

Reinforcement learning in Factories: The Auton Project

Andrew W. Moore

Robotics Institute and School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
awm@cs.cmu.edu

Abstract

Factories are fascinating test-beds for integrated learning systems. In recent years their sensory capabilities have, in many cases, been advanced and integrated so that data from all over the plant is available in real time over a LAN. Here we discuss how reinforcement learning, and related machine learning methods, can take advantage of this information to learn to improve performance, to adapt to change, and to exploit databases of historical records or similar processes in different plants.

This research involves very autonomous algorithms which monitor streams of sensory data, spot regularities, and design new control rules accordingly. Autonomous process control can lead to higher quality products, cheaper to manufacture products, and greatly increases flexibility in the running of plants. The modification of old processes and introduction of new processes has reduced reliance on calibration, manual redesign of controllers, and the time of process engineers.

We will talk about memory-based learning, in which the factory builds models of itself with data using nearest neighbor style approximations. We will then describe how these models can be exploited by certain classes of fast reinforcement learning algorithms. We'll outline new investigations into how factories can decide which models are important to learn, and how to actively seek data for those new models.

This arena may provide an important niche for numerical artificial intelligence to become a widespread method of choice throughout manufacturing industries. I will discuss experiences working with departments of two large corporations in putting these systems in place.

I will outline our ongoing project ("Auton") for an integrated autonomous system for robustly learning from process data and subsequently improving the process control. A number of issues arise in such a system, interesting from both a computational and statistical standpoint.

The software systems we have fielded have an unfortunate drawback: they probably have no cognitive or neural plausibility. I will conclude the talk by appealing for help in answering the question "Do these fielded learning systems have anything to say to Cognitive Science?".