DISSERTATION PROPOSAL

Su Jia

"Optimization In the Dark -- Exploration and Exploitation for Operations Problems"

Tuesday, November 17, 2020
11:00am EST
Zoom: https://cmu.zoom.us/j/97224447584?pwd=UXZIUk1rM0w1aC9ORUFjUjNVdi9RQT09

Spurred by the growing availability of data and recent advances in machine learning and optimization, there has been an increasing application of data analytics to problems in operations management. This dissertation focuses on applying techniques from machine learning, especially online learning, to fundamental operations problems such as pricing and personalization.

In the first chapter we study a novel variant of the classical optimal decision tree problem. The central task in active learning involves performing a sequence of tests to identify an unknown hypothesis that is drawn from a known distribution. This problem, known as optimal decision tree induction, has been widely studied for decades and the asymptotically best-possible approximation algorithm has been devised for it. We study a generalization where certain test outcomes are noisy. We design new approximation algorithms for both the non-adaptive setting, where the test sequence must be fixed a-priori, and the adaptive setting where the test sequence depends on the outcomes of prior tests.

In the second chapter we consider the Unimodal Multi-Armed Bandit problem where the goal is to find the optimal price under an unknown unimodal reward function, with an additional ``Markdown" constraint that requires the price exploration to be non-increasing. This markdown optimization problem faithfully models a single-product revenue management problem where, given infinite inventory, the objective is to adaptively reduce the price over a finite sales horizon to maximize expected revenues. Intuitively, the markdown constraint makes the tradeoff between exploitation and exploration harder to manage. Despite the numerous work in dynamic pricing under unknown models in recent years, an extremely basic question remains open: is there a ``separation" between markdown and non-markdown pricing? Our work presents the first affirmative answer to this question, by presenting tight regret bounds for markdown pricing, which are asymptotically higher than the known regret bounds for non-markdown dynamic pricing.

In the third chapter we propose an ongoing work, posed by Glance, a fast growing startup with one of the largest lockscreen content platforms in the world. The current personalization algorithm is based on a deep neural-network (DNN) model which predicts, for each pair of glance card and user, the expected click-through rate (``reward") when this glance card is assigned to this user, based on the features of the card and user, such as title, summary, language, past behavior etc. They then personalize the recommendation greedily solely based on the DNN prediction. In particular, they do not use the user feedback to update their ``belief" over the popularity of the card. The central question in this proposed work is, how much do we gain from utilizing the user feedback in online personalization? We plan to answer this question not only through the lens of theoretical analysis but also, more importantly, by field experiments.