

# Operations Management & Statistics

## Rotman Recruiting Seminar

FRI., JAN. 20, 2023 | 1:30 PM | ROOM LL 1020

### Fast Rates for Contextual Linear Optimization

Yichun Hu, PhD Candidate, Operations Research  
Cornell University

#### Abstract

Incorporating side observations in decision making can reduce uncertainty and boost performance, but it also requires we tackle a potentially complex predictive relationship. While one may use off-the-shelf machine learning methods to separately learn a predictive model and plug it in, a variety of recent methods instead integrate estimation and optimization by fitting the model to directly optimize downstream decision performance. Surprisingly, in the case of contextual linear optimization, we show that the naïve plug-in approach actually achieves regret convergence rates that are significantly faster than methods that directly optimize downstream decision performance. We show this by leveraging the fact that specific problem instances do not have arbitrarily bad near-dual-degeneracy. While there are other pros and cons to consider as we discuss and illustrate numerically, our results highlight a nuanced landscape for the enterprise to integrate estimation and optimization. Our results are overall positive for practice: predictive models are easy and fast to train using existing tools, simple to interpret, and, as we show, lead to decisions that perform very well. This is a joint work with Nathan Kallus and Xiaojie Mao.

#### Bio

I am a final-year Ph.D. candidate in Operations Research at Cornell University, advised by Professor Nathan Kallus at Cornell Tech. I obtained my B.S. in Mathematics and Applied Mathematics and B.A. in Economics from Peking University in 2017.

My research lies at the intersection of machine learning, stochastic optimization, and statistics. Specifically, I study the statistical limits of data-driven decision-making, working to understand when and how we can design fast and reliable personalized algorithms. My works draw inspiration from important real-world applications, including healthcare and online platforms. I strive to identify cutting-edge problems and generate real-world impact through industry collaborations. In Summer 2021, I worked as a research engineer intern in the Adaptive Experimentation team at Meta, where I developed new methods for multi-objective adaptive experiments. In summer 2020, I worked as a data scientist intern at Google Play, where I investigated causal models for customer retention.

**Note:** OM&S PhD students are reminded to stay in the classroom for the full duration of the seminar.