




Poster M9: *Optimistic Linear Programming gives Logarithmic Regret for Irreducible MDPs*

Ambuj Tewari & Peter L. Bartlett


Learner in an unknown MDP



Competitor knows the MDP



\geq



- OLP is based on Burnetas-Katehakis algorithm
- Is computationally more efficient
- Similar in flavor to Auer-Ortner algorithm (NIPS 2006)

- Doesn't have to know support sets of transition distributions
- Has a simpler analysis
- Has better (S vs. S^5) state space size dependence compared to AO