University of California, Irvine
Statistics Seminar

Challenges and Opportunities in Differentially Private Machine Learning

Thomas Steinke
Senior Research Scientist
Google DeepMind

4-5 p.m.
Thursday, April 10, 2025
6011 Donald Bren Hall

Machine Learning ingests vast amounts of training data. That data may contain sensitive information and there is a risk that this information is revealed when the trained model is used. Differential privacy provides a mathematical framework for measuring and controlling this information leakage. In this talk I will briefly survey the basics of differential privacy and its application to machine learning and then discuss state-of-the-art methods for training language models on sensitive data.

Bio:
Thomas Steinke is a Senior Research Scientist at Google DeepMind in Mountain View, California. He completed his PhD at Harvard University in 2016 advised by Prof. Salil Vadhan. From 2016 until joining Google in 2020, he was at IBM Research – Almaden. Thomas' research centers around differential privacy, with a particular focus on foundational definitions, algorithms, lower bounds (i.e., privacy attacks), and connections to machine learning (particularly, generalization) and information theory.