Efficient reinforcement learning for control

Lucian Busoniu (TU Delft)

September 24, 2010

Reinforcement learning (RL) can be used to learn how to control complex systems, characterized by highly nonlinear, continuous-variable, potentially stochastic and unknown dynamics. In order to deal with the continuous variables of these systems, compact, approximate representations of the functions of interest are required. This presentation highlights some of our work in the field of approximate RL for control.

The first part of the talk concerns the case in which a model of the dynamics is available, and presents an algorithm that is provably convergent and consistent (that is, asymptotically obtains the optimal solution as the approximation accuracy increases). We then move to the online model-free setting, in which a solution must be learnt from data obtained by interacting with the system. Since interaction is costly, such algorithms must learn fast, using as little data as possible. Several methods to obtain data-efficient algorithms are proposed, focusing in particular on an algorithm called online least-squares policy iteration. We demonstrate these methods in real-life learning control experiments. The presentation closes with an outlook of promising ideas to be explored.