

Bayesian Methods in Reinforcement Learning

Wednesday, June 20th, 2007

ICML-07 tutorial

Corvallis, Oregon, USA

Pascal Poupart (Univ. of Waterloo)

Mohammad Ghavamzadeh (Univ. of Alberta)


Yaakov Engel (Univ. of Alberta)

Motivation

- ***Why a tutorial on Bayesian Methods for Reinforcement Learning?***
- Bayesian methods sporadically used in RL
- Bayesian RL can be traced back to the 1950's
- Some advantages:
 - Uncertainty fully captured by probability distribution
 - Natural optimization of the exploration/exploitation tradeoff
 - Unifying framework for plain RL, inverse RL, multi-agent RL, imitation learning, active learning, etc.

Goal

- Add another tool in the toolbox of Reinforcement Learning researchers



Outline

- Intro to RL and Bayesian Learning
- History of Bayesian RL
- Model-based Bayesian RL
 - Prior knowledge, policy optimization, discussion, Bayesian approaches for other RL variants
- Model-free Bayesian RL
 - Gaussian process temporal difference, Gaussian process SARSA, Bayesian policy gradient, Bayesian actor-critique algorithms
- Demo: control of an octopus arm