms that use *compact* general models in line with more traditional AI methods, but still maintain theoretical guarantees about the amount of data they need.

The main vein of my research has been in the reinforcement learning (RL) [2] paradigm. RL covers a large class of problems where an agent is able to observe aspects of its environment and can take *actions* to change the *state* of the environment and receive *rewards*, with the long term goal of maximizing some function of these reward terms (e.g. their discounted sum). Within RL, my research has been focussed on "high level" representations (such as relational operators) with a strong emphasis on algorithms that learn such operators in a *sample efficient* manner. Sample efficiency in RL is a theoretical measure of algorithmic complexity that bounds the number of steps required for an agent to achieve near-optimal behavior (e.g. [1]). Traditionally, agents with such theoretical guarantees have learned only propositional models (e.g. MDPs, DBNs) of their environment. My research concerns agents that learn *compact* relational models of the environment (e.g. STRIPS operators, PDDL rules, stochastic action schemas). For example, in a domain where an agent stacks blocks, instead of learning that the action move(a, b, c) moves block a from block b to block c (the traditional RL approach), agents learn the more general form that moving *any* block X from Y to Z has the intended effect and then can apply this learned compact model to any block it wants to move. Unlike other systems for learning such operators, my research has shown ways to maintain sample-efficiency when learning these models, allowing us to make theoretical guarantees about the behavior of the corresponding agents, even in complex non-deterministic domains (see [3, 4]). This approach combines the best of both worlds: large state spaces can be represented by the compact operators, but their minimalist form allows them to be learned in a manner that preserves sample efficiency and near-optimal behavior.

I've also been active in research in more traditional RL as well as machine-learning and data-mining. I've designed and implemented reinforcement learning algorithms for Aibo robots using their visual parsing capabilities to identify object locations and colors and building a complex state space out of these observations. Outside of RL, I have collaborated on projects for learning attribute preferences from data, learning web-service descriptions within workflows, and in two different projects I've helped develop large scale text mining programs for monitoring data traffic and web sites, typically with variants of Latent Semantic Analysis. The common thread in all of these projects is the construction of compact (but expressive) models for complex (but limited) data. These dual goals blend the traditional AI focus on rich representations with the machine-learning dependency on realistic data and sample efficiency.

My future research goals are to continue to develop machine-learning algorithms that learn fast and generalize well. I am particularly interested in pushing my research further into areas such as robotics where very large amounts of real data needs to be processed, but compact models of the world still need to be learned. In such settings, while theoretical guarantees may not be of much use, I believe many of the techniques I have developed will show empirical benefits in the real world for agents that need to both put together a compact representation of their perceptually complex environment and simultaneously exploit that model to achieve their tasks.

References

[1] Lihong Li, Michael L. Littman, and Thomas J. Walsh. Knows what it knows: A framework for self-aware learning. In *ICML*, 2008.

[2] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction.* MIT Press, Cambridge, MA, March 1998.

[3] Thomas J. Walsh and Michael L. Littman. Efficient learning of action schemas and web-service descriptions. In *AAAI*, 2008.

[4] Thomas J. Walsh, Istvn Szita, Carlos Diuk, and Michael L. Littman. Exploring compact reinforcement-learning representations with linear regression. In *UAI*, 2009.