Linking Histories with Horizons: Modeling non-Markovian Reinforcement Learning Domains with Echo State Networks

Keith Bush, Colorado State University

A driving force in reinforcement learning research is the growing need for intelligent, autonomous control strategies that operate in real-world domains. Operating in the real-world assumes a high-dimensional, nonlinear, dynamic, and continuous-valued state-space that is only partially observable through signals containing some degree of noise; a mathematically disagreeable place, indeed. Applying reinforcement learning to such hard problems, therefore, requires function approximation techniques that learn even when the domain's simplifying assumptions---the popular ``gridworld'', for example---have been relaxed and replaced by parameterized, discrete-time dynamic systems which are both nonlinear and non-Markovian.

These function approximation requirements are, unfortunately, difficult to fulfill due to two mathematically competing interests. First, models of nonlinear dynamic systems are notoriously sensitive to changes in the system's dynamics. Second, to construct an efficient control strategy, reinforcement learning must repeatedly perturb system dynamics without respect to model sensitivity. In this presentation, I argue that reinforcement learning in non-Markovian domains requires a fresh modeling perspective; underlying models must be adaptive, accurate, and stable throughout the entire learning process.

A recent architecture with the potential to satisfy these criteria is the Echo State Network, or equivalent Liquid State Machine, developed, respectively, by Jaeger (2001) and Maass (2002). In this presentation, I explore the use of ESNs as models for reinforcement learning domains. Drawing on lessons learned from incorporating the ESN into reinforcement learning, I will propose current and future directions of ESN research.