

# AI AND THE FUTURE OF WORK

---

PETER CAPPELLI

GEORGE W. TAYLOR PROFESSOR OF MANAGEMENT

DIRECTOR – CENTER FOR HUMAN RESOURCES

THE WHARTON SCHOOL

# SOME DEFINITIONS

---

- What is AI
  - The kind of intelligence that only humans have – it's a moving target
- We have data science – what is that?
  - Statistics starts with hypotheses to test
  - Data Science starts with often messy data – what can we learn from it
- The big tool is machine learning –
  - Finds patterns in data and adapts to them to create a model that predicts or optimizes on some goal
- It requires a lot of data, more than most employers have....
  - Many observations – thousands – and the more attributes, the better

# WHAT IS IT WE ARE TRYING TO DO? *MAINLY MAKE BETTER PREDICTIONS....*

---

- 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)?

# INTO THIS COMES ALGORITHMS

---

- Old approach: Testing hypotheses about what predicts a good hire
  - Got answers for each technique – interviews predict this much, IQ predicts that much, etc.
  - That is statistics
- New approach: Building one model to fit the data –
  - Gather everything you know about past hires, put it together with a machine learning software tool and tinker until the software finds the model that predicts as much as possible about who is good.
  - One score/one predictor: how close does each candidate/option look to past successes?
  - Not clear which measures in the algorithm are doing what



# MACHINE LEARNING ALGORITHMS WILL PREDICT BETTER THAN ANYTHING WE HAVE

---

- In part because it takes away human “judgment” / bias
  - Not “seeing” gender, race, appearance, etc. Also consistent weighting
  - We could do that without machine learning, too
- It needs one overall measure of the outcome
  - If that measure is not good, the algorithm won’t work – Amazon’s experience
- Can we “explain” the results? “Why is her score higher?”
  - Can you explain it to a judge if your decisions have some unpleasant outcome?
- Can we deal with unfairness – “I’ve worked three weekends in a row”
- What if it predicts for the majority group but not for minority?



# SUPPOSE YOU ARE MAKING A PROMOTION DECISION\*

---

- You are the VP filling a director role...
  - One candidate is a manager in your group, someone you know and like
  - The other suggested by HR, also a manager in another group
- Both interview well and on paper seem comparable
  - Hiring algorithm score for your manager is 84% fit with attributes associated with good performance as a director; Score for the other is 92%
  - You are more comfortable with the internal candidate because you know them
  - Who will you pick?

\* Borrowed from Jeffrey Polzer HBS