

Principled Preference-Guided Multi-Objective Learning: Navigating and Uniformly Traversing the Pareto Front



Invited Speaker

Lisha Chen

University of Rochester

Date: June 5, 2026 (Fri)

Time: 9:30am (HKT)

Zoom Meeting: 801 137 0362

Biography

Lisha Chen's research develops mathematical and algorithmic foundations for large-scale signal processing and machine learning, with a focus on multi-objective and multi-level optimization and learning theory. Her work addresses structured learning problems involving competing objectives, hierarchical decision-making, and constraints arising from real-world scientific and engineering systems. Bridging theory and practice, her research has contributed to IBM products and patents and provides foundational tools for building machine learning systems that can reason over multiple objectives, and adapt to new deployment conditions. Dr. Chen has organized and delivered tutorials on multi-objective learning at ICASSP and AAAI, and served as a reviewer or program committee member for leading conferences and journals in machine learning, signal processing, and optimization. Her achievements have been recognized through the IEEE Signal Processing Society Scholarship, and the IBM AIRC Fellowship.

Abstract

Many real-world learning systems must balance multiple competing objectives while respecting user-specified preferences. This talk presents recent work on preference-guided multi-objective learning, with a focus on how to efficiently navigate and explore trade-offs along the Pareto front. I will first discuss a framework for navigating the Pareto set toward user-preferred solutions, allowing preferences to directly guide optimization. I will then introduce SURF, a complementary method that steers scalarization weights to uniformly traverse the Pareto front, improving coverage of diverse trade-off solutions. The methods have been applied to multi-lingual speech recognition, multi-objective reinforcement learning, and multi-objective LLM alignment. Together, these works provide principled and computationally tractable tools for incorporating user preferences into modern multi-objective learning systems.