Causal Inference Lecture 1

Kenneth (Kenny) Joseph

University at Buffalo
Department of Computer Science and Engineering
School of Engineering and Applied Sciences

Pass phrase: Cynthia Rudin

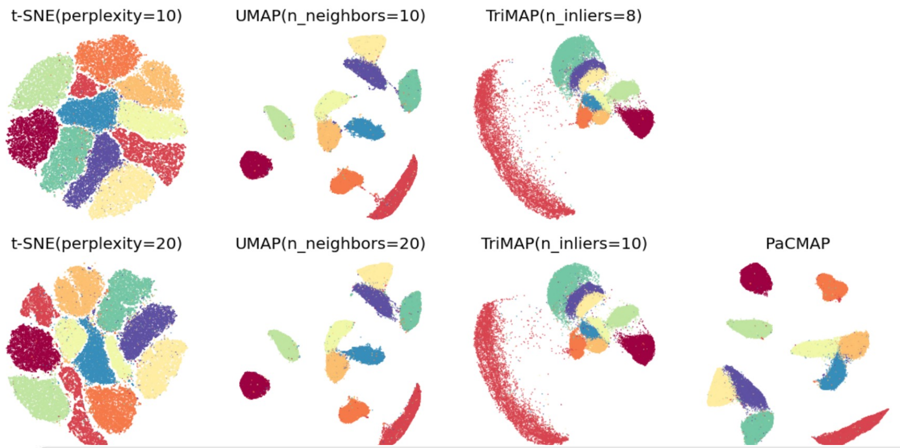

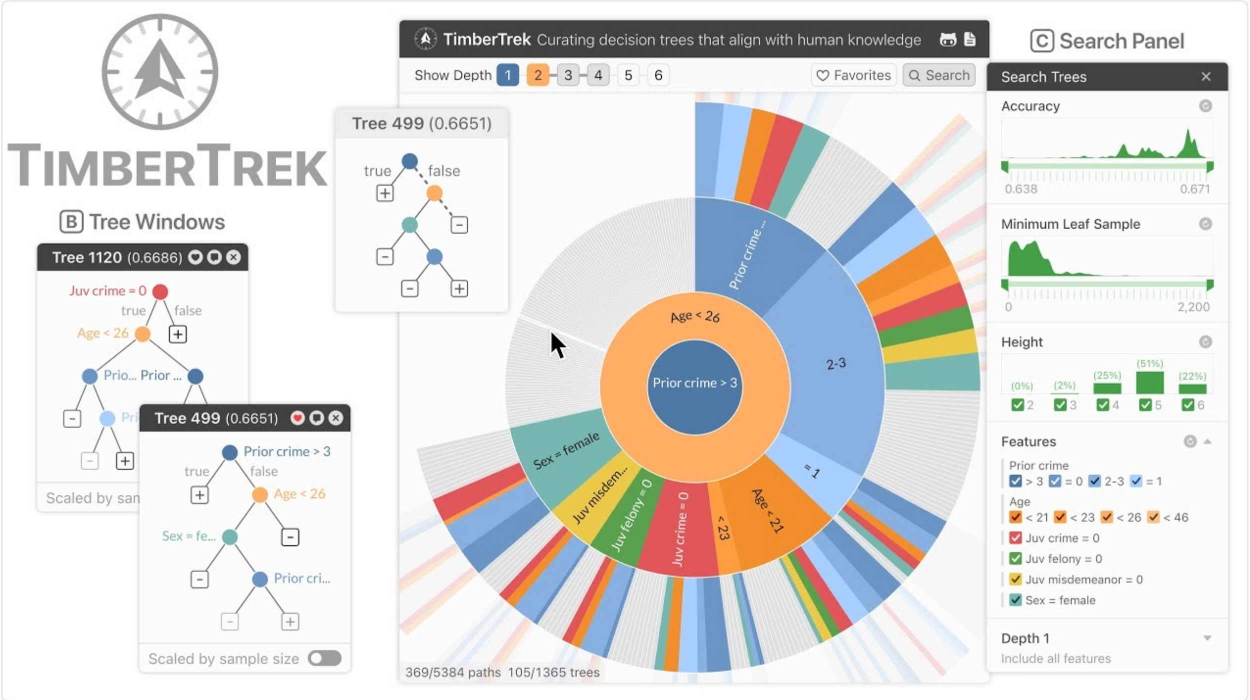

You have a *reading assignment* due Saturday

Discussion 1

1. *Listening Assignment:* How Race Was Made (Scene on Radio S2 E2)
2. *Watching Assignment:* The Root Of Algorithmic Bias And How To Deal With It

https://cse.buffalo.edu/rage-mlsoc/mlsoc_discussion/

Reading response: TQE format

For each in-class discussion (see the schedule for the dates), you will submit a summary of what you read. This response will be in the format of a TQE Response. Your submission should have three parts:

- **THOUGHTS** : What did this paper/video/podcast make you think about? What were the specific parts of the paper/video/podcast that made you think that? What were the main strengths/weaknesses of the paper/video/podcast? What did you like/dislike, and why?
- **QUESTIONS** What didn't you understand? What choices did the author make that you didn't understand/agree with? What were the aspects of the paper/video/podcast that you thought it got wrong?
- **EPIPHANIES** Does this paper/video/podcast help you think in a new way about a problem you're working on? Is there a part of the paper/video/podcast you found particularly confusing that you'd like help understanding? How does this paper/video/podcast link to some of the other papers/videos/podcasts we have discussed or other concepts you've learned in class?

- **THOUGHTS** : What did this paper/video/podcast make you think about? What were the specific parts of the paper/video/podcast that made you think that? What were the main strengths/weaknesses of the paper/video/podcast? What did you like/dislike, and why?
- **QUESTIONS** What didn't you understand? What choices did the author make that you didn't understand/agree with? What were the aspects of the paper/video/podcast that you thought it got wrong?
- **EPIPHANIES** Does this paper/video/podcast help you think in a new way about a problem you're working on? Is there a part of the paper/v
like help understanding? How does this p
papers/videos/podcasts we have discuss

e epiphany

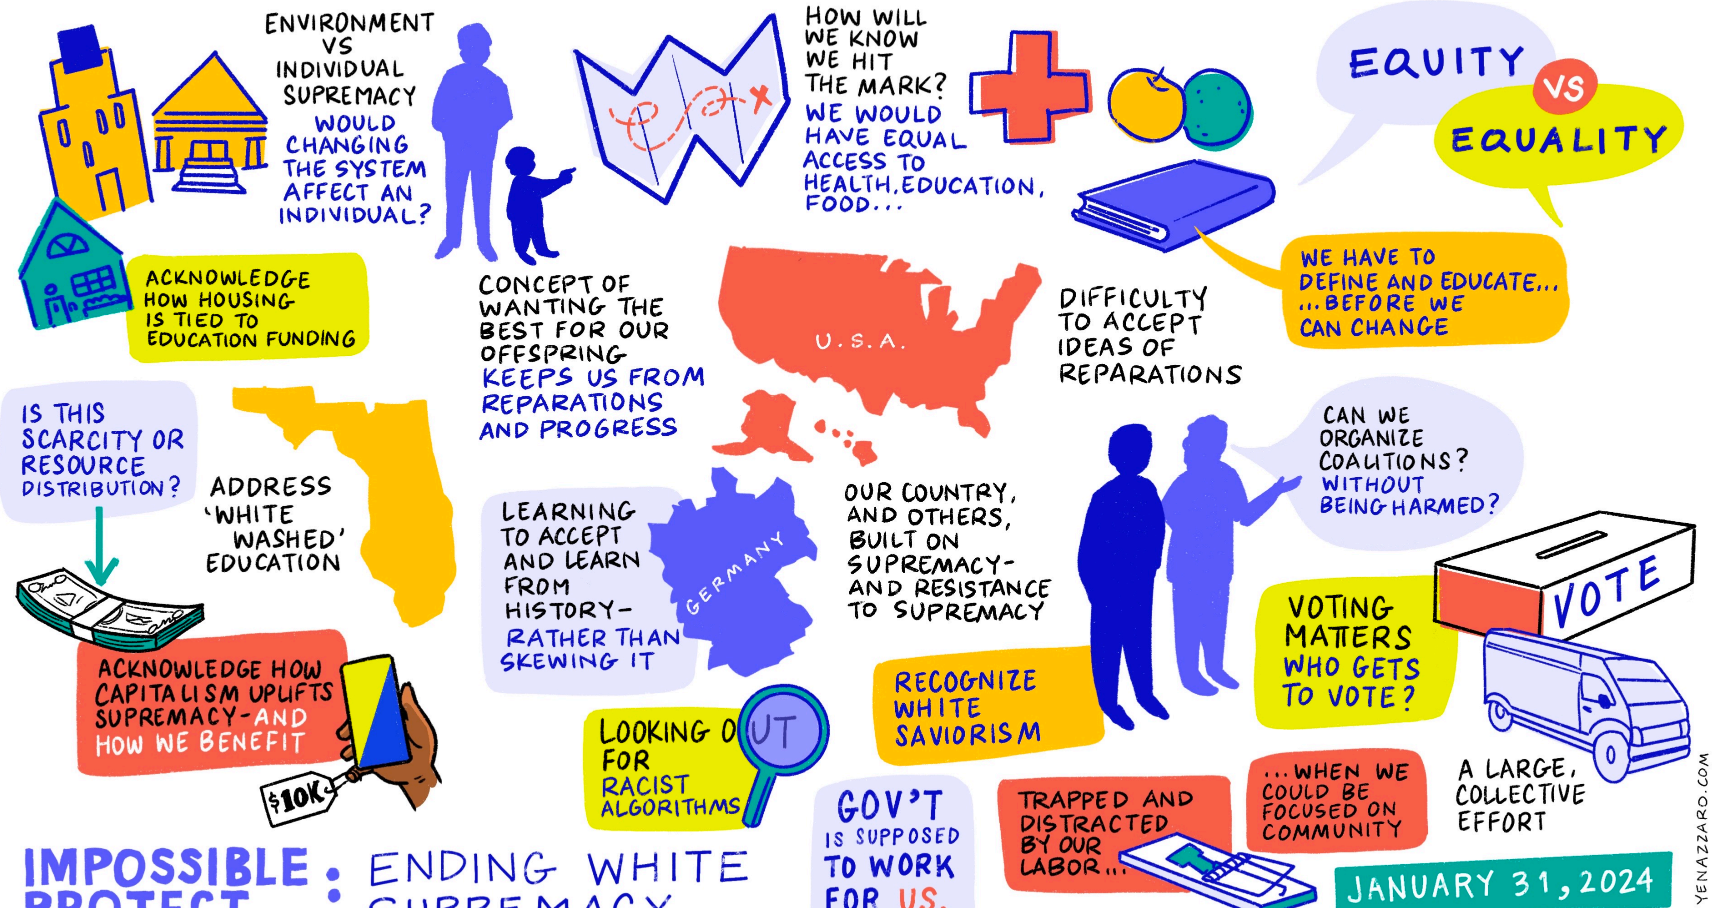

Kenny's (pedagogically-centered) TQE

- Thought:
  - I liked the discussion about limits to existing resources. This made me think about the fact that Atri and I have worked for several years on a project in the context of child welfare. It has become a useful site for us to think about resource allocation as a foil for debates about incremental vs. radical change
- Question:
  - I think it was good that tech played a limited role in the discussion. I am curious how you all felt about that. Was it humbling? Confusing? Right/wrong?
- Epiphany
  - I realized, and maybe this is obvious, that we should probably discuss the history of computing and AI to help us understand that advances in computing have come less from the goal of social progress, and more from the goal of centralizing power.

In two weeks you have a **group midpoint submission**

https://cse.buffalo.edu/rage-mlsoc/mlsoc_unitmid/

In three weeks you have a full group submission

Details later this week.

Units Overview

- **Unit 1**
  - **Tool:** Causal Modeling
  - **Domain:** Healthcare
- **Unit 2**
  - **Tool:** (CS) Theory
  - **Domain:** Policing and Criminal Justice
- **Unit 3**
  - **Tool:** Simulation
  - **Domain:** Mis/Disinfo

Questions/Comments?

Break!

What is Machine Learning?

Canonical Definition

"A computer program is said to **learn** from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E."

— Tom Mitchell, Professor at Carnegie Mellon University

Tom M Mitchell et al. "Machine learning. 1997". In: *Burr Ridge, IL: McGraw Hill* 45.37 (1997), pp. 870–877.

ML as a recipe creator

> *An algorithm is like a recipe. It takes "inputs" (the ingredients), performs a set of simple and (hopefully) well-defined steps, and then terminates after producing an "output" (the meal)*

A learning algorithm is a game of roulette on a 50 dimensional wheel that lands on a particular spot (a recipe) based completely on how it was trained, what examples it saw, and how long it took to search.

https://medium.com/@geomblog/when-an-algorithm-isn-t-2b9fe01b9bb5

ML as the production of intelligence from data

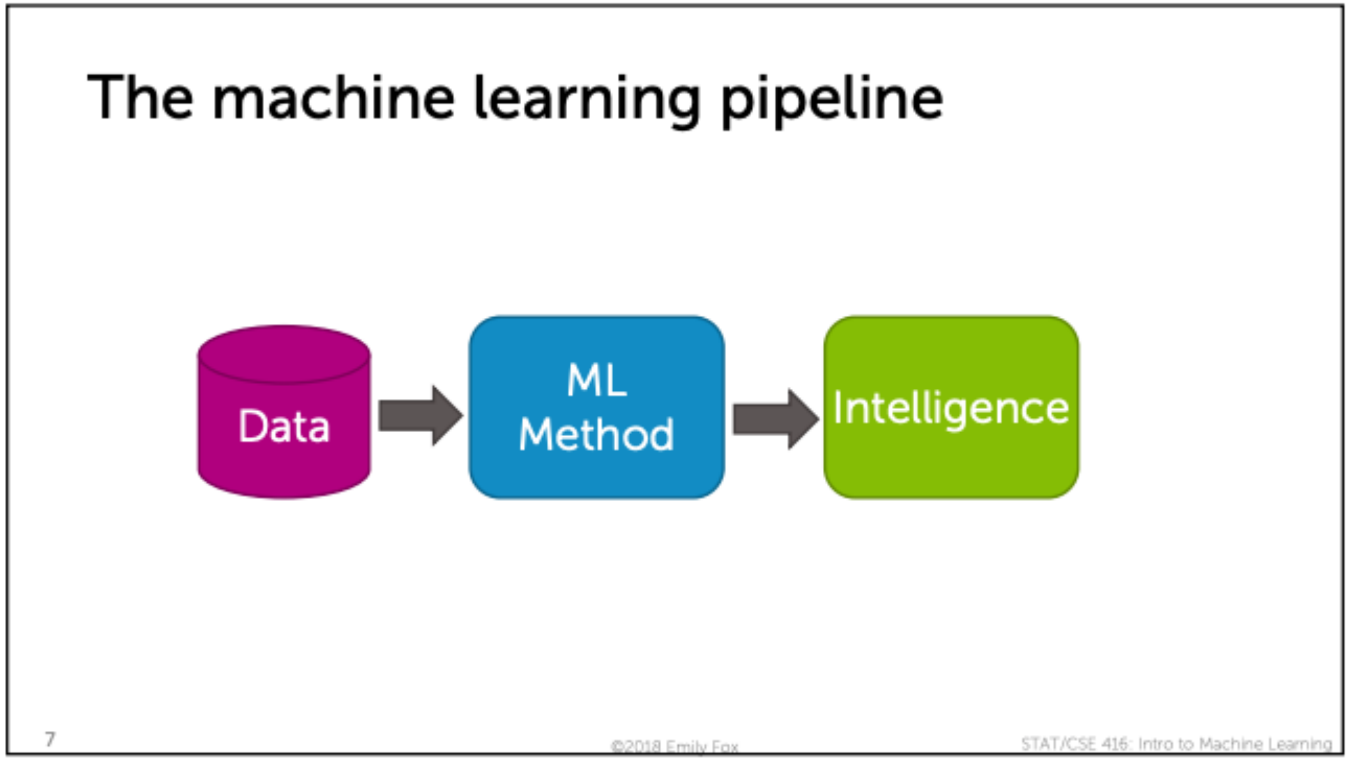

- I like this definition because it…
  - fits a number of different **learning paradigms**
  - is dead simple – we use ML to learn from data
- However, perhaps over-simplified… doesn't this match statistics too?

These images are taken from https://courses.cs.washington.edu/courses/cse416/18sp/lectures.html

ML as a composite of many things

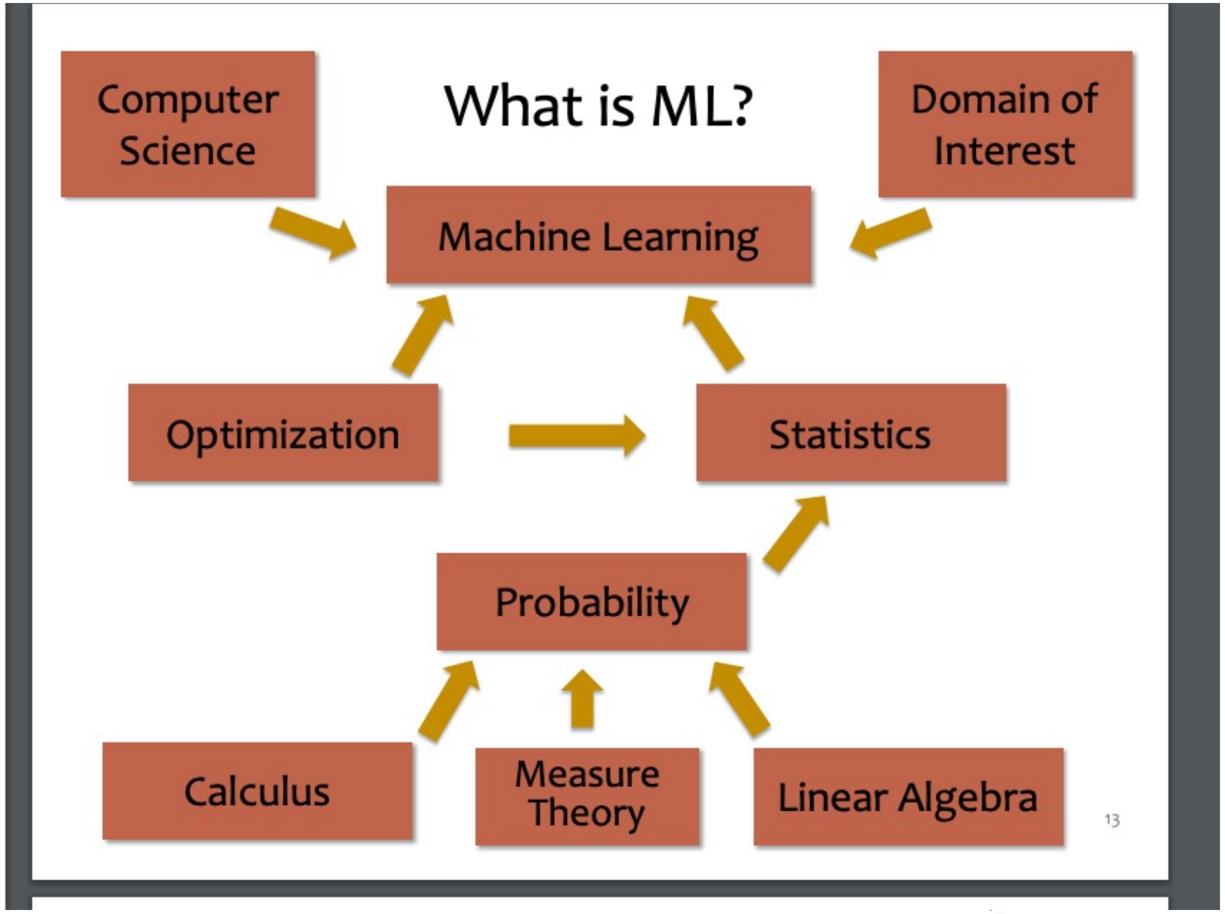

These images are taken from http://www.cs.cmu.edu/~mgormley/courses/10601/schedule.html

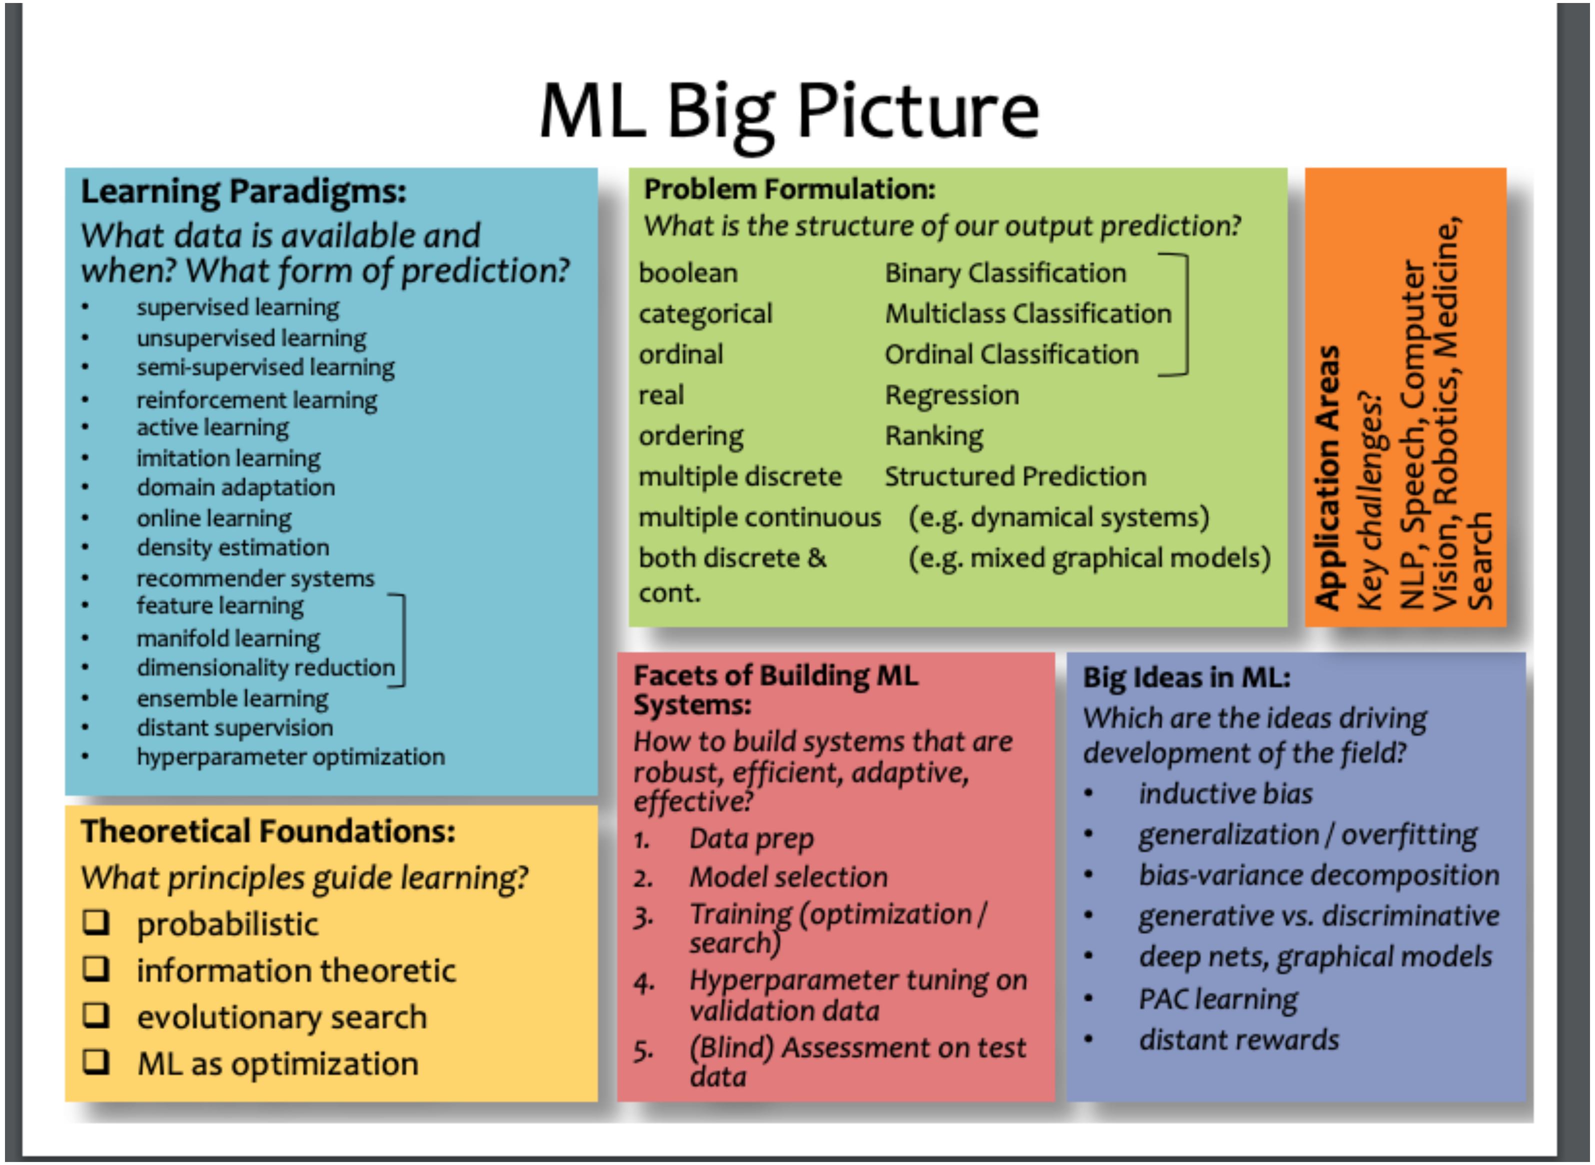

ML is many things to many people.

These images are taken from
http://www.cs.cmu.edu/~mgormley/courses/10601/schedule.html

ML as generalization of (training) data

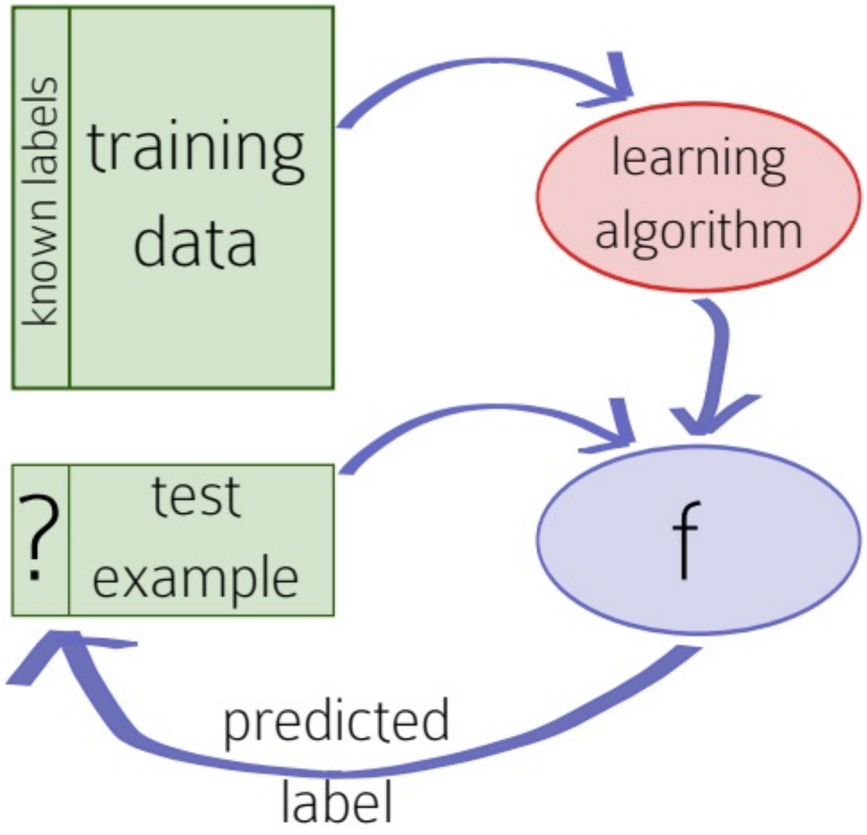

CIML, Chapter 1

ML as generalization of (training) data

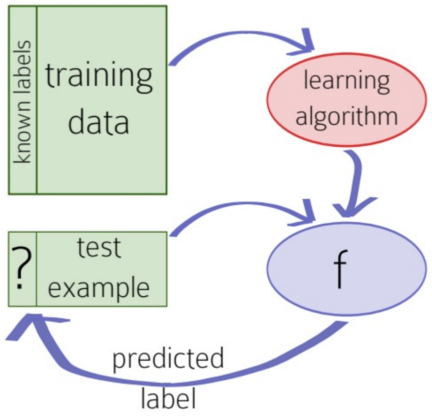

Figure 1.1: The general supervised approach to machine learning: a learning algorithm reads in training data and computes a learned function f. This function can then automatically label future text examples.

CIML, Chapter 1

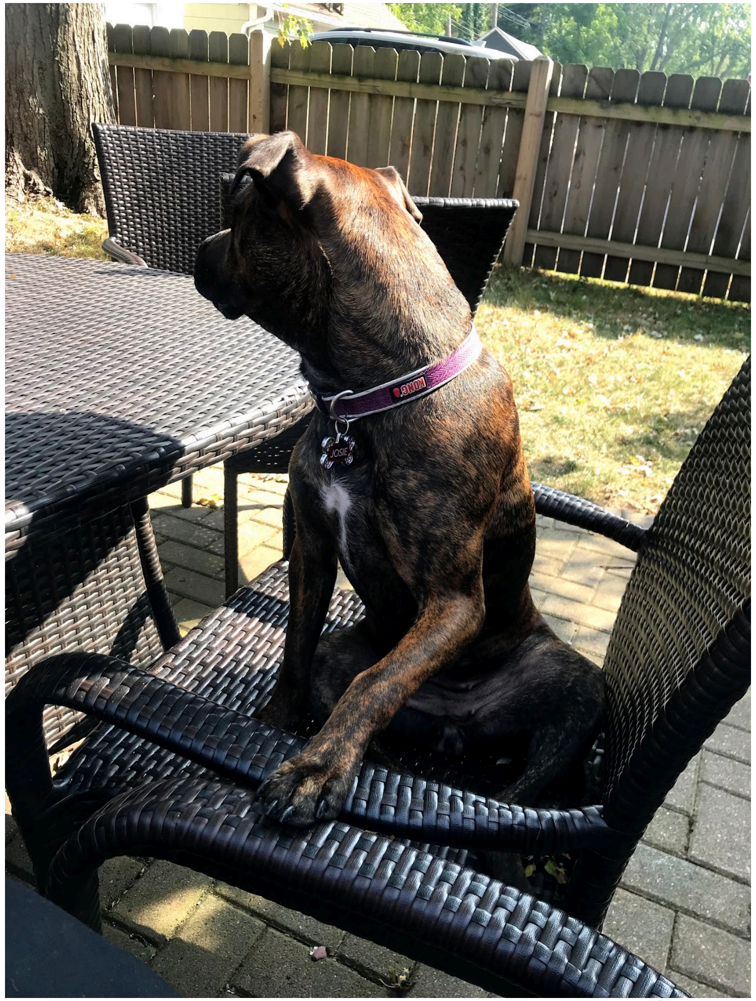

Josie looks like she understand chairs. She does not. She cannot **generalize** beyond her training data.

A boy saw 17 doctors over 3 years for chronic pain. ChatGPT found the diagnosis

Alex experienced pain that stopped him from playing with other children but doctors had no answers to why. His frustrated mom asked ChatGPT for help.

Why this course?

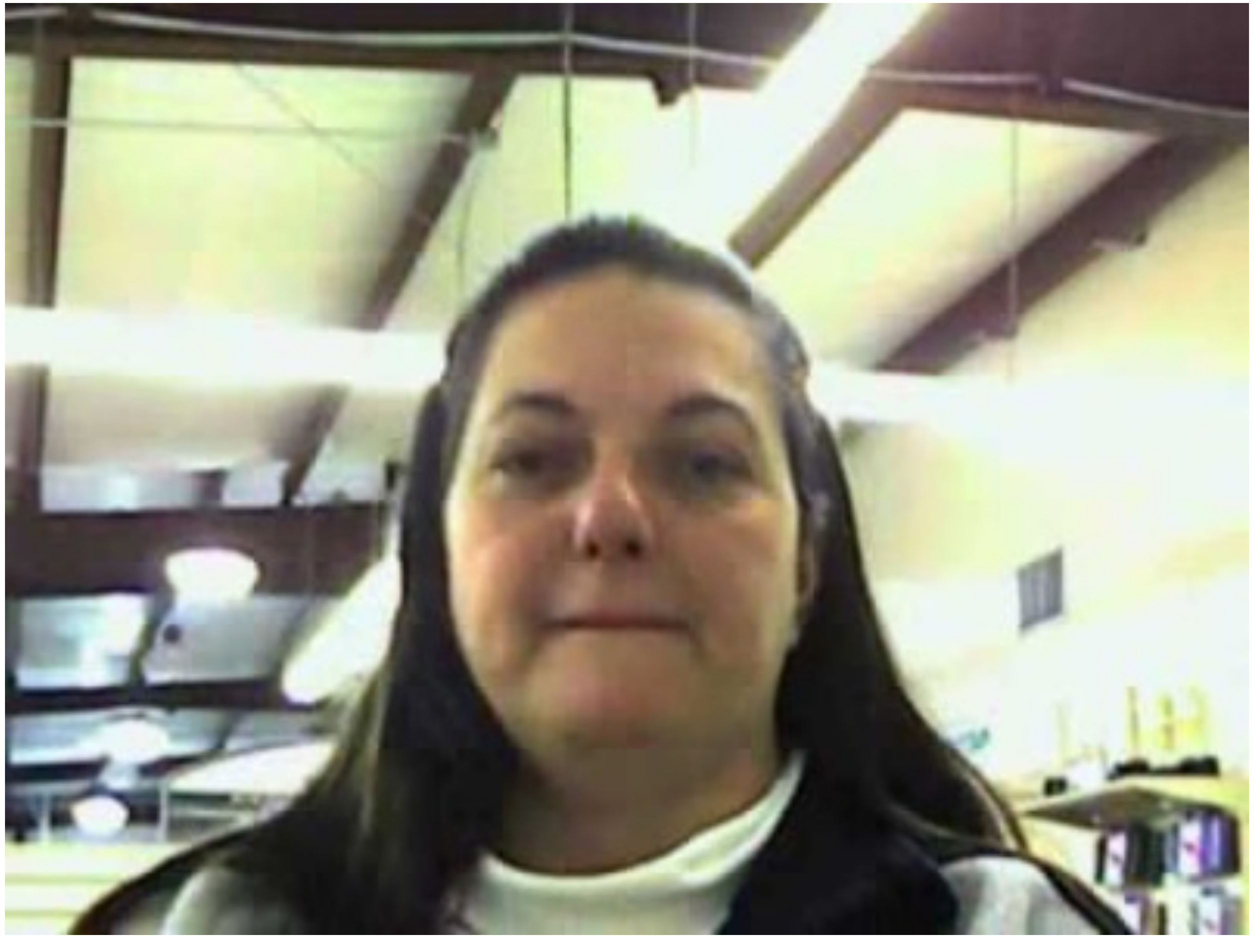

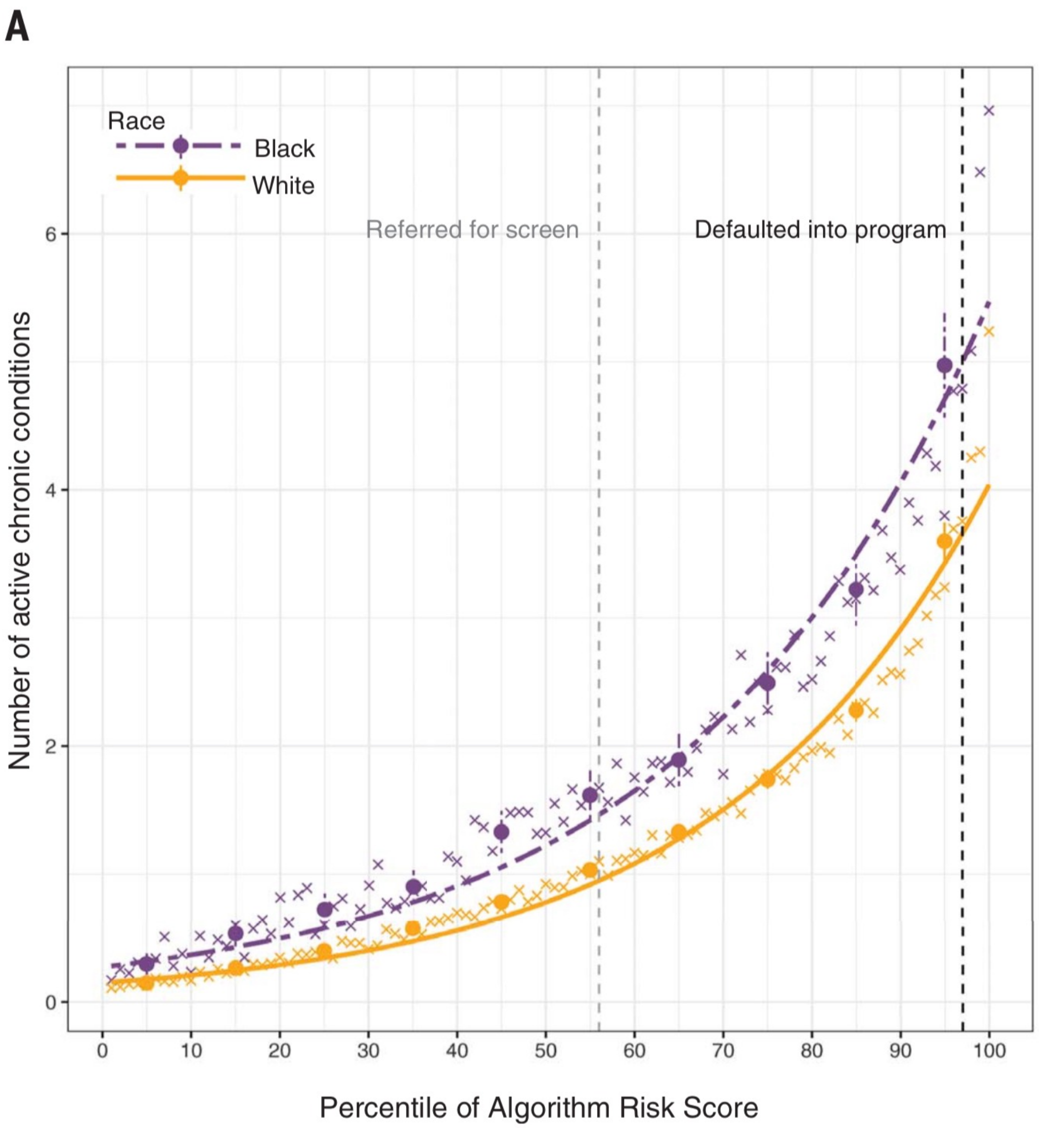

Fig. 1. Number of chronic illnesses versus algorithm-predicted risk, by race. (**A**) Mean number of chronic conditions by race, plotted against

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, *366*(6464), 447–453. https://doi.org/10.1126/science.aax2342

One Month, 500,000 Face Scans: How China Is Using A.I. to Profile a Minority

In a major ethical leap for the tech world, Chinese start-ups have built algorithms that the government uses to track members of a largely Muslim minority group.

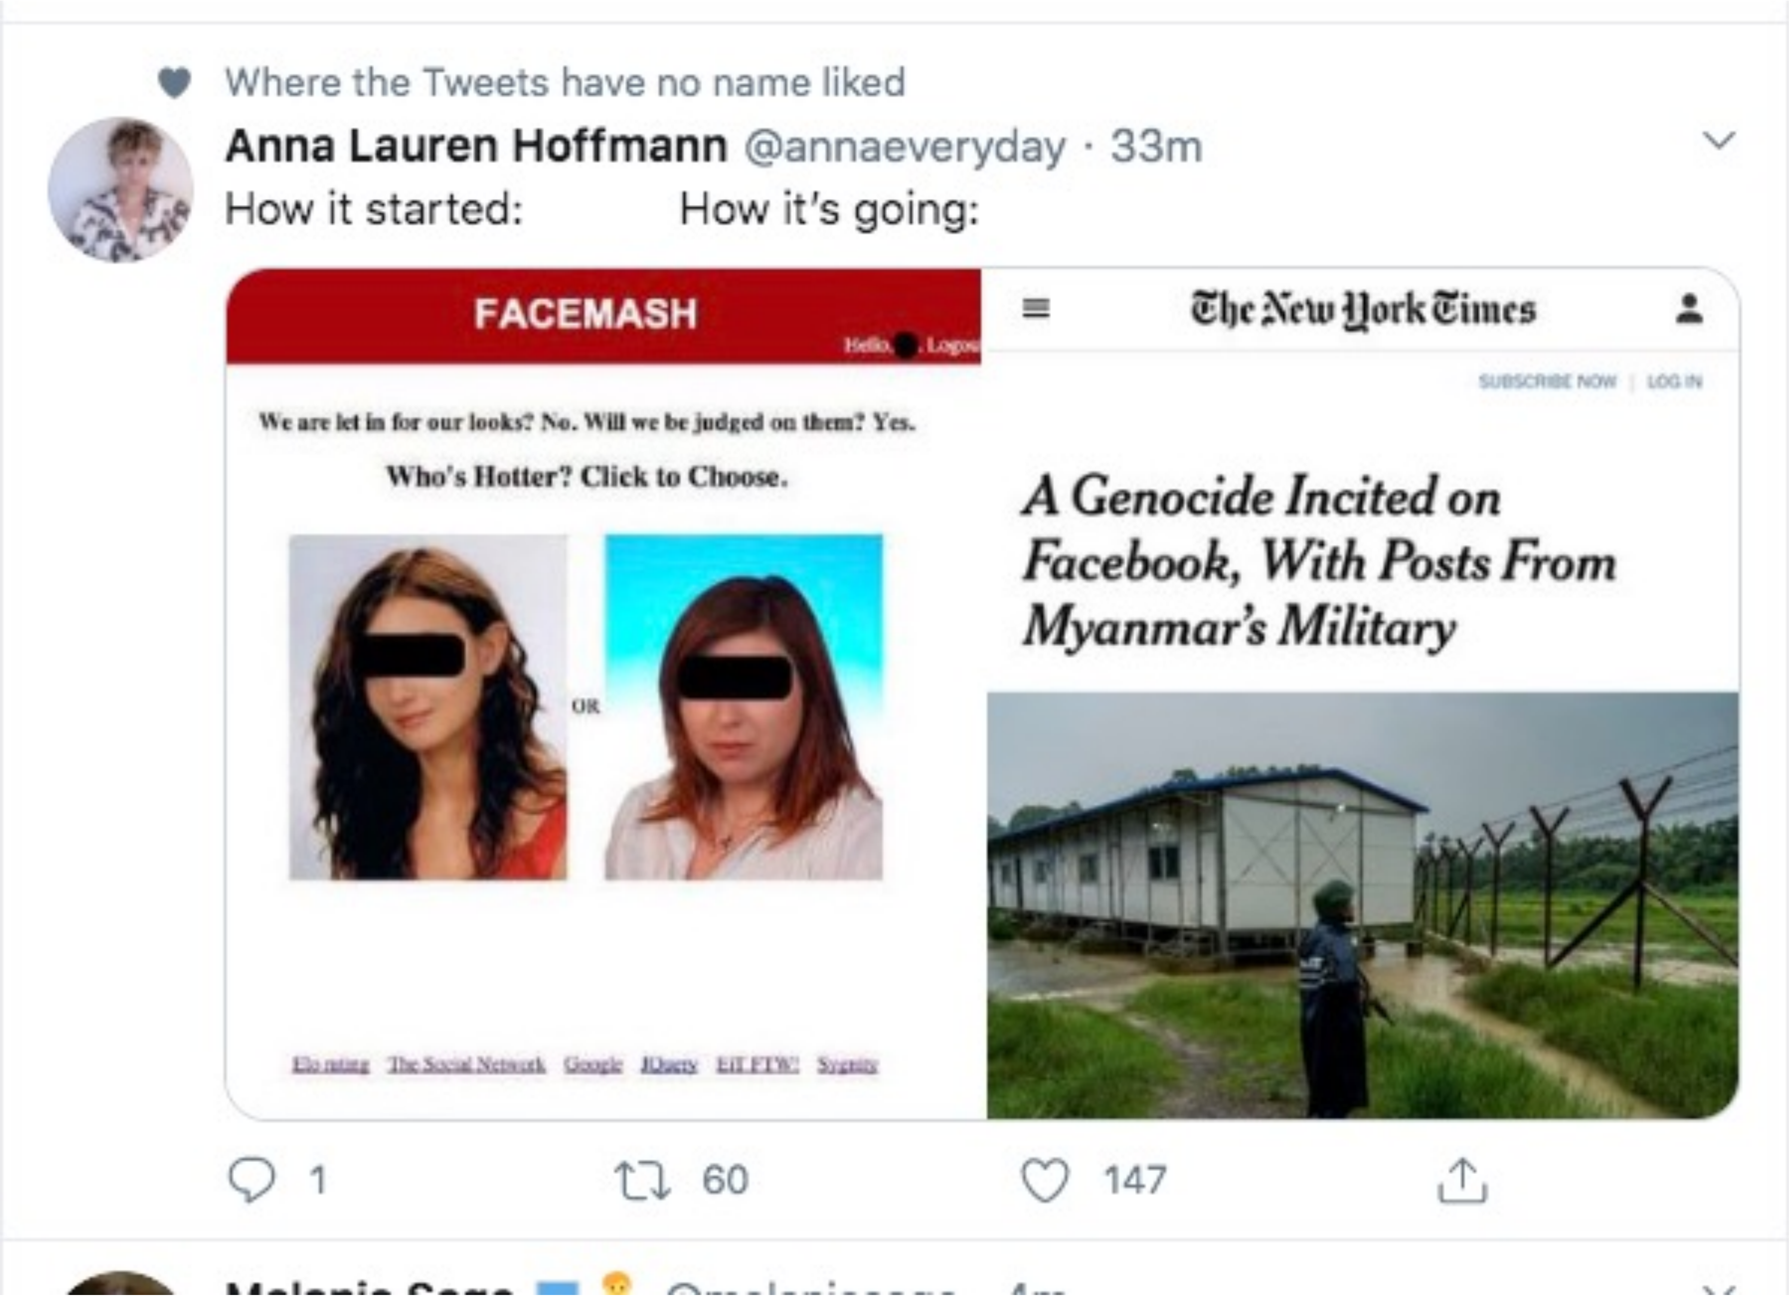
Where the Tweets have no name liked
Anna Lauren Hoffmann @annaeveryday · 33m
How it started:
How it's going:
FACEMASH
We are let in for our looks? No. Will we be judged on them? Yes.
Who's Hotter? Click to Choose.
OR
The New York Times
SUBSCRIBE NOW | LOG IN
A Genocide Incited on Facebook, With Posts From Myanmar's Military
1
60
147

MOTHERBOARD
TECH BY VICE

Flawed Algorithms Are Grading Millions of Students' Essays

Fooled by gibberish and highly susceptible to human bias, automated essay-scoring systems are being increasingly adopted, a Motherboard investigation has found

Research is scarce on the issue of machine scoring bias, partly due to the secrecy of the companies that create these systems. Test scoring vendors closely guard their algorithms, and states are wary of drawing attention to the fact that algorithms, not humans, are grading students' work. Only a

Meanwhile, it tended to underscore African Americans and, at various points, Arabic, Spanish, and Hindi speakers—even after attempts to reconfigure the system to fix the problem.

> **"The BABEL Generator proved you can have complete incoherence, meaning one sentence had nothing to do with another," and still receive a high mark from the algorithms.**

Slide by Kenny Joseph

Why?

Well, let's dig into one example