Abstract:

Optimal online portfolio selection, with respect to both regret and computational complexity, has been a classic open problem in online learning for over 30 years. Online learning quantum states with the log-loss is the quantum generalization of online portfolio selection and even more challenging. In this talk, I will introduce our recent achievements in developing computationally efficient algorithms with moderate regret rates for both problems. Specifically, I will introduce algorithms that may achieve the optimal regret rate while having per-iteration complexities that are nearly linear in dimension.