

Machine Learning and Economic Inequality

Maximilian Kasy

April 19, 2021

An exciting line-up

Day 1

- Abigail Adams-Prassl (Economics)
- Reuben Binns (Computer Science)
- Rediet Abebe (Computer Science)
- Joshua Loftus (Statistics)
- Paul Goldsmith-Pinkham (Economics)

Day 2

- Pauline Kim (Law)
- Jeremy Adams-Prassl (Law)
- Salome Viljoen (Law)
- Abigail Jacobs (School of Information)
- Moustapha Cisse (Computer Science)

How should we evaluate the social impact of AI?

- Automated decision-making is increasingly common in consequential settings: Hiring, consumer credit, bail setting, news feed selection, pricing, ...
- Much recent work is concerned with the “fairness” of such automated decisions: “Is unequal treatment rationalizable by profit maximization?”
- In this conference, the goal is to instead discuss the impact of AI on inequalities.
- This corresponds to two conflicting normative paradigms:
 1. **Just deserts** (e.g. Libertarianism):
“Does everyone get what they deserve, based on their merit?”
 2. **Consequentialism** (e.g. Utilitarianism, welfare economics):
“How does this policy / algorithm impact the wellbeing of those affected?”

Fairness versus equality

- Fairness is about **treating** people of the same “**merit**” independently of their **group** membership.
- Equality is about the (counterfactual / causal) **consequences** of an algorithm for the distribution of **welfare** of different **people**.

Examples when they are in conflict:

1. Increased surveillance / **better prediction** algorithms:
Lead to treatments more aligned with “merit.”
Good for fairness, bad for equality.
2. Affirmative action / **compensatory interventions** for pre-existing inequalities:
Bad for fairness, good for equality.

Means of prediction and agents of change

- ML and AI are largely about optimization.
 - But who's objective function is maximized?
 - Ownership and control of data and algorithms determine who gets to pick the objective.
 - Why is most of AI dedicated to maximizing ad-clicks?
 - How could alternative ownership structures look like?
- Who is the audience for work about the social impact of AI?
Who's perspective do we take?
 - Most commonly:
 - Corporations maximizing profit,
 - subject to constraints of avoiding litigation or bad press for discrimination.
 - How about addressing other stakeholders?
 - Workers and unions,
 - government regulators and policymakers,
 - consumers and NGOs, ...

Conference policies

- Talks are 40 minutes, followed by 10 minutes discussion time, and occasional breaks.
- Panel members can simply unmute themselves for questions and comments. Everyone else, please post your questions in the Q & A box, and Binta will read out your question or call on you to ask it.
- During talks, please limit yourself to clarifying questions. Reserve comments and criticisms for discussion time.
- Talks are live-streamed and recorded. They will be available for later watching on Youtube.

- Conference program:
https://maxkasy.github.io/home/ML_inequality_conference/
- Youtube channel:
<https://www.youtube.com/channel/UCB3VHmtU-Acta1o0wbzWaag>

Let's get started!