

Time-Constrained Robust MDPs

Tuesday, June 18, 2024 3:00 PM (20 minutes)

Robust reinforcement learning is essential for deploying reinforcement learning algorithms in real-world scenarios where environmental uncertainty predominates.

Traditional robust reinforcement learning often depends on rectangularity assumptions, where adverse probability measures of outcome states are assumed to be independent across different states and actions.

This assumption, rarely fulfilled in practice, leads to overly conservative policies.

To address this problem, we introduce a new time-constrained robust MDP (TC-RMDP) formulation that considers multifactorial, correlated, and time-dependent disturbances, thus more accurately reflecting real-world dynamics. This formulation goes beyond the conventional rectangularity paradigm, offering new perspectives and expanding the analytical framework for robust RL.

We propose three distinct algorithms, each using varying levels of environmental information, and evaluate them extensively on continuous control benchmarks.

Our results demonstrate that these algorithms yield an efficient tradeoff between performance and robustness, outperforming traditional deep robust RL methods in time-constrained environments while preserving robustness in classical benchmarks.

This study revisits the prevailing assumptions in robust RL and opens new avenues for developing more practical and realistic RL applications.

Presenter: ZOUTINE, Adil (SUPAERO)

Session Classification: Poster session