

Statistical Learning and Market Design

Xiaowu Dai

Abstract

We study the problem of decision-making in the setting of a scarcity of shared resources when the preferences of agents are unknown a priori and must be learned from data. Taking the two-sided matching market as a running example, we focus on the decentralized setting, where agents do not share their learned preferences with a central authority. Our approach is based on the representation of preferences in a reproducing kernel Hilbert space, and a learning algorithm for preferences that accounts for uncertainty due to the competition among the agents in the market. Under regularity conditions, we show that our estimator of preferences converges at a minimax optimal rate. Given this result, we derive optimal strategies that maximize agents' expected payoffs and we calibrate the uncertain state by taking opportunity costs into account. We also derive an incentive-compatibility property and show that the outcome from the learned strategies has a stability property. Finally, we prove a fairness property that asserts that there exists no justified envy according to the learned strategies.

This is a joint work with Michael I. Jordan.