

Reinforcement Learning Methods for Risk-Sensitive Sequential Decision Making

vendredi 1 août 2025 09:15 (1 heure)

This talk surveys recent developments in reinforcement learning (RL) methods for risk-aware model-free decision-making in Markov decision processes (MDPs). In the discounted setting, we adapt two popular risk neutral RL methods to account for risk aversion. The first approach minimizes a dynamic utility-based shortfall risk measure, while the other optimizes a specific quantile of the total discounted cost. We then present an RL framework for average-cost MDPs that incorporates dynamic risk measures. Together, these contributions represent a significant step toward scalable, risk-aware, model-free, sequential decision-making methods. The presentation will highlight the theoretical motivations, convergence guarantees, and empirical performance of these algorithms, offering insights into their applicability in finance and beyond.

Author: DELAGE, Erick (HEC Montréal)

Co-auteurs: DERMAN, Esther (Université de Montréal); M. GHAVAMZADEH, Mohammad (Amazon); M. HAU, Jia Lin (University of New Hampshire); LI, Jonathan (Telfer School of Management, University of Ottawa); MARZBAN, Saeed (HEC Montréal); M. PETRIK, Marek (University of New Hampshire); WANG, Weikai (HEC Montréal)

Orateur: DELAGE, Erick (HEC Montréal)

Classification de Session: Plenary session

Classification de thématique: Stochastic Programming