

Episodic Bayesian Optimal Control with Unknown Randomness Distributions

mercredi 30 juillet 2025 11:55 (25 minutes)

Stochastic optimal control with unknown randomness distributions has been studied for a long time, encompassing robust control, distributionally robust control, and adaptive control. We propose a new episodic Bayesian approach that incorporates Bayesian learning with optimal control. In each episode, the approach learns the randomness distribution with a Bayesian posterior and subsequently solves the corresponding Bayesian average estimate of the true problem. The resulting policy is exercised during the episode, while additional data/observations of the randomness are collected to update the Bayesian posterior for the next episode. We show that the resulting episodic value functions and policies converge almost surely to their optimal counterparts of the true problem if the parametrized model of the randomness distribution is correctly specified. With an approximation commonly used in statistical analysis, we further show that the asymptotic convergence rate of the episodic value functions is of the order $O(N^{-1/2})$, where N is the number of episodes given that only one data point is collected in each episode. We develop an efficient computational method based on stochastic dual dynamic programming (SDDP) for a class of problems that have convex cost functions and linear state dynamics. Our numerical results on a classical inventory control problem verify the theoretical convergence results, and numerical comparison with two other methods demonstrate the effectiveness of the proposed Bayesian approach.

Author: ZHOU, Enlu

Co-auteurs: SHAPIRO, Alexander; LIN, Yifan; WANG, Yuhao

Orateur: ZHOU, Enlu

Classification de Session: Mini-symposium

Classification de thématique: Sequential decision making under uncertainty