

WHEN ARE OFFLINE TWO-PLAYER ZERO-SUM MARKOV GAMES SOLVABLE?

Simon S. Du

Assistant Professor - CSE

ABSTRACT

We study what dataset assumption permits solving offline two-player zero-sum Markov game. In stark contrast to the offline single-agent Markov decision process, we show that the single strategy concentration assumption is insufficient for learning the Nash equilibrium (NE) strategy in offline two-player zero-sum Markov games. On the other hand, we propose a new assumption named unilateral concentration and design a pessimism-type algorithm that is provably efficient under this assumption. In addition, we show that the unilateral concentration assumption is necessary for learning an NE strategy. Furthermore, our algorithm can achieve minimax sample complexity without any modification for two widely studied settings: dataset with uniform concentration assumption and turn-based Markov game. Our work serves as an important initial step towards understanding offline multi-agent reinforcement learning.

DATE/TIME

 Nov. 15, 2022

 1:30 pm - 2:20 pm PST

 MEB 242



SPEAKER BIO

Simon S. Du is an assistant professor in the Paul G. Allen School of Computer Science & Engineering at the University of Washington. His research interests are broadly in machine learning, such as deep learning, representation learning, and reinforcement learning. Prior to starting as faculty, he was a postdoc at the Institute for Advanced Study of Princeton. He completed his Ph.D. in Machine Learning at Carnegie Mellon University. Simon's research has been recognized by a Samsung AI Researcher of the Year Award, an NSF CAREER award, an Nvidia Pioneer Award, a AAAI New Faculty Highlights, and a Distinguished Dissertation Award honorable mention from CMU.