THE UNIVERSITY OF CHICAGO | DEPARTMENT OF STATISTICS

MASTER'S THESIS PRESENTATION

YUNHONG WANG
Department of Statistics
The University of Chicago

Understanding the Sample Complexity of Offline Robust Reinforcement Learning With A General Uncertainty Set

WEDNESDAY, February15, 2023, at 4:30 PM
Zoom Meeting

ABSTRACT

Offline reinforcement learning, which aims to use historical observed data without interacting with the environment has become more and more important on both theoretical and practical sides. However, existing literature heavily relies on the assumption that there is no shift between the transition dynamics of the environment generating the data and the environment we want to deploy our policies and evaluate the performance. This raises a concern about the robustness of the output policy. In this paper, we consider a distributional robust finite-horizon offline reinforcement learning problem, with several different kinds of uncertainty measures with respect to the tabular Markov decision process. Our work extends the result of (Shi 22.) to the general f- divergence case, and shows that a model-based policy can achieve Hoeffding's type of performance bound by introducing the pessimism factor.