The Hong Kong University of Science and Technology

Seminar on Business Data Science

Department of ISOM

Ensuring Fairness in Sequential Decision Making:
Counterfactual Fairness in Reinforcement Learning

by
Dr. Zhenke WU
University of Michigan

Date: 19 January 2026 (Monday)
Time: 11:00am – 12:00nn
Venue: Room 1027 (LSK Business Building)

Abstract

As healthcare systems move to deploy reinforcement learning (RL) for dynamic resource allocation, significant concerns have been raised about the impact of these tools on health services access and equity. While RL may improve service personalization, black-box optimization based on historical data can inadvertently lead to withholding needed treatment from vulnerable subgroups. We investigated the impact on fairness in treatment allocation using clinical data from a study that tested an RL-supported intervention for managing chronic pain and opioid misuse risk. In this talk, I show how to (1) audit standard RL policies for allocation bias across patient subgroups, and (2) validate a deployable methodology for correcting identified disparities without compromising clinical outcomes. We show that standard methodologies seeking to achieve "fairness through unawareness" are insufficient. Our framework evaluates and patches these vulnerabilities prior to deployment, satisfying a key requirement for the ethical scaling of AI in healthcare. Although motivated by health applications, the proposed framework applies in broader operations management contexts.

Bio

Dr. Zhenke Wu completed a BS in Math at Fudan University and a PhD in Biostatistics from the Johns Hopkins University and then stayed at Hopkins for his postdoctoral training. He is currently a tenured Associate Professor of Biostatistics at University of Michigan, Ann Arbor, and a faculty affiliate in Michigan Institute for Data and AI in Society (MIDAS).

Dr. Wu's research currently focuses on study design, causal and reinforcement learning methods for evaluating sequential interventions that tailor to individuals' changing circumstances such as in interventional digital health studies, and valid statistical inference using AI/ML predictions/synthesized data.

Dr. Wu serves as an Associate Editor for Annals of Applied Statistics, Biostatistics, Journal of Royal Statistical Society: Series A (JRSS-A) and a standing Statistical Consultant and Reviewer for New England Journal of Medicine - Artificial Intelligence (NEJM-AI).

All interested are welcome!
For details, please contact ISOM Department.