

HONG KONG INSTITUTE FOR DATA SCIENCE POSTDOC ASSOCIATES AND STUDENTS SEMINAR

Robust Satisficing MDPs

Mr Haolin RUAN

Ph.D. Student, School of Data Science
City University of Hong Kong



Date: 26 February 2024 (Monday)

Time: 3:00pm - 4:00pm

Zoom meeting ID: 985 7555 6083

Seminar link: <https://cityu.zoom.us/j/98575556083>



Abstract

Despite being a fundamental building block for reinforcement learning, Markov decision processes (MDPs) often suffer from ambiguity in model parameters. Robust MDPs are proposed to overcome this challenge by optimizing the worst-case performance under ambiguity. While robust MDPs can provide reliable policies with limited data, their worst-case performances are often overly conservative, and so they do not offer practical insights into the actual performance of these reliable policies. This paper proposes robust satisficing MDPs (RSMDPs), where the expected returns of feasible policies are softly-constrained to achieve a user-specified target under ambiguity. We derive a tractable reformulation for RSMDPs and develop a first-order method for solving large instances. Experimental results demonstrate that RSMDPs can prescribe policies to achieve their targets, which are much higher than the optimal worst-case returns computed by robust MDPs. Moreover, the average and percentile performances of our model are competitive among other models. We also demonstrate the scalability of the proposed algorithm compared with a state-of-the-art commercial solver.

Biography

Mr. Haolin Ruan is a fourth-year Ph.D. student from the School of Data Science at the City University of Hong Kong, supervised by Dr. Chin Pang Ho. He received his bachelor's degree in computer science and technology from Jinan University in 2016, and his master's degree in statistics from Jinan University in 2020. His research interests include robust optimization, reinforcement learning, and optimization algorithms.