

GREEN FAMILY LECTURE SERIES BY DR. CYNTHIA DWORK



Cynthia Dwork Harvard University

Cynthia Dwork, Gordon McKay Professor of Computer Science at Harvard, Affiliated Faculty at Harvard Law School and Department of Statistics, and Distinguished Scientist at Microsoft, is renowned for placing privacy-preserving data analysis on a mathematically rigorous foundation. She has also made seminal contributions in algorithmic fairness, cryptography, and distributed computing, and is the recipient of numerous awards including the IEEE Hamming Medal and the ACM Paris Kanellakis Theory and Practice Award. Dwork is a member of the US National Academy of Sciences and the US National Academy of Engineering, and is a Fellow of the American Academy of Arts and Sciences and the American Philosophical Society.

Fairness, Justice, and ..Algorithms?

Thursday, July 14, 2022 5:00PM
Moore Hall 100, UCLA

Examples of unfairness in Artificial Intelligence abound in the popular press, despite relentless growth in the power of machine learning. The emerging field of *algorithmic fairness* explores such questions as:

- What does it mean for an algorithm to be fair, and how can we translate this into conditions that can be mathematically specified and checked?
- What can cause algorithms to be unfair in the first place, and how can this be prevented?
- When, and for what kinds of fairness, is there a tradeoff between fairness and accuracy? How do the equations change in long lived systems?

We will survey the landscape of algorithmic fairness, highlighting some recent work in fair prediction that *improves* accuracy and has surprising applications to settings in which fairness itself is not a concern.

This lecture will be accessible to a broad scientific audience.

Differential Privacy and the US Census

Monday, July 18, 2022 4:30 PM
Moore Hall 100, UCLA
Reception to follow at IPAM

Traditional "anonymized" data can often be easily attacked to break privacy, as was dramatically demonstrated in the "Netflix Challenge" some years ago. Differential Privacy is a mathematically rigorous definition of privacy tailored to the statistical analysis of data. Differentially private systems provide useful statistics to the well-intentioned data analyst while simultaneously delivering a mathematical bulwark against arbitrarily powerful privacy adversaries. Its unique mathematical guarantees have led to its extensive deployment in industry, and led the US Census Bureau to adopt differential privacy as the disclosure avoidance methodology for the 2020 Decennial Census.

This lecture will be accessible to a general public audience.

