

The University of Chicago Booth School of Business

Operations Management/Management Science Workshop

Tuesday, October 11, 2016

Professor Shipra Agrawal, Columbia University

Title: Online optimization and learning under long-term convex constraints and objective

Abstract:

Sequential decision making situations in real world applications often involve multiple long term constraints and nonlinear objectives. Some examples from online advertising include budget constraints, nonlinear under-delivery penalties, need for diversity in allocation, and managing risk on revenue. However, most standard models like multi-armed bandits, contextual bandits, online learning, etc., simply optimize the sum of revenue/costs/profits received in every step, and therefore the decisions obtained do not satisfy these long-term considerations. Some full information models like online packing and covering problems allow linear knapsack type constraints on consumption.

In this talk, I will discuss a series of recent results in which we have introduced novel techniques to handle a very general form of convex constraints and objective in sequential decision making, while providing provably near-optimal guarantees. Our primal-dual techniques are efficient to implement and employ fundamental concepts of Fenchel duality in convex programming, while using online learning methods as blackbox. Our techniques are generalizable, as we demonstrate by their successful application to a variety of sequential decision making models: multi-armed bandits, contextual bandits, online packing, and online stochastic convex programming.

References:

This talk is based on the following recent papers:

Shipra Agrawal and Nikhil R. Devanur, Bandits with concave rewards and convex knapsacks, in EC 2014, ACM conference on Economics and Computation, June 2014.

Shipra Agrawal and Nikhil R. Devanur, Fast algorithms for online stochastic convex programming, in SODA 2015

Shipra Agrawal, Nikhil R. Devanur, Lihong Li, An Efficient Algorithm for Contextual Bandits with Knapsacks, and an Extension to Concave Objectives, COLT 2016.

Shipra Agrawal and Nikhil R. Devanur, Linear Contextual Bandits with Knapsacks, NIPS 2016.

All the referenced papers are available at: <http://www.columbia.edu/~sa3305/>

Speaker Bio:

Shipra Agrawal is an Assistant Professor in the Department of Industrial Engineering and Operations Research, and member of Data Science Institute, at Columbia University. Her research spans several areas of optimization and machine learning, including data-driven optimization under partial, uncertain, and online inputs, and related concepts in learning, namely multi-armed bandits, online learning, and reinforcement learning. She is also interested in prediction markets and game theory. Application areas of her interests include internet advertising, recommendation systems, revenue management and resource allocation problems.

Shipra received her PhD in Computer Science from Stanford University in June 2011 under guidance of Prof. Yinyu Ye, and was a researcher at Microsoft Research India from July 2011 to August 2015.