USC Viterbi
School of Engineering
Ming Hsieh Department of Electrical and Computer Engineering

Seminar

Data Markets and Learning: Privacy Mechanisms and Personalization

Alireza Fallah
PhD Candidate
Department of Electrical Engineering and Computer Science
Laboratory for Information and Decision Systems (LIDS)
Massachusetts Institute of Technology

Monday, March 27, 2023
10:00am – 11:00am
EEB 132
Zoom Link: https://usc.zoom.us/j/93606233291?pwd=dGQxNWRZVmE1bzZvRVVYRTd1Mk1VQT09

Abstract: The fuel of machine learning models and algorithms is the data usually collected from users, enabling refined search results, personalized product recommendations, informative ratings, and timely traffic data. However, increasing reliance on user data raises serious challenges. A common concern with many of these data-intensive applications centers on privacy — as a user's data is harnessed, more and more information about her behavior and preferences is uncovered and potentially utilized by platforms and advertisers. These privacy costs necessitate adjusting the design of data markets to include privacy-preserving mechanisms.

This talk establishes a framework for collecting data of privacy-sensitive strategic users for estimating a parameter of interest (by pooling users' data) in exchange for privacy guarantees and possible compensation for each user. We formulate this question as a Bayesian-optimal mechanism design problem, in which an individual can share her data in exchange for compensation but at the same time has a private heterogeneous privacy cost which we quantify using differential privacy. We consider two popular data market architectures: central and local. In both settings, we use Le Cam's method to establish minimax lower bounds for the estimation error and derive (near) optimal estimators for given heterogeneous privacy loss levels for users. Next, we pose the mechanism design problem as the optimal selection of an estimator and payments that elicit truthful reporting of users' privacy sensitivities. We further develop efficient algorithmic mechanisms to solve this problem in both privacy settings.

Finally, we consider the case that users are interested in learning different personalized parameters. In particular, we highlight the connections between this problem and the meta-learning framework, allowing us to train a model that can be adapted to each user's objective function.

Bio: Alireza Fallah is a Ph.D. candidate at the department of Electrical Engineering and Computer Science (EECS) and the Laboratory for Information and Decision Systems (LIDS) at Massachusetts Institute of Technology (MIT). His research interests are machine learning theory, data market and privacy, game theory, optimization, and statistics. He has received a number of awards and fellowships, including the Ernst A. Guillemin Best MIT EECS M.Sc. Thesis Award, Apple Scholars in AI/ML Ph.D. fellowship, MathWorks Engineering Fellowship, and Siebel Scholarship. He has also worked as a research intern at the Apple ML privacy team. Before joining MIT, he earned a dual B.Sc. degree in Electrical Engineering and Mathematics from Sharif University of Technology, Tehran, Iran.

Host: Dr Mahdi Soltanolkotabi, soltanol@usc.edu