

Yufeng Zheng

yufeng.zheng@mail.utoronto.ca

Education

University of Toronto

Ph.D. Candidate, Operations Management and Statistics

Ontario, Canada

2023–present

University of California, Berkeley

M.Sc., Industrial Engineering and Operations Research

California, United States

2019–2020

Shanghai Jiao Tong University

B.Eng., Industrial Engineering and Operations Research

Shanghai, China

2015–2019

Research Interests

Large language models and AI agents; reinforcement learning; online learning and adaptive experimentation; operations management; revenue management and pricing; data-driven decision-making; simulation and optimization.

Research

Publications

LOLA: LLM-Assisted Online Learning Algorithm for Content Experiments

Marketing Science, 2025; *third author*. Developed an LLM-assisted adaptive experimentation framework for content delivery. Analyzed 17,681 Upworthy A/B headline tests and integrated LLM predictions with Upper Confidence Bound allocation to improve engagement under limited traffic.

Robust Situational Reinforcement Learning in Face of Context Disturbances

ICML, 2023; *second author*. Developed the Robust Situational Markov Decision Process framework for reinforcement learning under contaminated context transitions. Introduced a softmin-smoothed robust Bellman operator and evaluated the method in robot control and inventory management tasks.

Mind Your Step: Continuous Conditional GANs with Generator Regularization

NeurIPS SyntheticData4ML Workshop, 2022; *co-first author*. Proposed a Lipschitz-style generator regularization term for conditional GANs to promote smoothness across neighboring conditions and reduce mode collapse. Evaluated the method on synthetic tasks and conditional time-series generation.

An Adaptive Deep RL Method for Non-Stationary Environments with Piecewise Stable Context

NeurIPS, 2022; *third author*. Studied reinforcement learning with piecewise-stable latent contexts. Proposed Segmented Context Belief Augmented Deep RL, combining change-point detection with belief-state inference for grid-world, MuJoCo, and real-time communication tasks.

A Doubly Stochastic Simulator with Applications in Arrivals Modeling and Simulation

Operations Research, 2021; *first author*. Proposed DS-WGAN, combining doubly stochastic process structure with Wasserstein GAN estimation for modeling and simulating non-stationary random arrival processes. Proved consistency and convergence rates and evaluated performance on synthetic and real-world data.

Demand Prediction, Predictive Shipping, and Product Allocation for Large-scale E-commerce

M&SOM Data-driven Research Challenge Finalist, 2018; *second author*. Analyzed a 130GB logistics and order dataset for demand prediction, predictive shipping, and warehouse product allocation. Used representation learning to model competitive and complementary product relationships.

Ongoing Research Projects

Formal Math Theorem Proving via LLM and Reinforcement Learning

University of Toronto. Fine-tune large language models with reinforcement learning to improve Lean-code generation for formal theorem proving.

Enhancing Discrete Event Simulation with Advanced Language Model Techniques

UC Berkeley. Develop methods for coding discrete-event simulations with LLMs, prompt engineering, and retrieval-augmented generation.

Revenue Management and Pricing for Machine Learning Products

University of Toronto. Develop algorithms for dynamic pricing and retraining of machine-learning products, balancing performance gains against retraining costs.

Queueing Foundation Model

University of Toronto. Develop a foundation-model framework for queueing, service, and resource-constrained discrete-event systems by representing arrivals, service, departures, routing, resource states, and policy actions as event-state-action trajectories. Study causal event Transformer models for forecasting, simulation, routing, scheduling, capacity planning, policy evaluation, and system identification.

Continuous-Time Assortment Dynamic Games

University of Toronto. Study dynamic pricing and assortment games in which multiple sellers compete on a shared platform, customers arrive according to a Poisson process, and each seller maximizes profit over continuous time. Develop equilibrium computation and PINN-like numerical methods for the resulting dynamic game.

Analyzing Equity in Private Tutoring Markets Post-Regulation

University of Toronto. Use mean-field game theory to study exam competition, regulation, fairness, and access in private tutoring markets.

Research Experience

Microsoft Research Asia

Beijing, China

Research Intern

2022–2023

- Developed a pretrained video foundation model for embodied agents, focusing on generalization and few-shot learning in new real-world environments.
- Implemented a reinforcement-learning control system for HVAC optimization, targeting energy efficiency and operational performance.
- Executed a multi-agent reinforcement learning algorithm for multi-SKU inventory management and supply-chain operations.

Relevant Coursework

University of Toronto: Online and Data-Driven Decision Making; Advanced Topics in Linear and Nonlinear Optimization; Machine Learning for Mathematical Optimization; Automated Reasoning with Machine Learning.

UC Berkeley: Mathematical Programming; Applied Stochastic Processes; Supply Chain and Logistics Management; Production Planning and Logistics Models; Network Flows and Graphs; Statistical Learning Theory; Deep Reinforcement Learning; Introduction to Machine Learning.

Shanghai Jiao Tong University: Calculus; Probability and Statistics; Linear Algebra; C++ Programming; Industrial Engineering; Microeconomics; Operations Research; Logistics and Supply Chain; Production Planning and Control; Stochastic Models.

Technical Skills

Programming and modeling: Python, PyTorch, MATLAB, C++, R, Gurobi, CPLEX.

Research strengths: LLM and agent workflows, reinforcement learning, machine learning, optimization, simulation, stochastic modeling, data-driven operations.