Reinforcement Learning with Regularization Networks

Tobias Jung (UT Austin)

January 21, 2011

Reinforcement learning addresses the most intriguing class of problems faced by living creatures and artificial agents alike: that of making (or learning how to make) optimal decisions in a complex world without knowing the exact rules by which the world will respond to the decisions made. Reinforcement learning is a universal methodology that is widely applicable: it is useful for any task that involves taking a sequence of actions and where the outcome of one action influences the utility of subsequent actions. Practical applications abound and range from business and operations research to optimal control and robotics. In fact, reinforcement learning possibly comes nearest to what artificial intelligence is all about, which is (simply put): "How can we build intelligent machines."

Reinforcement learning has its roots in classical dynamic programming. Central to this methodology is the concept of a value function which measures the utility and desirability of states in the world (similar to an evaluation function in board games). The optimal value function is obtained by solving a functional equation (called Bellman's equation). Unfortunately, for all problems of practical interest, we can solve this equation only approximately, borrowing various techniques ranging from function approximation and statistical regression to pattern recognition and linear programming.

In this talk I will discuss regularization networks as a modern approach to function approximation in reinforcement learning. Powerful nonparametric methods (such as regularization networks and the related Gaussian process regression) expand the solution directly in the data, thus allowing the parametrization to automatically adapt itself to the complexity of the function we are trying to estimate. Combining regularization networks and least-squares-based policy evaluation, we are able to develop fast and efficient reinforcement learning algorithms that can scale to high-dimensional state-spaces without requiring manual tuning or engineering of basis functions. As for practical applications, I will consider various challenging real-world tasks, such as RoboCup-Keepaway, where I will demonstrate that this solution can achieve a performance superior to convential methods.

Speaker's Bio:

Tobias Jung recently joined the University of Liege as a postdoc researcher. Before coming to Liege, he was postdoc researcher at UT Austin. His general research interests lie at the intersection of machine learning and decision making/optimal control with a particular focus on reinforcement learning.