Yale University

Computer Science Talk

How to handle Biased Data and Multiple Agents in Machine Learning?

Host: Yang Cai

Manolis Zampetakis

Tuesday – March 8, 2022
4:00 p.m.

Zoom Presentation

Abstract: Modern machine learning (ML) methods commonly postulate strong assumptions such as: (1) access to data that adequately captures the application environment, (2) the goal is to optimize the objective function of a single agent, assuming that the application environment is isolated and is not affected by the outcome chosen by the ML system. In this talk I will present methods with theoretical guarantees that are applicable in the absence of (1) and (2) as well as corresponding fundamental lower bounds. In the context of (1) I will focus on how to deal with truncation and self-selection bias and in the context of (2) I will present a foundational comparison between two-objective and single objective optimization.

Bio: Manolis Zampetakis is currently a post-doctoral researcher at the EECS Department of UC Berkeley working with Michael Jordan. He received his PhD from the EECS Department at MIT where he was advised by Constantinos Daskalakis. He has been awarded the Google PhD Fellowship and the ACM SIGEcom Doctoral Dissertation Award. He works on the foundations of machine learning (ML), statistics, and data science, with focus on statistical analysis from systematically biased data, optimization methods for multi-agent environments, and convergence properties of popular heuristic methods.