Thesis title and Abstract

Name: V VARAGAPRIYA

Department: Mathematics

Title

CONSTRAINED MARKOV DECISION PROCESSES UNDER UNCERTAIN RUNNING COSTS AND TRANSITION PROBABILITIES

Abstract

In this thesis, we consider a constrained Markov decision process (CMDP) where running costs and transition probabilities are uncertain. We model the uncertain running costs and transition probabilities using robust optimization, chance-constrained optimization, and distributionally robust chance-constrained optimization frameworks.

We start with a robust CMDP problem with known transition probabilities, and uncertain running cost vectors that are known to belong to an uncertainty set. We consider polytopic, ellipsoidal, and semidefinite cone uncertainty sets and equivalently reformulate the robust CMDP problem as a linear programming (LP), second-order cone programming (SOCP), and semidefinite programming (SDP) problem, respectively. As an application, we propose a variant of a machine replacement problem and perform numerical experiments on randomly generated instances of different sizes. Furthermore, we study a robust CMDP problem under uncertain transition probabilities and known running cost vectors. Under the assumption that the uncertainty is

driven by a single parameter belonging to an interval, we equivalently reformulate the problem into a bilinear programming (BP) problem. By exploiting the structure of the BP problem, we construct an LP-based algorithm to find its global optimal solution. We propose a sufficient condition under which an optimal policy of the robust CMDP problem is unaffected by uncertainty. The numerical experiments are performed on a machine replacement problem and on randomly generated CMDP problems of various sizes using LP-based algorithm as well as Gurobi solver. We observe that in some cases, the LP-based algorithm outperforms the Gurobi solver. We extend this work to the case where the uncertainty in the transition probabilities is driven by multiple parameters belonging to a polytopic uncertainty set. Using the duality theory of LP problem, we equivalently reformulate the problem into a BP problem and perform numerical experiments on randomly generated CMDP problems of various sizes, using the Gurobi solver.

As an alternative approach, we use a chance-constrained optimization framework to address uncertainties in the running costs and transition probabilities. This leads to the formulation of a joint chance-constrained Markov decision process (JCCMDP). We first consider a JCCMDP problem where running cost vectors are defined using random vectors and transition probabilities are known. We assume that the random cost vectors are elliptically distributed and the dependence among the random constraints is driven by a Gumbel–Hougaard copula. Under these assumptions, we present two SOCP approximations such that their optimal values provide upper and lower bounds to the optimal cost of the JCCMDP problem. The numerical experiments are performed on a queueing control problem, a budget optimization problem in advertising platforms, and randomly generated CMDP problems of various sizes. We also consider the case when the distributions of running cost vectors are not known. In this case, we use classical probability inequalities, namely, one-sided Chebyshev, Bernstein, and Hoeffding inequalities, to construct convex programming problems, each of whose optimal value provides an upper bound to the optimal cost of the associated JCCMDP problem. We propose an LP problem whose optimal value provides a lower bound to the optimal cost of the associated JCCMDP problem. In addition, we study a JCCMDP problem where running costs are known and transition probabilities have random perturbations. Under certain conditions on the random perturbation vector, we construct convex programming problems using Bernstein and Hoeffding inequalities, each of whose optimal value gives an upper bound to the optimal cost of the associated JCCMDP problem. Similar to the case of random costs, we propose an LP problem whose optimal value provides a lower bound to the optimal cost of the associated JCCMDP problem. The numerical experiments are performed on a queueing control problem and randomly generated CMDP problems of various sizes.

At the end, we consider a distributionally robust chance-constrained optimization framework to address uncertainties in the running costs and transition probabilities. This approach is used when the exact probability distribution is not known and the optimization model considers the worst-case scenario of the underlying distribution. This leads to the formulation of a distributionally robust joint chance-constrained Markov decision process (DRJCCMDP). We first consider a DRJCCMDP problem under known transition probabilities and random running cost vectors whose distribution is only partially known. The only information we have about the distribution is that it belongs to an uncertainty set which is constructed using the full or partial information available on the first two moments. We consider three different moments-based uncertainty sets, and for each case, we present convex approximations whose optimal values provide upper and lower bounds to the optimal cost of the original problem. Furthermore, we study a DRJCCMDP problem under random transition probabilities and known running costs. We assume that the transition probability vector follows a discrete distribution and has a finite support. Using the estimates of the first two moments, we construct two moments-based uncertainty sets similar to the case of random costs. We show that the DRJCCMDP problem can be equivalently reformulated either as a mixed integer bilinear programming (MIBP) problem or a mixed integer semidefinite programming (MISDP) problem with bilinear constraints. In both cases, we perform numerical experiments on randomly generated CMDP problems of various sizes.