Hiring as Exploration

Danielle Li (MIT & NBER), Lindsey Raymond (MIT) and Peter Bergman (Columbia & NBER)

NBER Summer Institute

IT and Digitization

17 July 2017

Modern hiring ML favors traditionally successful groups

- Firms increasingly use algorithms to make HR decisions
  - Modern hiring ML typically relies on “supervised learning”
  - By design, supervised learning identifies and selects from traditionally successful groups

- Is this the right approach?

- Hiring as a dynamic learning problem
  - **Exploit** – selecting who you think is actually good (this is what supervised learning does)
  - **Explore** – selecting who you know less about to learn (this is what we do)

Exploration based ML improves quality and increases diversity

Who gets to interview at elite firms?

- Our setting
  - Professional services hiring at a Fortune 500 firm
  - 95% of applicants rejected just on the basis of their CV
  - We design algorithms to choose first round interviews
  - Goal: maximize the share of interviewed applicants who receive and accept an offer

- What we do
  - Build and train an “Upper Confidence Bound” (UCB) contextual bandit algorithm
  - Generate interview decisions on administrative test data
  - Compare with a) actual human decisions and b) traditional supervised learning algorithm
  - Focus on hiring rates (“quality”) and demographics (“diversity”) of selected candidates

Firms' recruiting practices are inside the Pareto frontier

- UCB algorithm selects higher quality and more diverse candidates
  - Increases hiring yield from 10% (human recruiters) to 25%
  - Doubles share of interviewed candidates who are Black or Hispanic (from 10% to 23%)

- But algorithmic design matters
  - Traditional supervised learning also improves hiring rates
  - But cuts the share of interview slots for Black and Hispanics to under 5%

- But what about?
  - Selective labels problem – we address using matching on covariates and an IV approach
  - Other measures of applicant quality – while limited data, no evidence that candidates humans select are better on other dimensions

There's more in the paper!

- Learning about (simulated) changes in applicant quality
  - UCB more effective at detecting changes in the quality of minority groups
- Blinding of race/gender variables
  - We can achieve similar results without using explicit demographic variables
- More information
  - Longer talk in Labor Studies: Monday, July 20th, 1pm
  - Questions? lraymond@mit.edu