EPSTEIN INSTITUTE SEMINAR ▪ ISE 651

Does Your Training Data Violate My Privacy? A Near-Optimal Model Discrimination Method With Non-Disclosure

ABSTRACT - As we rely more on artificial intelligence to support human decisions, we are prone to serious societal risks including unfair/discriminatory outcomes and the violation of individuals' privacy. To limit such adverse impacts, regulatory entities should ensure that the data used by industries and technological platforms comply with specific privacy standards. Thus, it is crucial to develop procedures for auditing machine learning platforms to assess their compliance with regulations and guidelines. In this work, we study a hypothesis testing problem in which the auditor has to determine which of the two datasets has been used to train a given machine learning model. Making the initial steps towards answering this question in full generality, we first consider the case of a well-specified linear model with squared loss. We provide matching upper and lower bounds on the sample complexity (up to a constant factor). We then extend this result in two directions: (i) for the general parametric setup in asymptotic regime; (ii) for generalized linear models in the small-sample regime. Our test statistic is based on the classical Newton step on the empirical risk objective function.

This is a joint work with Dmitrii Ostrovskii (USC-ISE), Mohamed Ndaoud (USC-Math), and Adel Javanmard (USC-Marshall).

Dr. Meisam Razaviyayn
Assistant Professor
Dept. of Industrial & Systems Engineering, USC

SPEAKER BIO – Meisam Razaviyayn is an Assistant Professor of Industrial and Systems Engineering at the University of Southern California (with courtesy appointments at the Computer Science and Electrical Engineering Departments). Prior to joining USC, he was a postdoctoral research fellow at Stanford University. He received his PhD in Electrical Engineering with minor in Computer Science at the University of Minnesota. Meisam Razaviyayn is the recipient of IEEE Data Science Workshop Best Paper Award in 2019, the Signal Processing Society Young Author Best Paper Award in 2014, and the finalist for Best Paper Prize for Young Researcher in Continuous Optimization in 2013 and 2016. His research interests include the design and analysis of large scale optimization algorithms arise in modern data science era.

USC Viterbi School of Engineering
Daniel J. Epstein Department of Industrial and Systems Engineering

TUESDAY, MARCH 2, 2021
3:30 PM – 4:50 PM
ZOOM/ONLINE *PLEASE EMAIL OWH@USC.EDU FOR PASSWORD*