From Monolithic Reinforcement Learning to Distribution-aware Optimization: Adventures in Analyzing Transportation Systems

BIOGRAPHY

Cathy Wu is the Class of 1954 Career Development Professor at MIT, holding appointments in LIDS, CEE, and IDSS. She holds a Ph.D. in EECS from UC Berkeley, and B.S. and M.Eng. in EECS from MIT, and completed a Postdoc at Microsoft Research. She designs machine learning methods to solve optimization problems that impact daily life, with a particular focus on urban mobility. She is broadly interested in AI for sociotechnical systems and evidence-based decision making. Cathy is the recipient of the NSF CAREER (2023), the Ole Madsen Mentoring Award (2025), the IEEE ITS Best Dissertation Award (2019), and the CUTC Milton Pikarsky Memorial Award (2018). She serves on the Board of Governors for the IEEE ITSS, is an Associate Editor or Area Chair for ICML, NeurIPS, ICRA, and Transportation Research Part C, and served as Program Co-chair for RLC 2025. She is also the inaugural Chair and Co-founder of the REproducible Research In Transportation Engineering (RERITE) Working Group.

ABSTRACT

The engineering of modern transportation systems is severely constrained by the cost of repeatedly modeling and solving complex optimization problems. Yet this challenge presents a timely opportunity: to leverage artificial intelligence (AI) as a partner to optimization to enable more responsive, innovative, confident and accountable analysis. As a case study of using AI to inform transportation policy, the talk presents the first prospective impact assessment of city-scale eco-driving, which uses deep reinforcement learning (RL) to reveal that optimizing vehicle speeds at intersections can reduce emissions by 11–22% without sacrificing throughput or safety. At the same time, significant gaps remain before AI for Optimization is ready for practitioners; for instance, studies like the case study expose the brittleness of deep RL to small changes in network structure or demand. This motivates two emerging research directions: (1) task-space RL, which explicitly reasons about the training and generalization structure of problem distributions, leading to 10-30x more sample-efficient methods, and (2) learning-guided optimization, which uses machine learning to specialize classical optimization to problem distributions, leading to 2-10x faster methods. Together, these approaches demonstrate how AI for Optimization can transform how we design, operate, and govern complex systems—from transportation to logistics and beyond.