

YICHUN HU

114 Feeney Way, Ithaca NY 14853
yh767@cornell.edu || yichunhu.github.io

ACADEMIC POSITIONS

Cornell University, SC Johnson Graduate School of Management
Assistant Professor of Operations, Technology, & Information Management

Ithaca, NY
07/2023 -

EDUCATION

Cornell University/Cornell Tech
Ph.D. in Operations Research, Ph.D. minor in Applied Mathematics
Advisor: Nathan Kallus

New York, NY
08/2017 - 05/2023

Peking University
B.S. in Mathematics and Applied Mathematics
B.A. in Economics

Beijing, China
09/2013 - 06/2017
09/2014 - 06/2017

RESEARCH INTERESTS

Data-driven Decision-Making, Adaptive Experimentation, Online Platforms, Contextual Bandit and Reinforcement Learning, Stochastic Optimization.

RESEARCH PAPERS

Publications

1. **Y. Hu**, N. Kallus, X. Mao (2022) Smooth Contextual Bandits: Bridging the Parametric and Non-differentiable Regret Regimes. *Operations Research* 70(6):3261-3281.
 - **Finalist, INFORMS Applied Probability Society 2020 Best Student Paper Competition.**
 - Preliminary version appeared in 33rd Conference on Learning Theory (*COLT*), 2020.
2. **Y. Hu**, N. Kallus, X. Mao (2022) Fast Rates for Contextual Linear Optimization. *Management Science* 68(6):4236-4245 (*fast track*).
3. M. Garrard, H. Wang, B. Letham, Z. Wang, Y. Huang, **Y. Hu**, C. Zhou, N. Zhou, E. Bakshy (2021) Practical Policy Optimization with Personalized Experimentation. *NeurIPS Workshop on Causal Inference Challenges in Sequential Decision Making*.

Working Papers

4. **Y. Hu**, N. Kallus, M. Uehara. Fast Rates for the Regret of Offline Reinforcement Learning. *Minor Revision at Mathematics of Operations Research*.
 - Preliminary version appeared in 34th Conference on Learning Theory (*COLT*), 2021.
5. **Y. Hu**, N. Kallus. DTR Bandit: Learning to Make Response-Adaptive Decisions with Low Regret. *Moderate Revision at Journal of the American Statistical Association*.

TEACHING EXPERIENCE

Cornell University, Instructor
NCC 5010: Data Analytics and Modeling

Spring 2024

Cornell University, Teaching Assistant
ORIE 5751: Learning, Inference, and Decision Making from Data
CS 5785: Applied Machine Learning
ORIE 4360: A Mathematical Examination of Fair Representation
ORIE 3510: Introduction to Engineering Stochastic Processes I
ORIE 5600: Financial Engineering with Stochastic Calculus I

Spring 2023
Fall 2019
Fall 2018
Spring 2018
Fall 2017

SELECTED HONORS

FODSI Postdoctoral Fellowship (declined)	2023
Rising Scholar, Chicago Booth	2022
Finalist, Applied Probability Society Best Student Paper Competition, INFORMS	2020
Sherri Koenig Stuewer Graduate Fellowship, Cornell University	2018
Excellent Graduate Award, Peking University	2017

GRANTS AND FUNDINGS

Cornell-HKUST Global Strategic Collaboration Award (CO-PI: Ruohan Zhan. Duration: 01/01/2024-12/31/2024. \$10k total.)

INDUSTRY EXPERIENCE

Facebook (Meta)	<i>Menlo Park, CA (Remote)</i>
Research Engineer Intern, Core Data Science (Adaptive Experimentation)	05/2021-08/2021
• Worked on multi-objective adaptive experimentation algorithms.	
Google	<i>Mountain View, CA (Remote)</i>
Data Scientist Intern, Google Play	05/2020-08/2020
• Worked on causal methods to analyze the impact of app usage on the retention rate of Google Play Pass.	

PROFESSIONAL SERVICE

- **Journal Reviewer:** Operations Research (3), Journal of Machine Learning Research, Transactions on Machine Learning Research, IEEE Transactions on Information Theory, Omega
- **Conference Reviewer & PC Member**
 - International Conference on Machine Learning (ICML) 2020-2023
 - International Conference on Learning Representations (ICLR) 2021, 2024
 - International Conference on Artificial Intelligence and Statistics (AISTATS) 2021-2024
 - Conference on Neural Information Processing Systems (NeurIPS) 2021-2023
 - ACM conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO) 2022
 - NeurIPS Workshop on Causal Inference Challenges in Sequential Decision Making 2021
- **Session Chair**
 - INFORMS 2020 General Session (Stochastic Bandits)
 - INFORMS 2023 General Session (Learning and Experimentation in Markets)
 - INFORMS 2023 General Session (Learning and Decision-Making for Platforms)
- Cornell University ORGA (Operations Research Graduate Association) Tech Liaison 2019-2020

INVITED PRESENTATIONS

Fast Rates for Contextual Linear Optimization	
INFORMS Annual Meeting, Phoenix, AZ	10/2023
University of Illinois Urbana-Champaign - Industrial & Enterprise Systems Engineering, Champaign, IL	02/2023
Massachusetts Institute of Technology - Sloan School of Management, Cambridge, MA	02/2023
University of Oxford - Saïd Business School, Virtual	01/2023
University of Toronto - Rotman School of Management, Toronto, Canada	01/2023
Columbia University - Industrial Engineering and Operations Research, New York, NY	01/2023
University of Illinois Chicago - Information and Decision Sciences, Chicago, IL	01/2023
Johns Hopkins University - Carey Business School, Baltimore, MD	01/2023
Duke University - Fuqua School of Business, Durham, NC	01/2023
Carnegie Mellon University - Tepper School of Business, Pittsburgh, PA	12/2022
University of British Columbia - Sauder School of Business, Vancouver, Canada	12/2022
Cornell University - SC Johnson College of Business, Ithaca, NY	12/2022
Northwestern University - Industrial Engineering & Management Sciences, Evanston, IL	12/2022
University College London - School of Management, Virtual	11/2022
London Business School, Virtual	11/2022
University of Texas at Austin - McCombs School of Business, Austin, TX	11/2022
INFORMS Annual Meeting, Indianapolis, IN	10/2022
Cornell ORIE Young Researchers Workshop, Ithaca, NY	10/2022

NYC Operations Day (Poster), New York, NY	04/2022
ORIE PhD Colloquium at Cornell Tech, New York, NY	04/2022
INFORMS Optimization Society Conference, Greenville, SC	03/2022
Cornell ORIE Young Researchers Workshop (Poster), Ithaca, NY	10/2021

Fast Rates for the Regret of Offline Reinforcement Learning

RL Theory Seminar, Virtual	11/2021
INFORMS Annual Meeting, Anaheim, CA	10/2021
16th INFORMS Workshop on Data Mining and Decision Analytics, Anaheim, CA	10/2021
34th Annual Conference on Learning Theory, Boulder, CO	08/2021

DTR Bandit: Learning to Make Response-Adaptive Decisions with Low Regret

INFORMS Annual Meeting, Virtual	11/2020
15th INFORMS Workshop on Data Mining and Decision Analytics, Virtual	11/2020

Smooth Contextual Bandits: Bridging the Parametric and Non-differentiable Regret Regimes

INFORMS Healthcare Conference, Toronto, Canada	07/2023
Rising Scholars Conference (Chicago Booth), Virtual	11/2022
33rd Annual Conference on Learning Theory, Virtual	07/2020
INFORMS Annual Meeting, Seattle, WA	10/2019
14th INFORMS Workshop on Data Mining and Decision Analytics, Seattle, WA	10/2019
Cornell ORIE Young Researchers Workshop, Ithaca, NY	10/2019

OTHERS

- *Programming*: Python (PyTorch), C/C++, R, Julia.
- *Languages*: Chinese (native), English (fluent), Japanese (N2)