



UMass Lowell Computer Science Colloquium Announcement

Speaker: Michael T. Rosenstein
Department of Computer Science, UMass Amherst

Date: October 30, 2002

Time: 3:30pm - 4:30pm

Place: Olsen Hall 311
Refreshments are served at 3:00pm

SUPERVISED ACTOR-CRITIC REINFORCEMENT LEARNING

Reinforcement learning (RL) methods are attractive for solving optimal control problems when accurate models are unavailable. For many such problems, however, RL alone is impractical and the associated learning problem must be structured somehow to take advantage domain knowledge. In this talk I'll describe the use of domain knowledge in the form of a supervisor that generates control inputs (i.e., decisions or actions) in parallel with a reinforcement learning system. The basic challenge for the learning system is to combine two sources of feedback: error information derived from the supervisor and evaluative reinforcement supplied by the environment.

Almost all RL methods that incorporate supervisory information do so by modifying a "value function" in some fashion, and embedded within the value function is an implicit representation of the needed control policy. The alternative described in this talk is the combination of supervised learning with an actor-critic architecture. Actor-critic architectures--with separate data structures for the control policy and the value function--are particularly suited for decision problems that involve real-valued actions. In this talk I'll describe recent efforts at extending the standard actor-critic framework to model the supervisor in a general way. I'll show specific examples where the actor updates its policy in accordance with gradient information from both critic and supervisor. I'll also describe ongoing work to make these techniques suitable for remote operation of robot manipulators, such as NASA's Robonaut.

Colloquium Coordinator: Jie Wang, wang@cs.uml.edu