
MDPs with Unawareness in Robotics

Nan Rong Joseph Y. Halpern Ashutosh Saxena
Computer Science Department
Cornell University
Ithaca, NY 14853
{rongnan | halpern | asaxena}@cs.cornell.edu

Abstract

We formalize decision-making problems in robotics and automated control using continuous MDPs and actions that take place over continuous time intervals. We then approximate the continuous MDP using finer and finer discretizations. Doing this results in a family of systems, each of which has an extremely large action space, although only a few actions are “interesting”. We can view the decision maker as being unaware of which actions are “interesting”. We model this using MDPU, MDPs with unawareness, where the action space is much smaller. As we show, MDPU can be used as a general framework for learning tasks in robotic problems. We prove results on the difficulty of learning a near-optimal policy in an MDPU for a continuous task. We apply these ideas to the problem of having a humanoid robot learn on its own how to walk.

This poster from the UAI 2016 conference was given as an invited presentation at the Bayesian Modeling Applications Workshop.