AI in HR Management: Challenges and a path forward

CCC Workshop on Fairness and Economics, Harvard University

Prasanna (Sonny) Tambe, Wharton School, U. Pennsylvania

Practical issues that HR poses for AI

Portions of these slides are from a working paper:

AI in human resources management: Challenges and a path forward (Cappelli, Tambe, and Yakubovich 2018)

Available at:
`https://ssrn.com/abstract=3263878`

Informed by discussions during two workshops on digitization and analytics conducted with senior HR practitioners (in 2018 and 2019)

Prediction tasks in HR

- Recruiting – identifying possible candidates and persuading them to apply: are we securing good candidates?
- Selection – choosing which candidate should receive job offers: Are we offering jobs to those who will be the best employees?
- On-boarding – the initial process of bringing an employee into an organization, which includes a large number of administrative tasks
- Training – what should we recommend for you? Do our interventions improve performance?
- Performance management – can we identify good and bad performance: Do our practices improve job performance?
- Advancement – who gets promoted: Can we predict who will perform best in new roles? Can we make recommendations for your career?
- Retention – can we predict who is likely to leave and manage the level of retention?
- Employee benefits – Can we identify which benefits matter most to employees and what the effects of those benefits are (e.g., do they improve recruiting and retention)?

Contrast HR with marketing or finance

When compared with marketing or finance:

- The questions in marketing and finance are relatively clear
- Outcomes are easily measurable (e.g. sales or clicks)
- User activity generates very large data sets (number of clicks, number of purchases)
- In terms of social norms, it is generally accepted that companies are going to try to sell more of their product or to make more money

Example 1: HR norms and legal frameworks

Source: *NY Times*, Sep 18, 2018

Example 2: HR norms and legal frameworks

- Supplying ad engines with HR input data (job ads) leads to instant EEOC violations (Lambrecht and Tucker 2018)

FIVE challenges HR poses for machine learning

Focus is on the biggest problems that arise when transporting a prediction model to an HR context:

$$y = X\beta + \epsilon$$

1. Defining the dependent variable

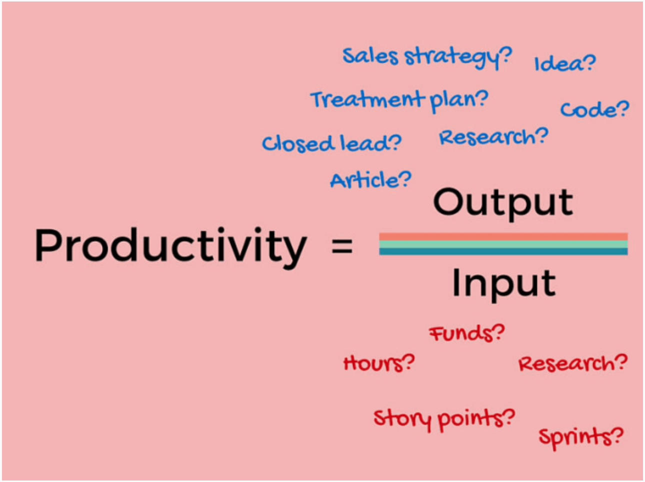

Very little agreement on what makes a good employee

What makes a **good** employee?

- Many dimensions to this construct, job performance is interdependent, and not everyone agrees what "good" means.
- Even if we know the answer, digital traces of what we would like to measure may not exist.
- Traditional measures, such as performance appraisal scores, have been **roundly criticized** in the literature (e.g. Coen and Jenkins 2002).

This has a paralyzing effect

Almost no HR practitioners we spoke with could define what an ideal measure *should* look like, let alone how a proxy measure could be generated from the data that is available.

2. Attribution and bias

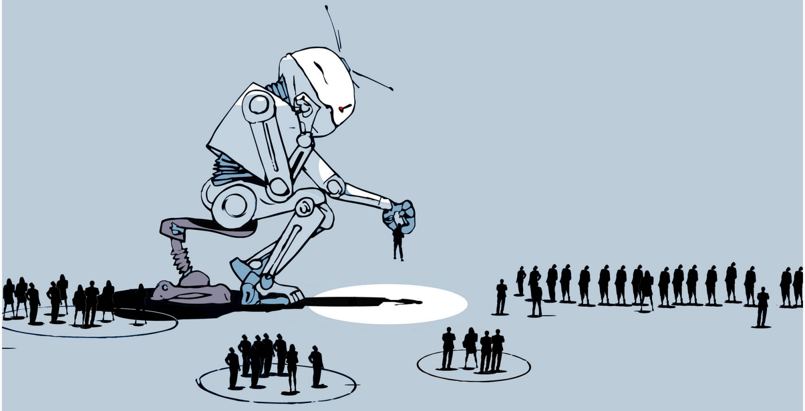

Image source: *The Wall Street Journal*

The economics of discrimination

- Economic attention to discrimination has a long history (Becker 1957; Phelps 1972; Arrow 1973)
  - Divided into 1) taste-based and 2) statistical discrimination (which in turn can be first or second degree)

- With free entry, taste-based discrimination may be competed out of the market, and should not be observable in equilibrium. More attention has been paid to the latter.

The economics of discrimination

- Key hiring problem: Information asymmetries in labor markets. Workers know their abilities but employers do not.
- Statistical discrimination is an optimal response to this signal v. noise problem (Aigner and Cain 1977).
- Group variables are used as a **proxy** for unobservable or unknowable individual characteristics (i.e. productivity).

$$y_i = x_i + \overline{x} + \epsilon_i$$

Measurement error and bias in input variables

- Depending on the covariance structure of the predictor variables, measurement error in the skill variable propagates bias to aggregate (group) variables
- With noisy measures of productivity, there is attenuation of the individual signal and weight on the *group* indicator

Creates a tradeoff between efficiency and equity

- Expected productivity is equal for the marginal applicant, but more false positives for the disadvantaged groups.
- Economically efficient but inequitable.

How does this translate to EEOC guidelines?

- **Many** open legal questions here
- Closest existing guidance is UGESP, Uniform Guidelines on Employee Selection Procedures
- Selection tests (including scoring models) cannot cause adverse impact to protected groups
- Determined by the rate of selection of each group relative to the highest rate selection group (4/5 rule)

3. The primacy of the paper trail

Algorithms are biased but HR managers are also biased (affinity bias, confirmation effects, endowment effects, etc.).

- **Ban the box**: high-profile campaign to remove the box asking if applicants have a criminal record. Makes it easier for ex-offenders to get a job and makes them less likely to re-offend.
- Effect of throwing away this information:
  - Evidence of decreased employment opportunities for low-skill black and hispanic workers (Doleac and Hansen, 2016; Agan and Starr 2017).
  - In the absence of the box, employers **infer this information** from demographic variables.

3. The primacy of the paper trail

Key point: For most civil action suits, the burden of proof lies with the plaintiff. In cases regarding termination of an employee's contract, the burden of proof is reversed.

This makes explainability very important

- Plaintiffs' burden of production:
  - Belong to a protected class
  - Qualified for the job
  - Experienced an adverse employment outcome
  - Job went to someone not of the protected class
- A principal HR function is to meet the burden of production when faced with allegations of wrongful termination based on race, gender, age.
- Contrast this with the explainability and fairness challenges that arise in an algorithmic context

This makes the paper trail very important

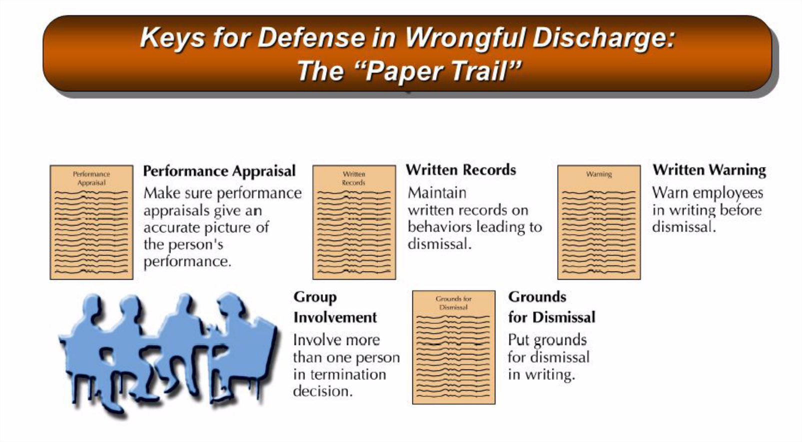

Source: *Human Resource Management*, Tenth Edition, Mathis and Jackson

4. Implementation: Employees vs. algorithms

- Employees may not react well to algorithmic decisions, especially if the news is bad.

- Like with any system, employees may adjust their behavior once they learn the incentives, making data inauthentic.

- A great deal of coordination within the firm is based on "relational" contracts.
  - "You let me leave early yesterday, so I don't mind staying a little later today ..."

The era of scientific management was not known for worker happiness

5. The biggest problem: Assembling the data

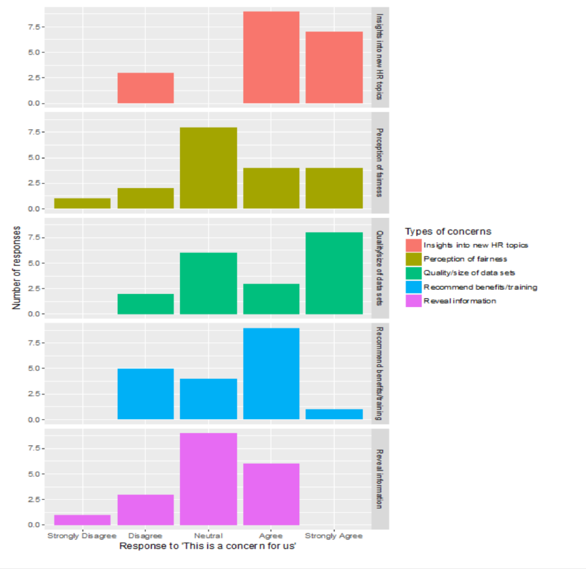

It is hard to get HR data at scale

- There are few observations per worker and algorithms can perform poorly when predicting rare outcomes.
  - Firms often use different HR vendors for different tasks, and it can be hard to merge data across sources.
  - Unlike accounting, there is no "standard" list of HR variables.

HRIS systems are very challenging to work with

HRIS	HCM	HRMS
Recruiting / ATS	HRIS	HCM
Core HR	Onboarding	Payroll
Benefit Admin / OE	Performance	Time & Labor
Absence Management	Position Control	
Compensation	Succession	
Training	Salary Planning	
Workflow	Global	
Self-Service	Analytics	
Reporting		

- HR information systems are data islands and this poses a tremendous challenge. Difficult to integrate systems when project goals and even outcome variables are poorly defined.
- Further compounded by compliance differences across borders.
- **Combined with the lack of a robust DV, this makes exploration and learning “costly”.**

Vendors can help

- Vendors have the ability to aggregate data from many employers to generate superior performance, which can alleviate the small data issue (and the sample bias issue).

- But challenges remain:
  - For each employer, a question that often arises is "how distinct is our context"?
  - This raises questions related to prediction accuracy and there are legal (privacy) issues as well.

The role of the GDPR

- The **GDPR** does not prohibit machine learning, but imposes a heavy burden on compliance
  - Article 22: Must opt-in to automated decision making.
  - Article 17: Right to be forgotten.
  - Article 13: Has a right to a meaningful explanation of logic.

- In sum: The guard rails for HR data usage that exist even in the most progressive companies are formidable and becoming increasingly difficult to navigate.

Potential paths forward

Causality	Legally defensible
	Explainable
Randomization	Easily understood heuristic
	Can help to debias algorithms
	Often perceived as fair
Employee contribution	Appeal process?
	AI councils?

Key battles for the deployment of AI in HR

- Challenges from the left and right in the use of ML for HR applications
  - defining the dv
  - attribution and bias
  - explainability vs. the paper trail
  - employees vs. algorithmic decision-makers
  - assembling the data

Thank you. Comments and questions are welcome at tambe@wharton.upenn.edu.