
MARINA GARROTE-LOPEZ, University of British Columbia
Algebraic Optimization of Sequential Decision Problems

In this talk, we study the optimization of the expected long-term reward in finite partially observable Markov decision processes over the set of stationary stochastic policies. We focus on the case of deterministic observations, where the problem is equivalent to optimizing a linear objective subject to quadratic constraints. We characterize the feasible set of this problem as the intersection of a product of affine varieties of rank one matrices and a polytope, which allows us to obtain bounds on the number of critical points of the optimization problem. Finally, we will explain some experiments in which we solve the KKT equations or the Lagrange equations over different boundary components of the feasible set to solve the optimization problem and compare the result to the theoretical bounds and to other constrained optimization methods.