Societal Computing Week 1

Prof. Kenny Joseph

Who am I?

- Buffalo born (Buffalo sports fan)
- Undergrad CS + Chinese (Minor)
- Grad Societal Computing
- Free time:
  - See slide right
  - Also running/biking, listening to bad music
- My other day job – studying people using computers

Some reminders/notes

- You have assignments due at 11:59PM on Wednesday both this week and next week
- My office hours:
  - On Fridays, by appointment.
  - kjoseph@buffalo.edu
- I will discuss political topics. I will do so as objectively as possible
- I have discovered that my chair is very squeaky. Sorry.

What are we gonna do?!

- Lecture 1.1: This! Intro, Details, Motivation
- Lecture 1.2: What is societal computing?
- Rest of this week: AI/ML/DM through a Societal Computing Lens
- Next week: Social Media – the good, the bad, and TikTok

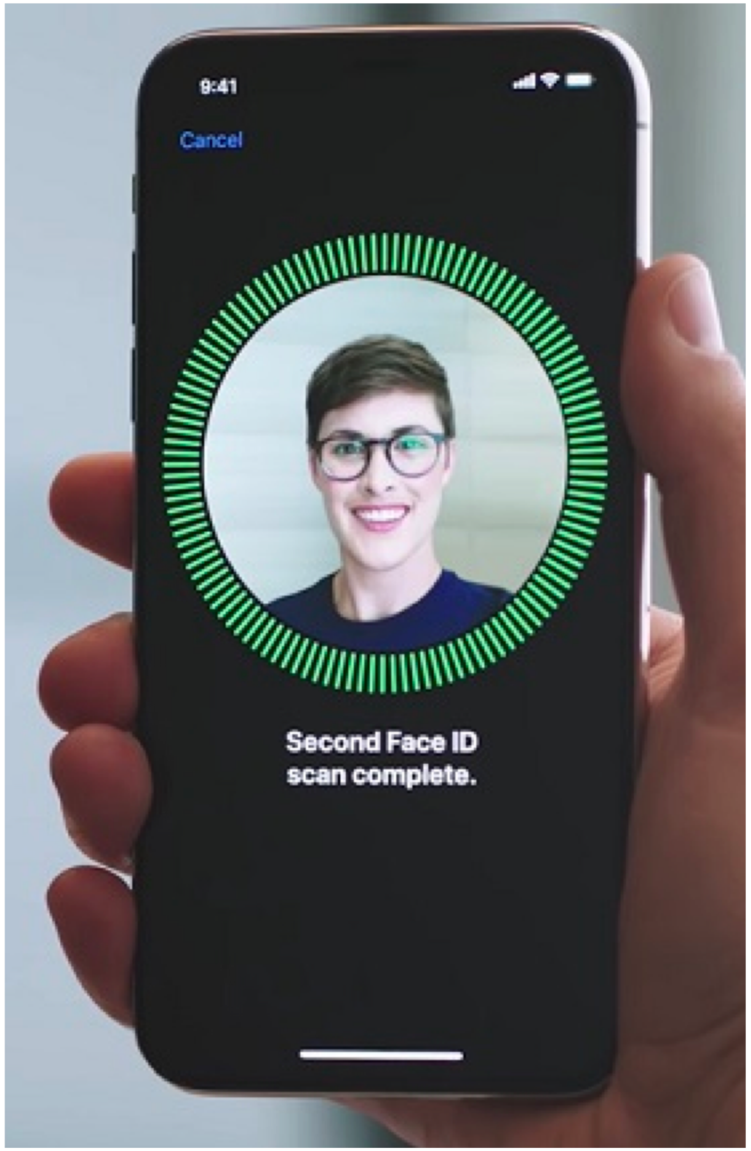

https://www.youtube.com/watch?v=t4DT3tQqgRM

Who does my technology serve? How do I know?

Rachel Thomas
@math_rachel

"Is this ML model doing good?" is the wrong question to be asking.

We need to ask: "how is this ML model shifting power?" @riakall #NeurIPS2019

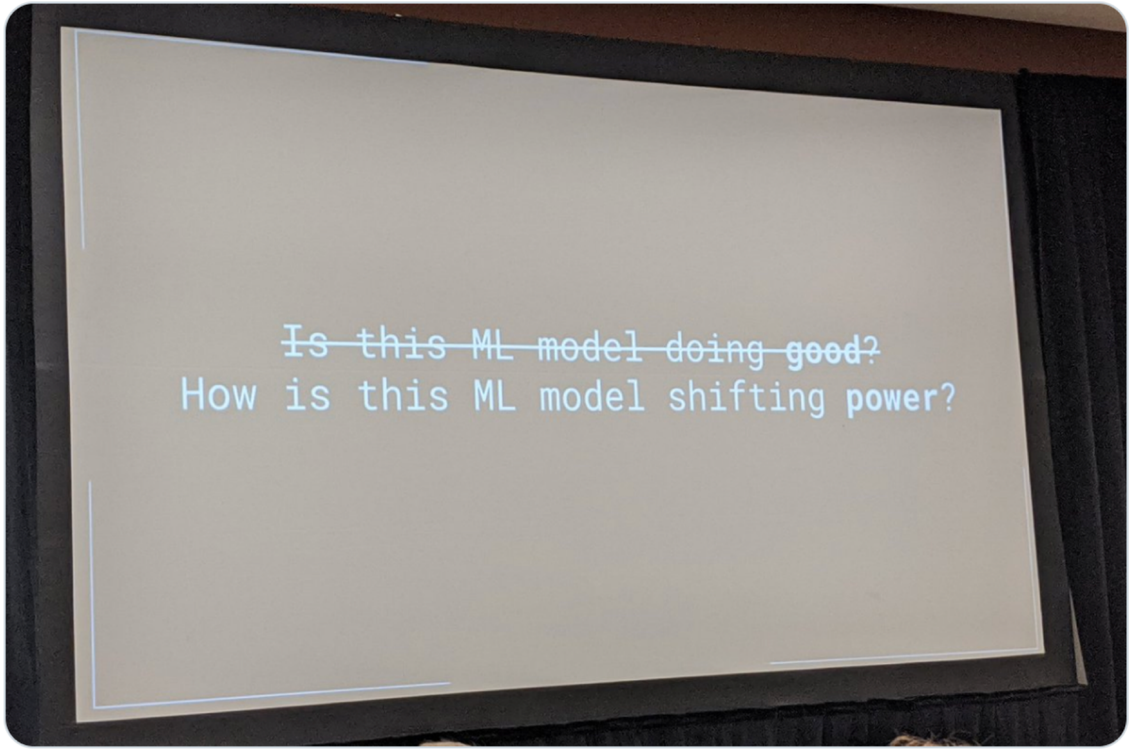

Computing **for**, **in**, and ideally, **with** and **by** society

Women in Statistics and Data Science
@WomenInStat

Sometimes I wonder about the reply guys.

I like to imagine that they wake up in the morning and say "Maybe I'll explain water to a fish today? Nah, too wet; I'll stick with explaining statistics to statisticians."

> **Women in Statistics and Data Science** @WomenInStat · Jul 25
> I have an 82 cent joke. If my male colleague told it he'd be paid a dollar.

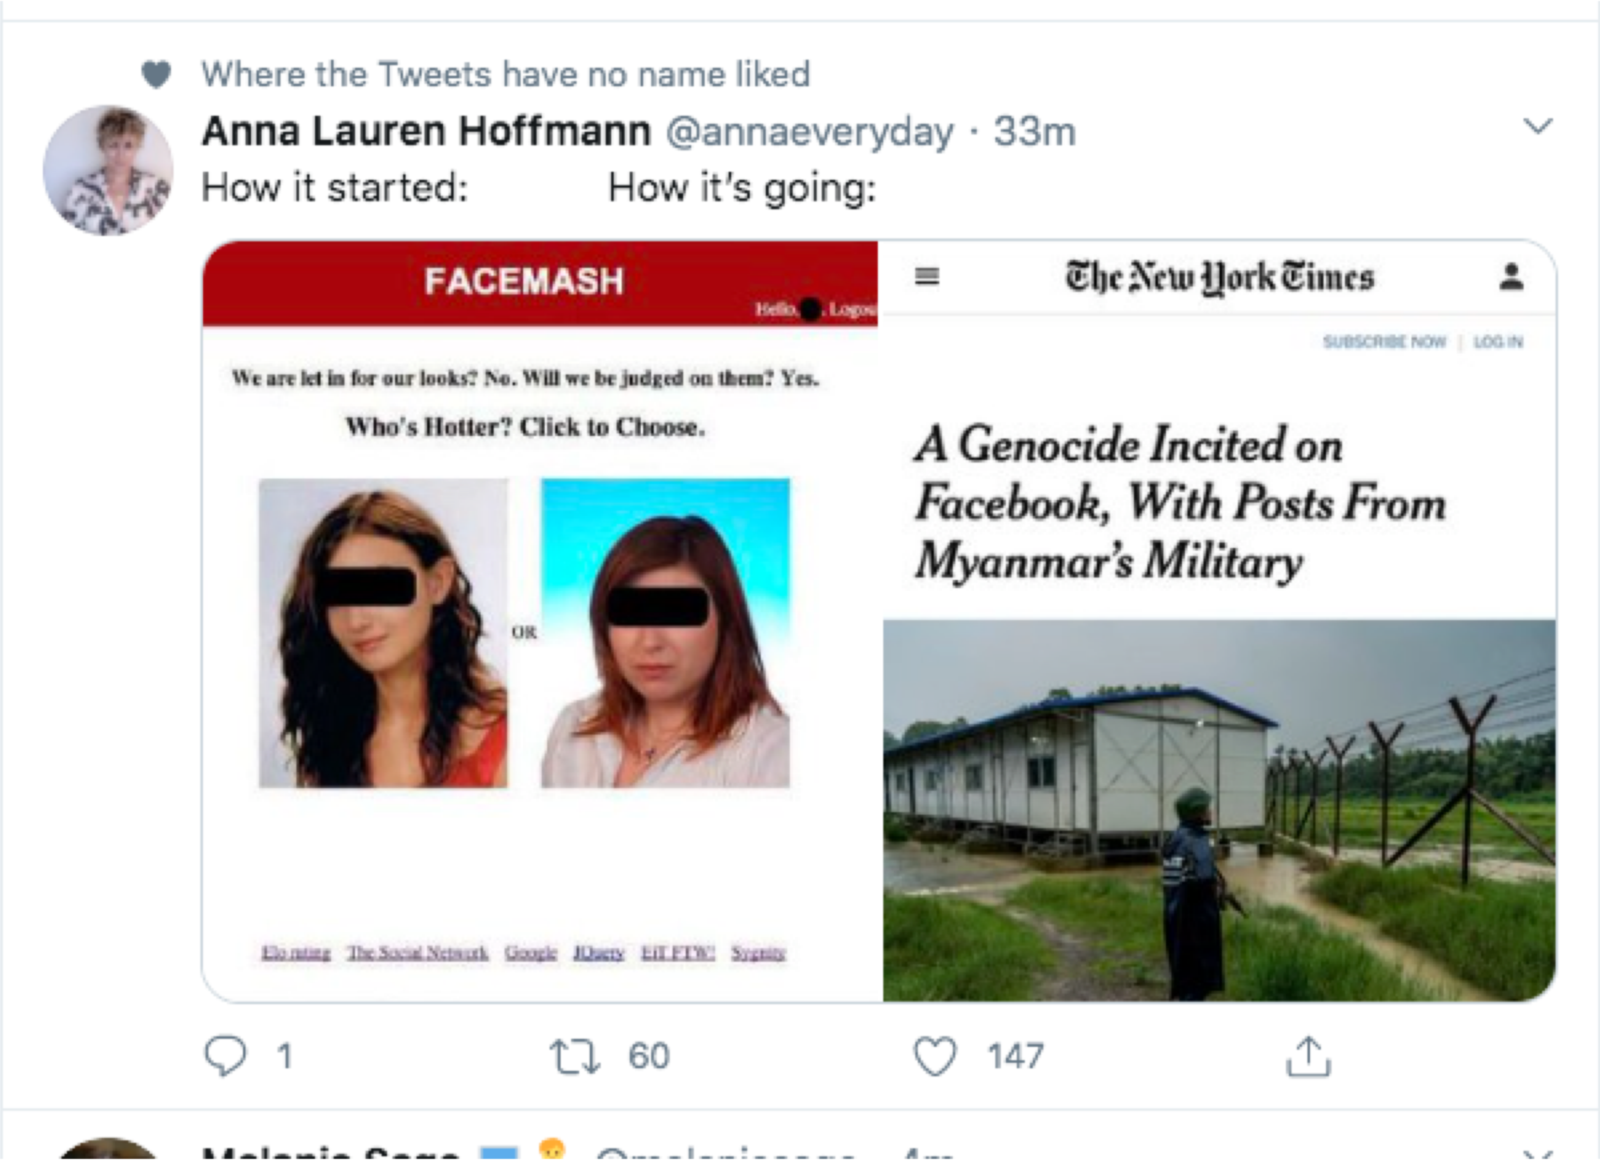

Where the Tweets have no name liked
Anna Lauren Hoffmann @annaeveryday · 33m
How it started:
How it's going:
FACEMASH
We are let in for our looks? No. Will we be judged on them? Yes.
Who's Hotter? Click to Choose.
OR
The New York Times
SUBSCRIBE NOW | LOG IN
A Genocide Incited on Facebook, With Posts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

What is societal computing?

Computing **for**, **in**, and ideally, **with** and **by** society

Computing **for** society

- Takes as a starting point that technology can solve certain societal problems
- Focus is on:
  - What problems can I identify in this world?
  - How do we build tools to solve these problems?

Computing **for** society

https://www.who.int/violence_injury_prevention/road_safety_status/2013/facts/en/

Computing **for** society

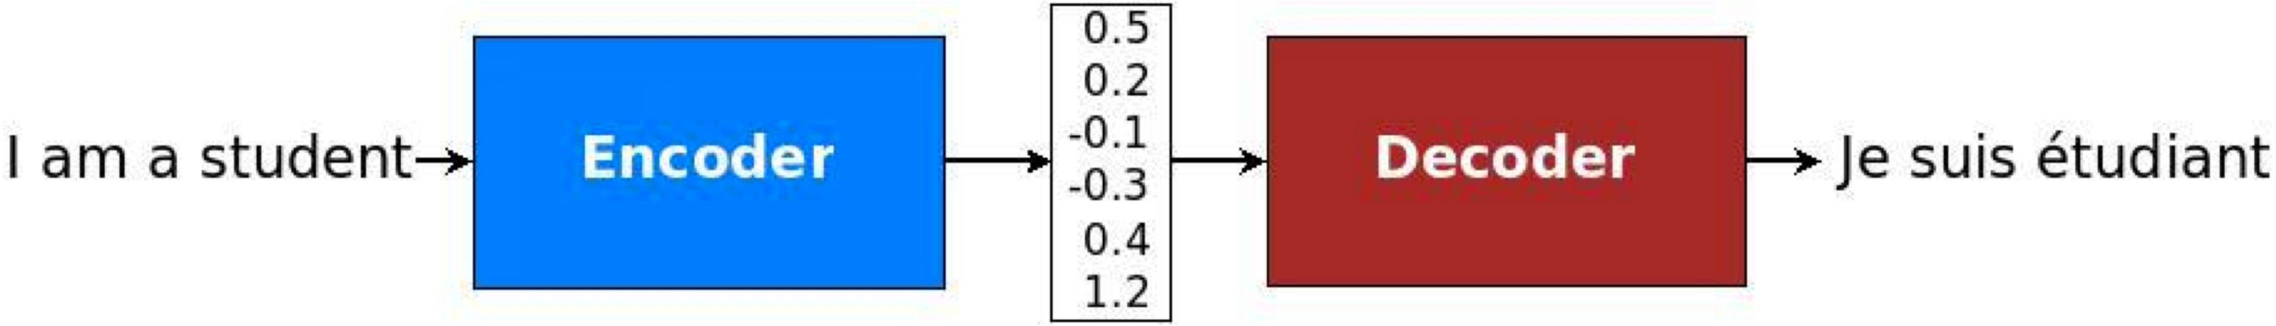

https://youtu.be/e0RLWNVxfb0?t=8

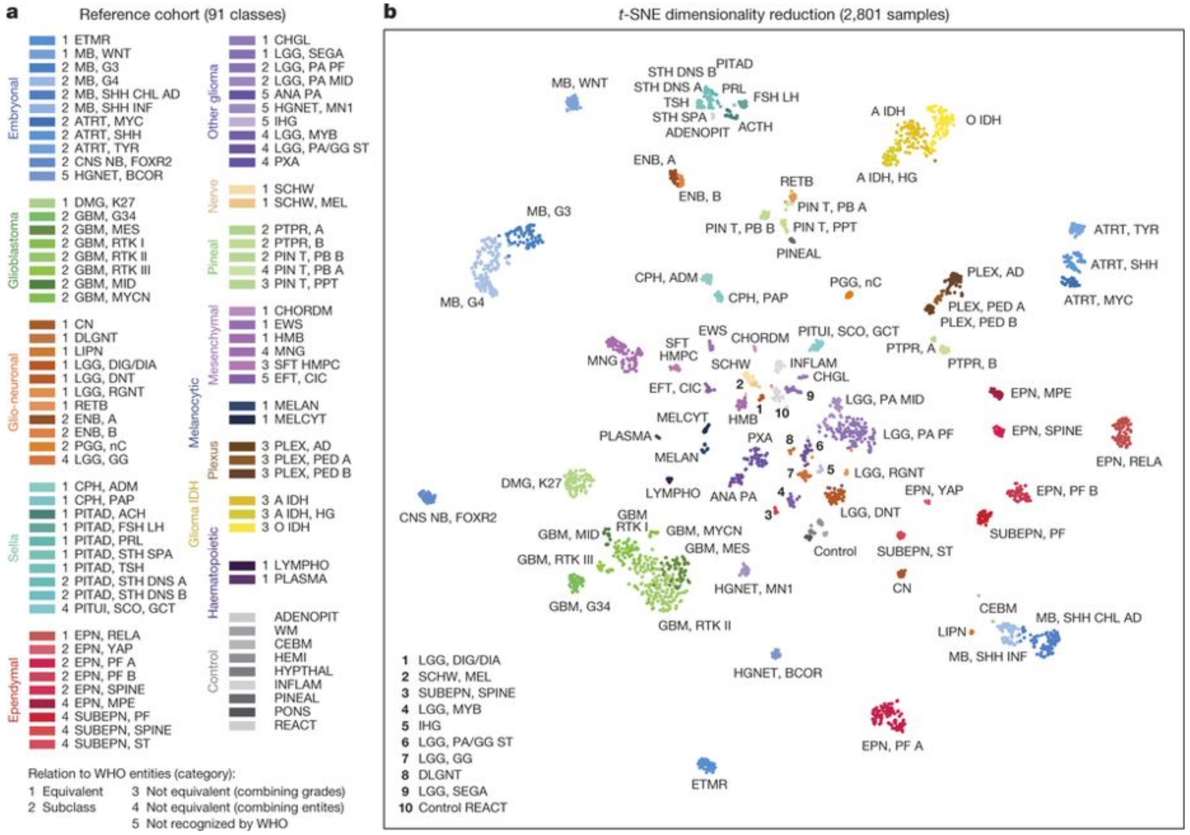

https://blogs.plos.org/speakingofmedicine/2018/11/28/better-medicine-through-machine-learning-whats-real-and-whats-artificial/

https://www.nature.com/articles/nature26000/figures/1

Computing in society

- Realizes that technology rarely, if ever, “just works”
- Focus is on:
  - Who was(n’t) this technology built for?
  - What problems might arise in the real world that the designers didn’t think about?
  - How can we identify these issues?

Computing **in** society

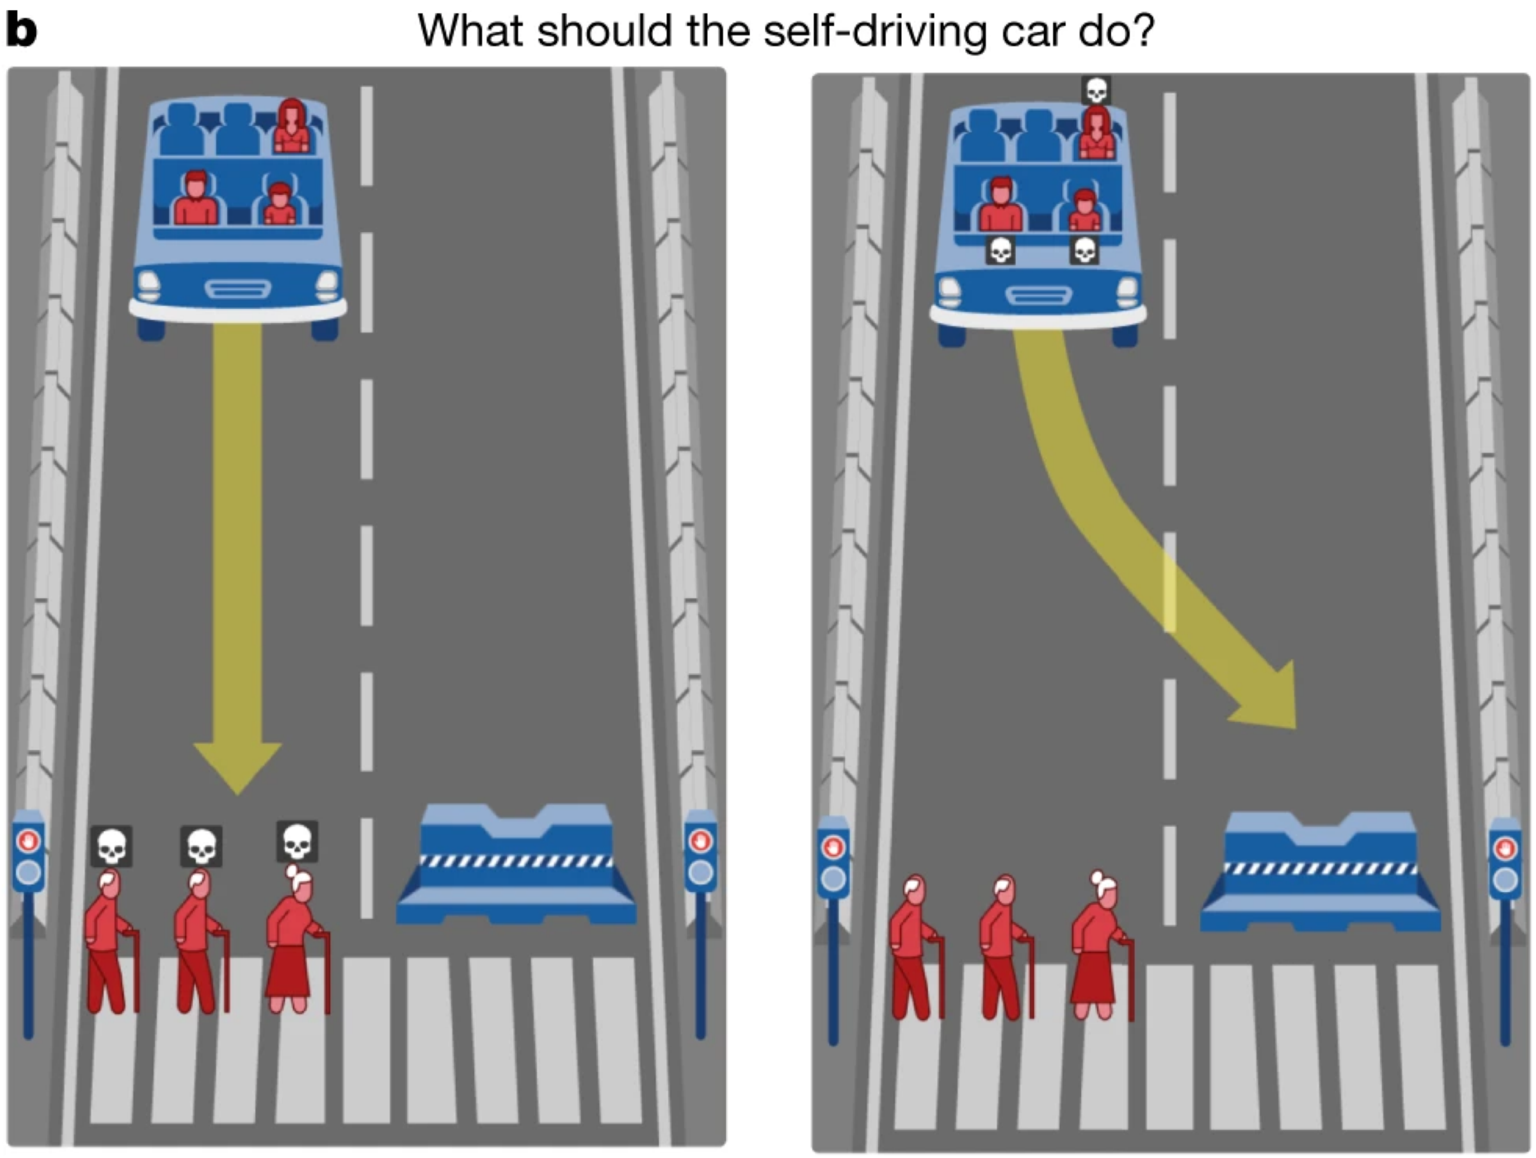

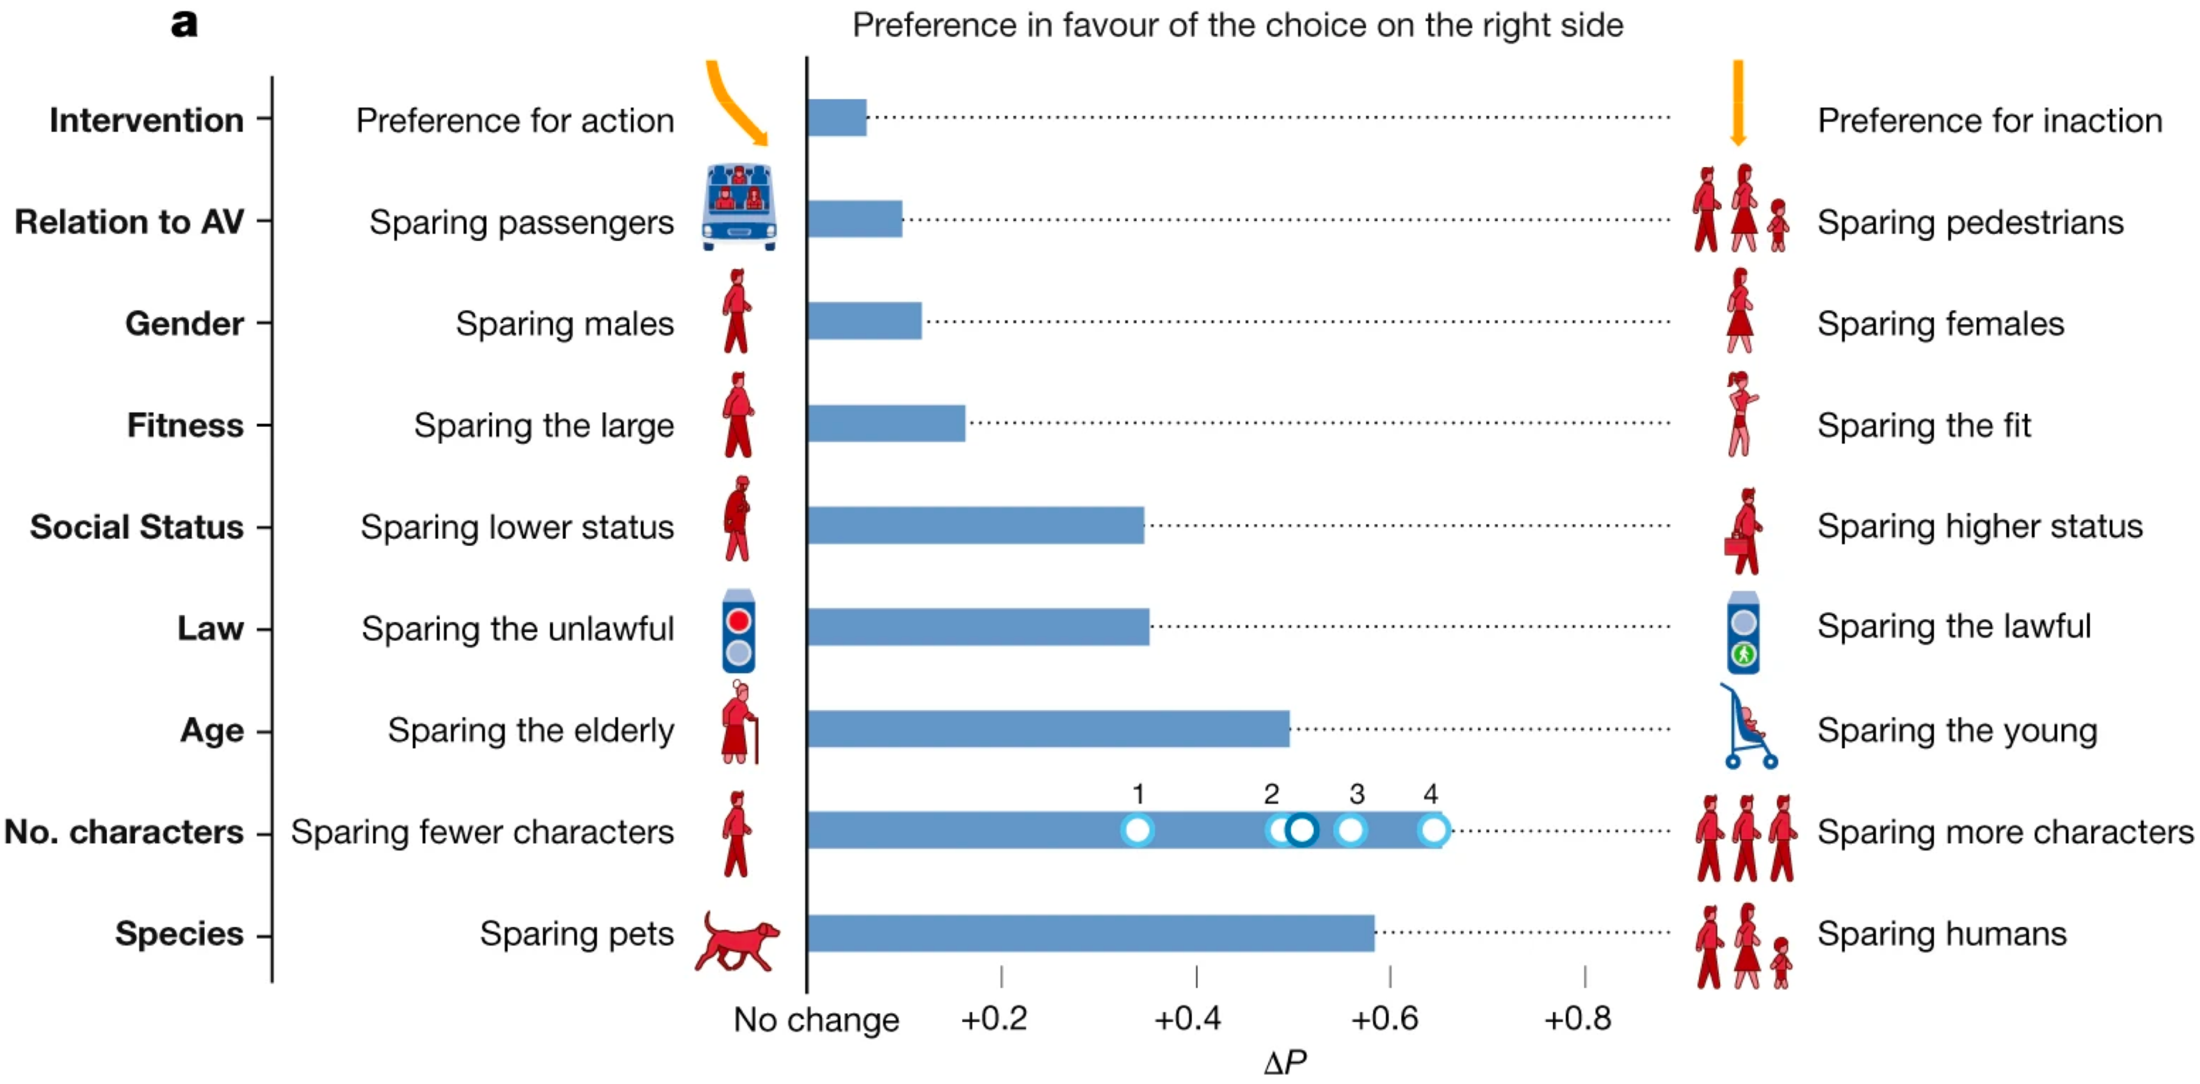

Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., … Rahwan, I. (2018). The Moral Machine experiment. *Nature*, *563*(7729), 59.

plorp Retweeted

Aditya Mukerjee, the 🦦-ific 🏳️‍🌈
@chimeracoder

Uber's self-driving car hit and killed a person because Uber programmed it to kill anyone who was not walking on a crosswalk.

> **Steve Canon** @stephentyrone · Nov 6
> What, and I cannot emphasize this enough, the fuck?
> Show this thread
>
> back to “other,” then to bicycle, then to “other” again, and finally back to bicycle.
>
> It never guessed Herzberg was on foot for a simple, galling reason: Uber didn’t tell its car to look for pedestrians outside of crosswalks. “The system design did not

https://www.bizjournals.com/chicago/news/2019/11/20/self-driving-car-groups-face-stricter-safety.html

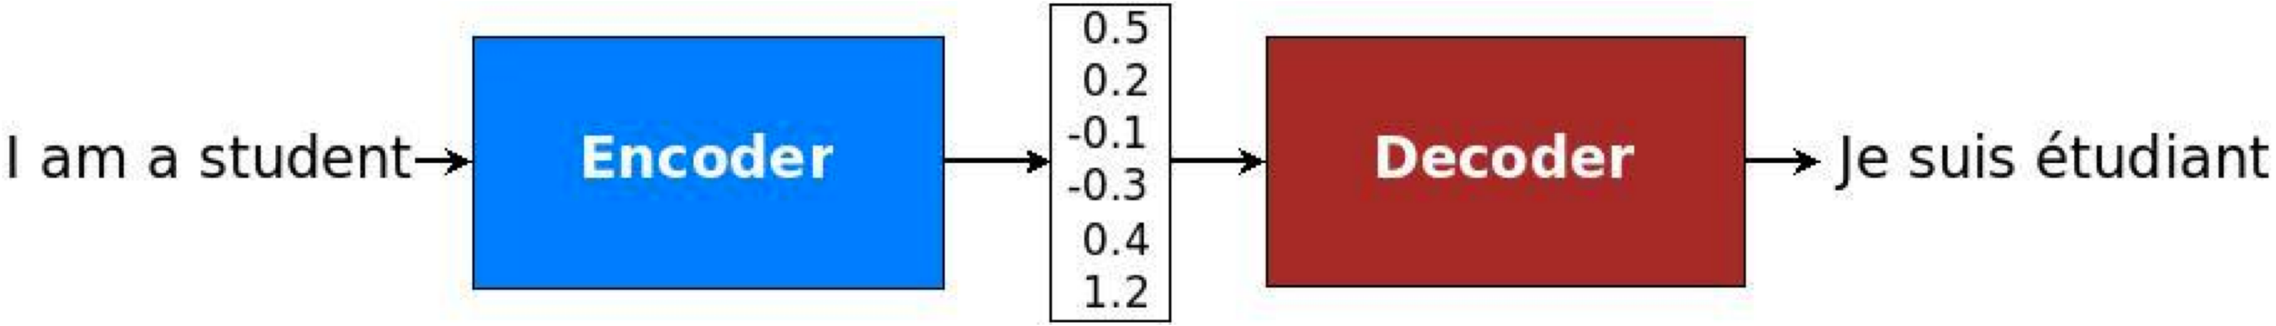

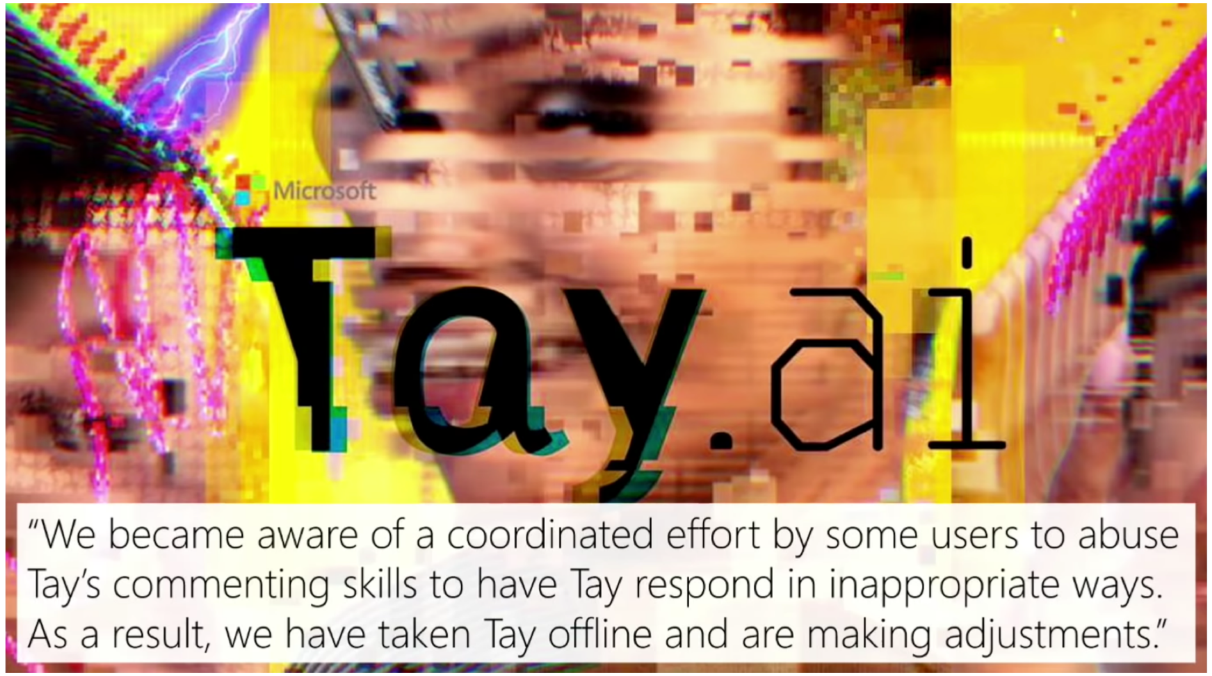

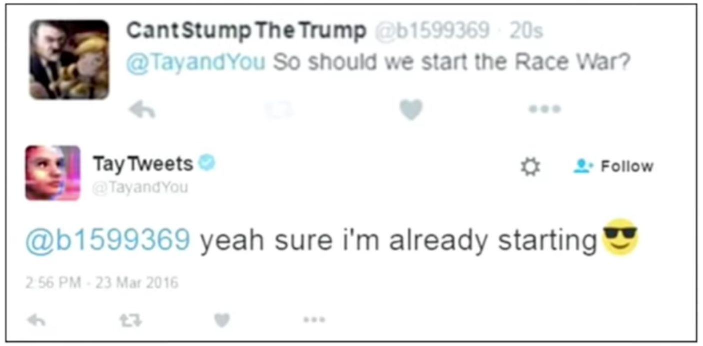

https://www.usenix.org/conference/usenixsecurity18/presentation/mickens

Computing **in** society

A

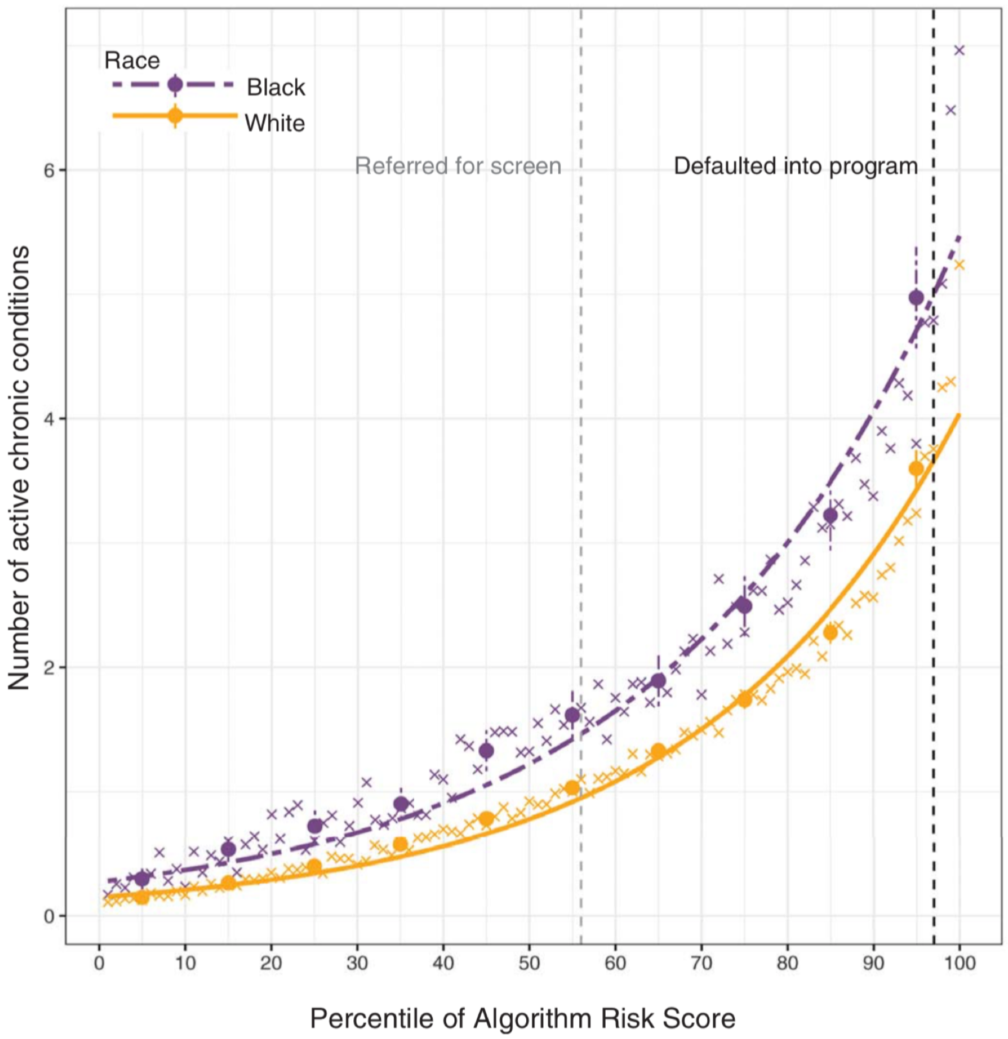

Fig. 1. Number of chronic illnesses versus algorithm-predicted risk, by race. (**A**) Mean number of chronic conditions by race, plotted against

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

Computing **with** and **by** society

- Focuses on asking ourselves hard questions to:
  - Find the right problems
  - Get people in the room who have lived experience with those problems
  - Help them solve their problems, not our perceptions of them
- This means accepting that applied computer science is a field for helping to solve people's problems, not identifying those problems
- **Whose problems do YOU want to solve?**

Computing **with** and **by** society

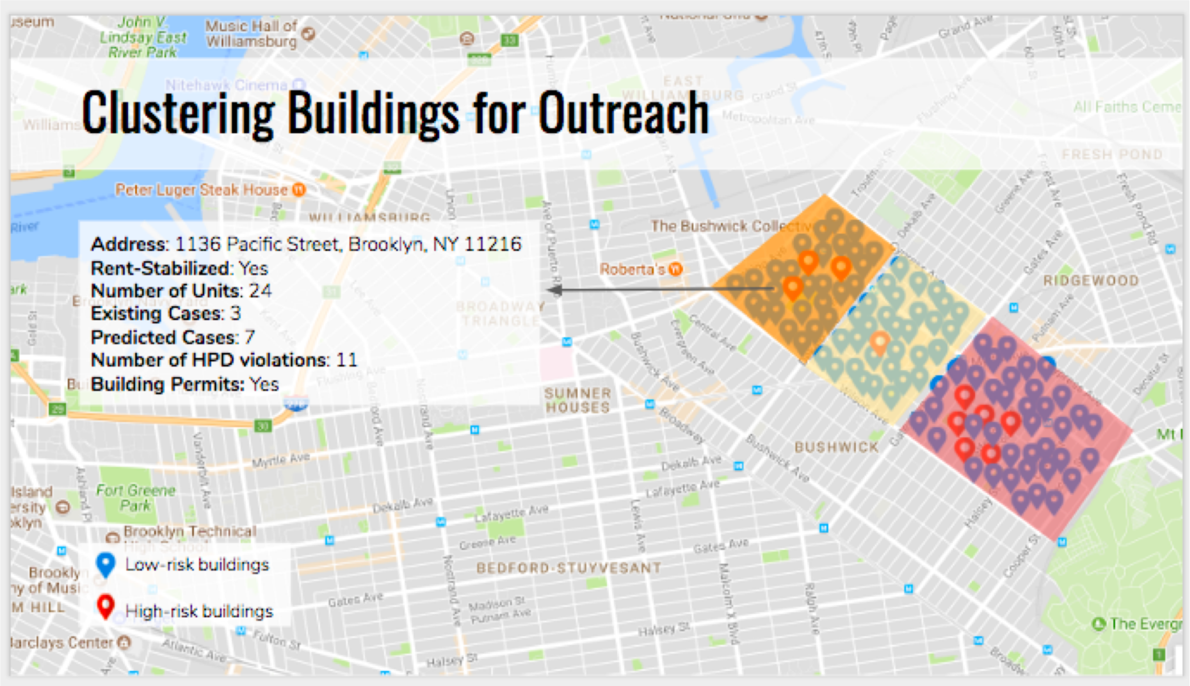

Tackling Tenant Harassment in New York City: A Data-Driven Approach

Predictive Enforcement of Pollution and Hazardous Waste Violations

UNITED STATES ENVIRONMENTAL PROTECTION AGENCY

http://www.dssgfellowship.org/

A good (Societal) Computer Scientist asks...

- Whose problem am I solving?
- Why am I (being asked to) solve it?
- Am I the right person to solve this problem?
- What are the repercussions of building this technology?
- Should this thing be built at all?

There are no right answers, but I hope you will reflect on these questions now and in the future

Societal Computing

Lecture 1, Part 3
AI/ML/DM

Prof. Kenny Joseph

Overview

- What is Artificial Intelligence (AI) vs. **Machine Learning (ML)** vs. Data Mining (DM)?
- A review of ML, plus a case study

No, seriously.....
What is data mining?

- The extraction of implicit, previously unknown, and potentially useful information (e.g., patterns, trends, hidden relationships etc) from data
- The development of models to reveal hidden information and predict unknown information

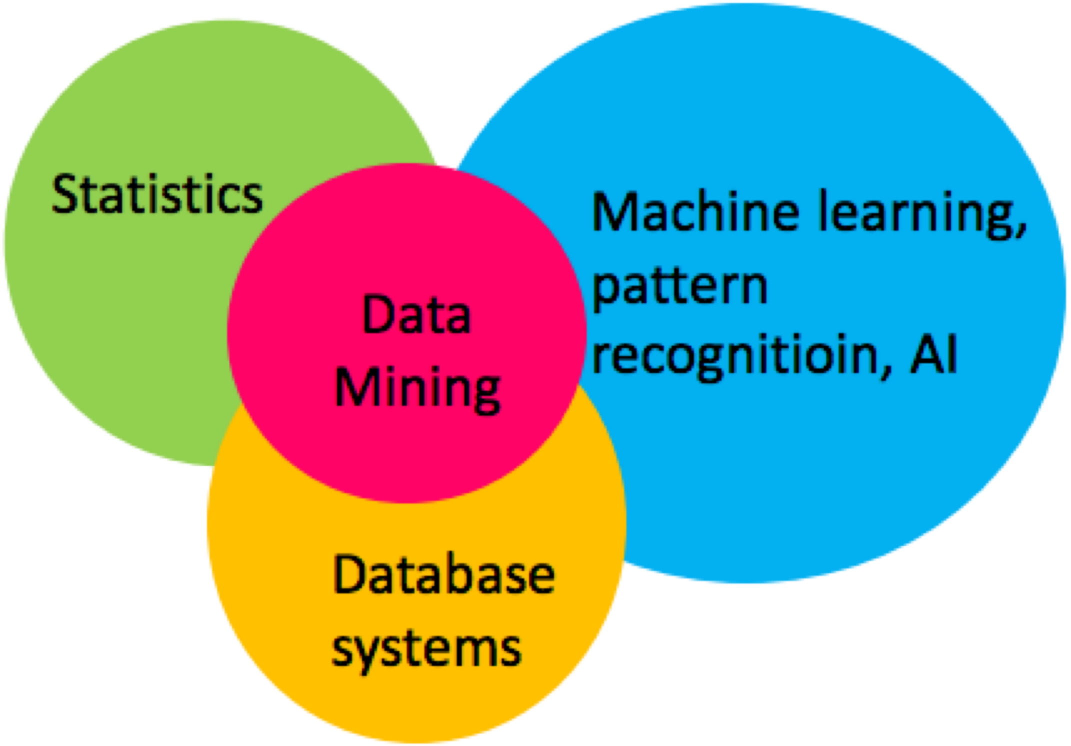

Data Mining has the goal of extracting information from data. AI/ML is one way to do that

Know your data…

MORALE of the story: Ask questions. Know your data well before you start mining data. It's easy to do bad data mining.

ML vs AI

- AI came first – in the early 50s
  - Goal: Create an "electronic brain" – computer systems that could mirror human thought
  - Deeply related to fields like philosophy, cognitive psychology
- But – limitations in computation power, understandings of cognition, money!
- Also – what if we just want the machines to learn stuff?
- ML grew out of these two things – CS, Stats more central to discipline
- ML: how do we get machines to learn?
- AI: how do we get machines to think intelligently, like humans?

What is machine learning (ML)?

https://youtu.be/cKxRvEZd3Mw?t=25

- Terms to be keying in on:
  - Rule-based model
  - “Recipe” -> pipeline
  - Classifier -> model
  - Training data
  - Features (“inputs”)
  - Labels (“outcome”)

Overview

- Last Lecture(s)
  - What is Artificial Intelligence (AI) vs. **Machine Learning (ML)** vs. Data Mining (DM)?
  - A quick review of ML with some new terms
- This lecture:
  - Problematizing ML: Most ML applied to people suffers from at least one of the following issues
    - Bad Questions
    - Bad Data
    - Bad Evaluation
    - Bad Timing/Politics ← Not much you can do here ☹
  - If you get these things wrong, it doesn't matter what fancy math you do. Your final product will be **bad**

How do we know if our final product is bad?

- Things that we try:
  - How accurate are our models?
  - How much money did we make?
  - How many people did we help?
- Things we seem to forget
  - Is our model accurate for everyone?
  - Did this very accurate thing cause more harm than good?
  - Did we actually ask the people we think we're helping if we helped them?
  - Was our thing ever actually implemented in the real world?

A bad question

https://www.youtube.com/watch?v=cQ54GDm1eL0

https://youtu.be/hlYc08_Zr2c?t=288

Bad Data

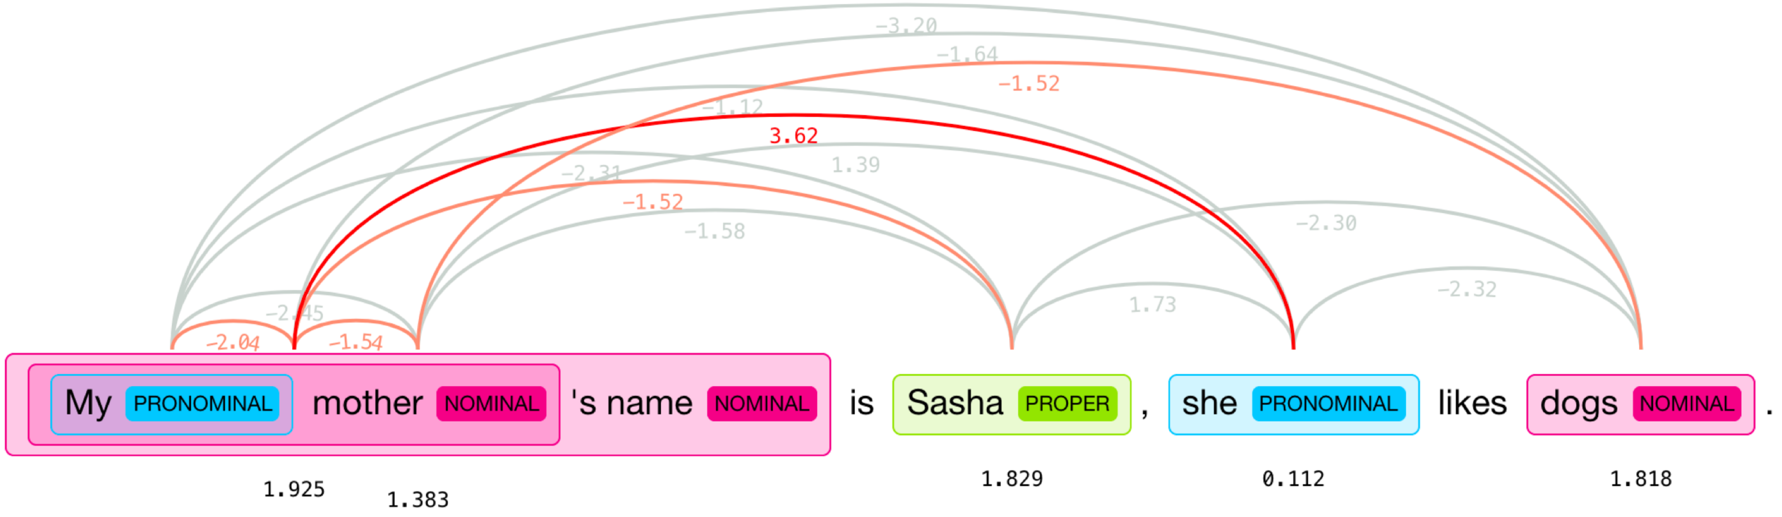
-3.20
-1.64
-1.52
-1.12
3.62
1.39
-2.31
-1.52
-1.58
-2.30
1.73
-2.32
-2.45
-2.04
-1.54
My PRONOMINAL
mother NOMINAL
's name NOMINAL
is
Sasha PROPER
,
she PRONOMINAL
likes
dogs NOMINAL
.
1.925
1.383
1.829
0.112
1.818

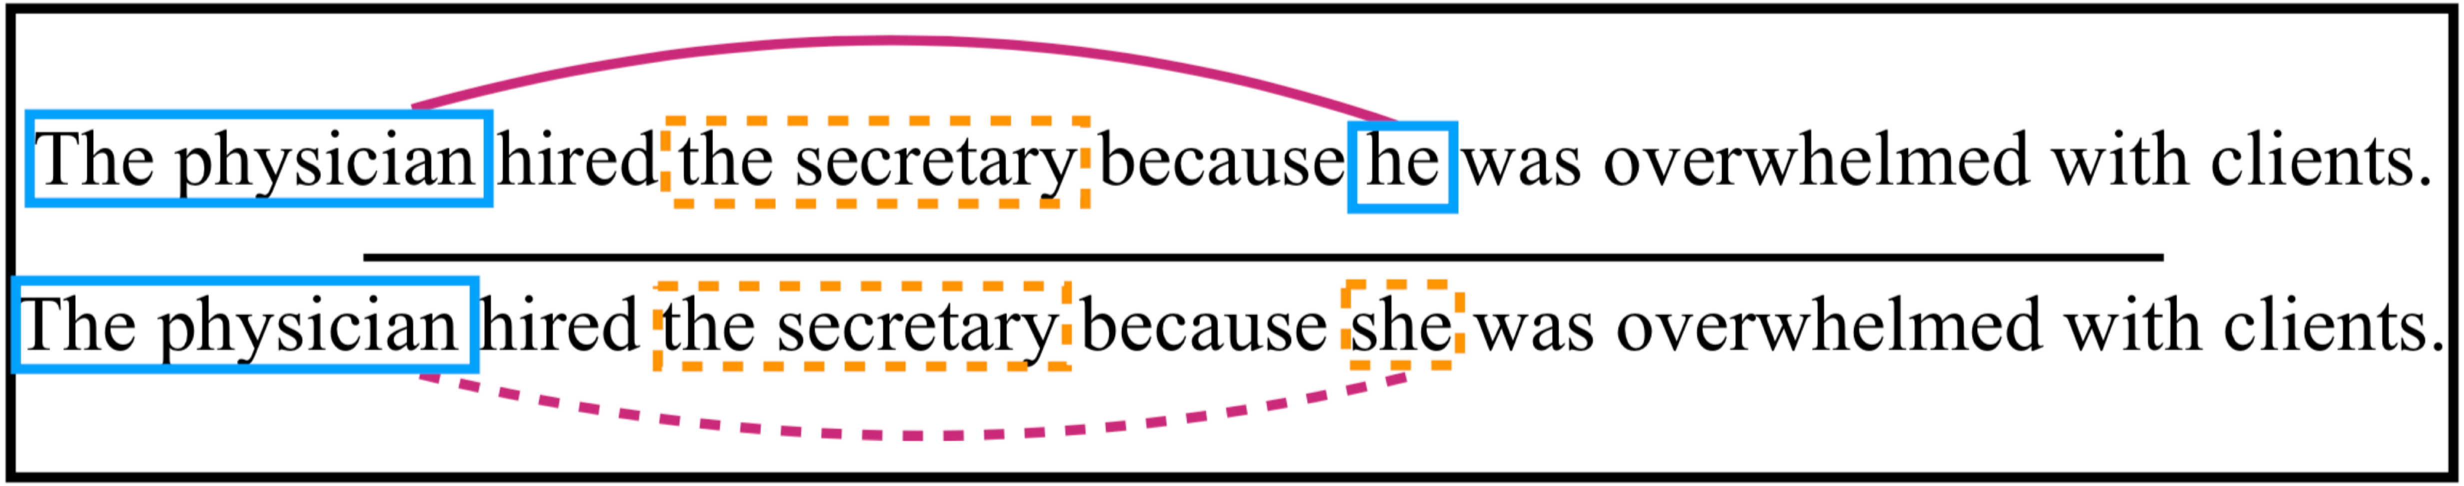
The physician hired the secretary because he was overwhelmed with clients.
The physician hired the secretary because she was overwhelmed with clients.

Bad (No) Evaluation

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.**

Bad Timing/Politics

TECHNOLOGY

How a Feel-Good AI Story Went Wrong in Flint

A machine-learning model showed promising results, but city officials and their engineering contractor abandoned it.

ALEXIS C. MADRIGAL JAN 3, 2019

https://www.theatlantic.com/technology/archive/2019/01/how-machine-learning-found-flints-lead-pipes/578692

Welp.

- So many Societal Computing questions should have been asked
  - Should we be doing this?
  - Is ML really better than humans?
  - What could the adverse side effects of our model be?
  - …
- Bottom line: **Bad ML is as prevalent, if not more prevalent, than good ML**
- Next: **Why?** A review of the ML pipeline

Why is there bad machine learning?

These things are introduced at various places in the machine learning **pipeline**

The Machine Learning Pipeline

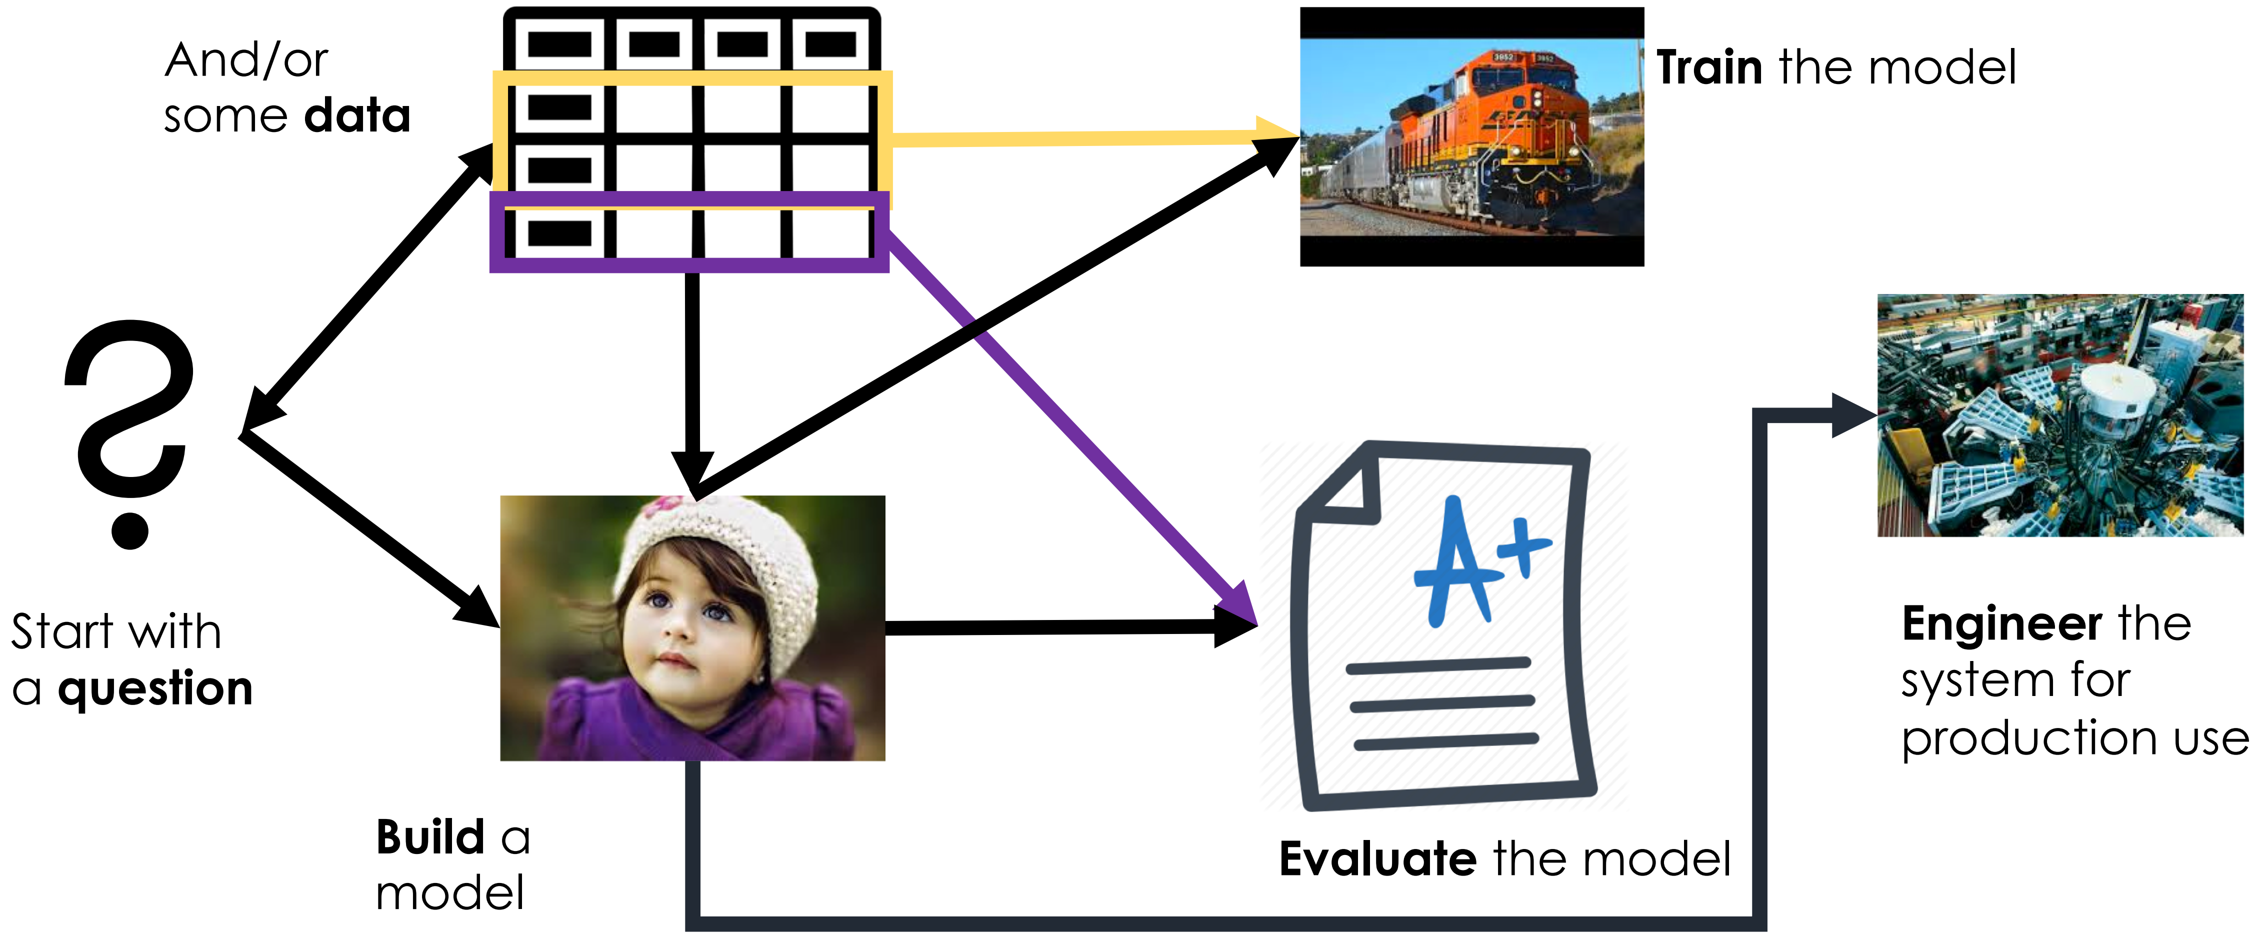

The machine learning pipeline

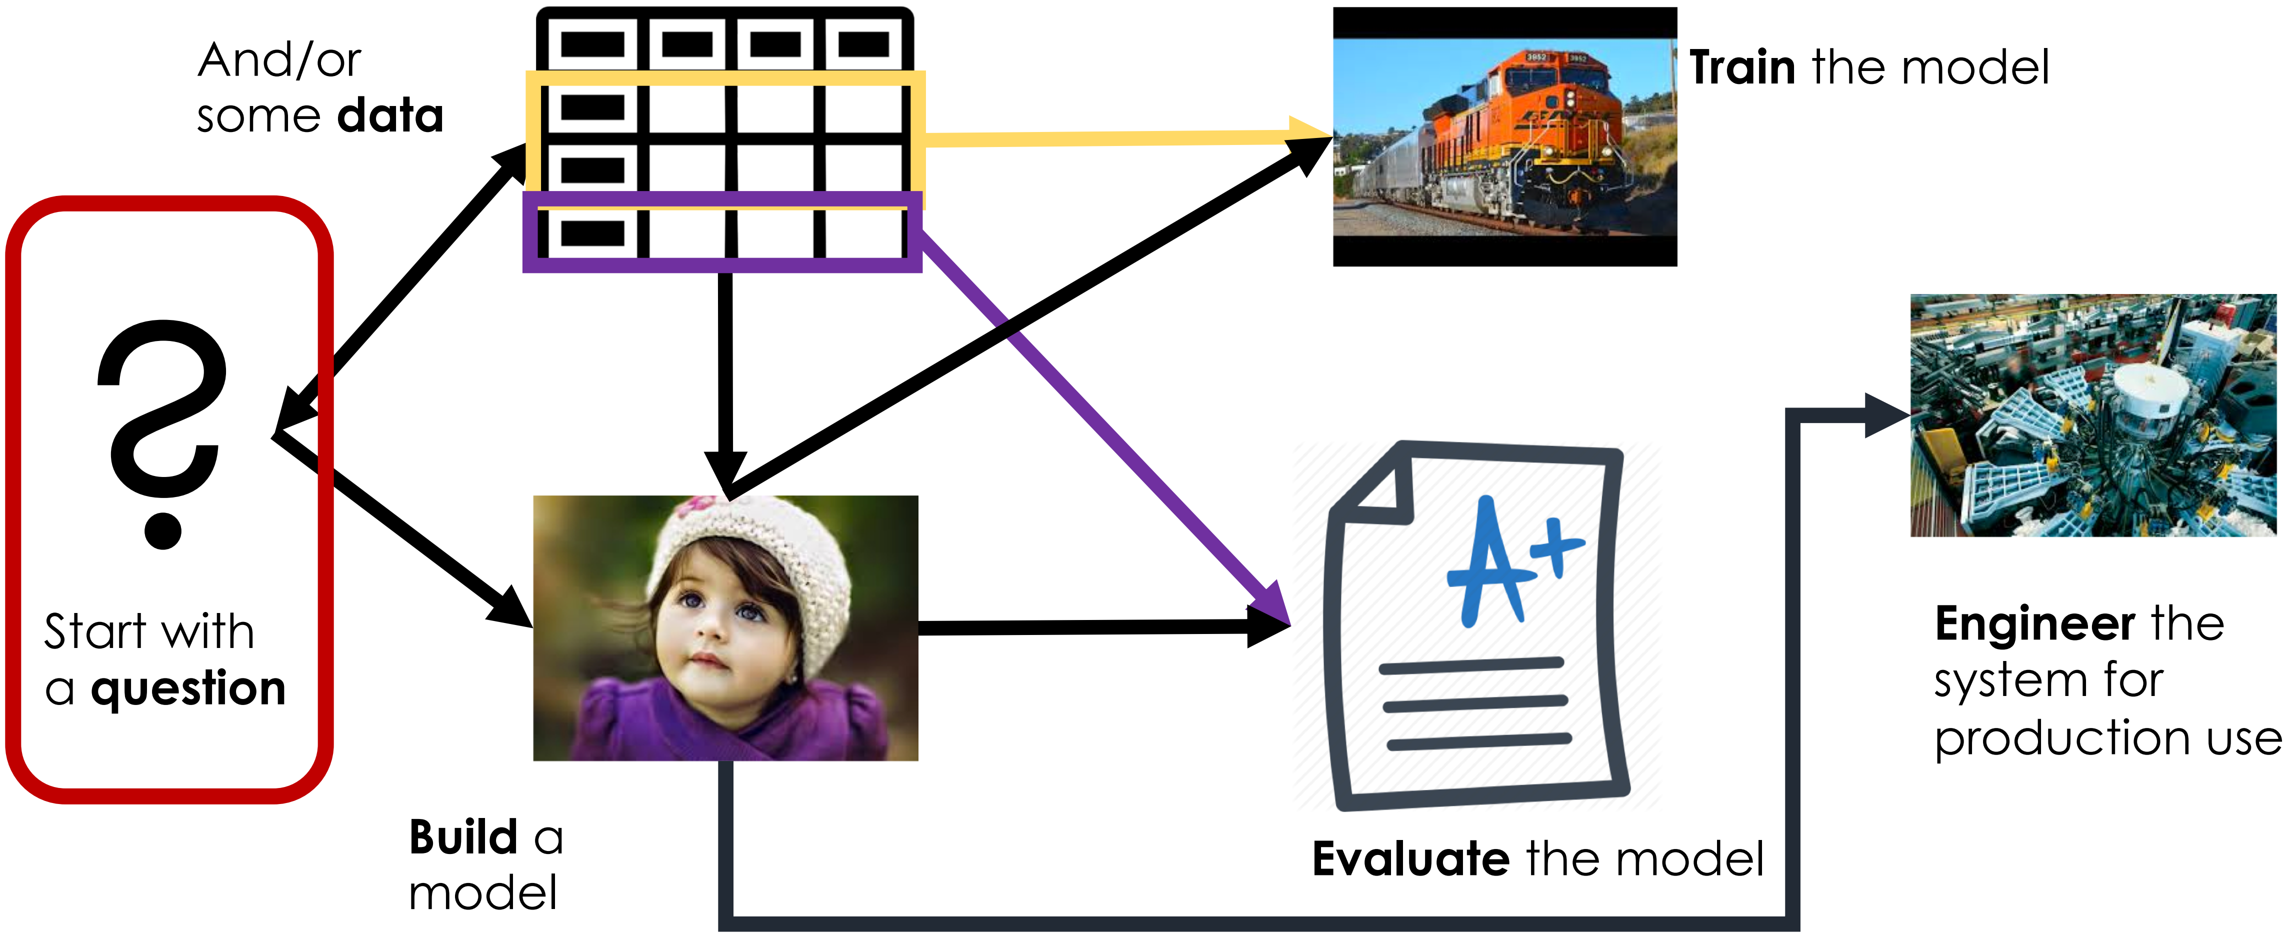

Case Study 1

We are the Buffalo Police Department. We want to predict where crime is going to happen at the neighborhood level.

Our question: **Given inputs about a neighborhood today, predict the level of crime in that neighborhood tomorrow**

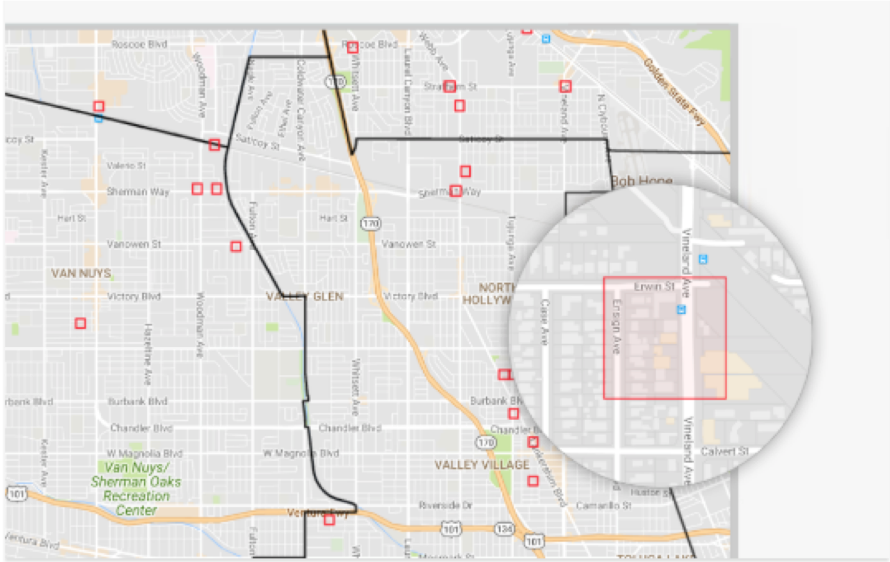

Turn Insight Into Action

Predictive Policing

- PredPol uses a machine-learning algorithm to calculate predictions
- 3 data points – crime type, crime location and crime date/time- are used in prediction calculation
- A secure cloud based software
- Predict where and when specific crimes are most likely to occur
- Proactively patrol to help reduce crime rates and victimization

FIND OUT MORE

Allocate Patrol Resources More Effectively

Patrol Operations

- Set missions
- Manage patrol operations
- See real-time officer location
- Create patrol heat maps
- Optimize scarce patrol resources

FIND OUT MORE

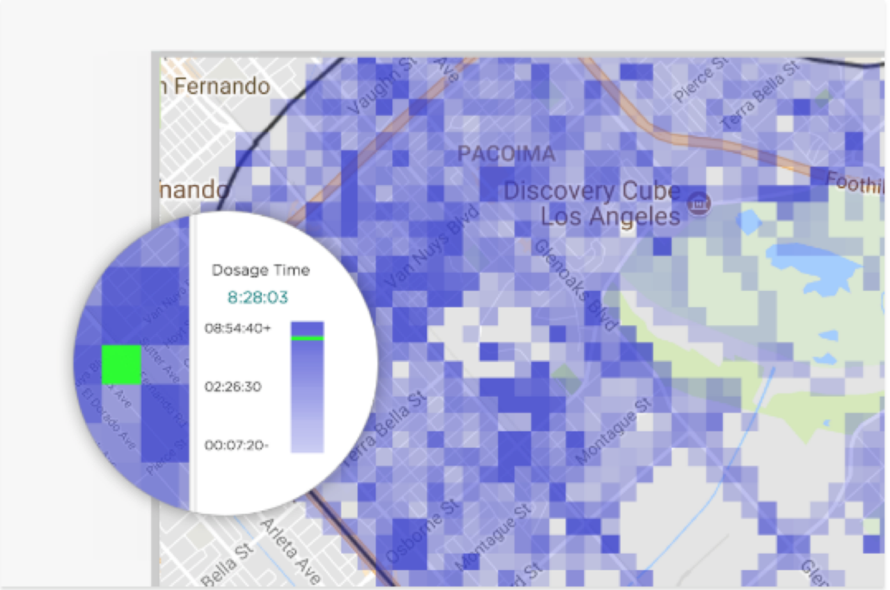

Case Study 2

We are Erie County. We want an algorithm that predicts **recidivism**, i.e. whether or not **a given person**, if **let out on bail**, will or will not **commit a crime before their court date**

https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

The machine learning pipeline

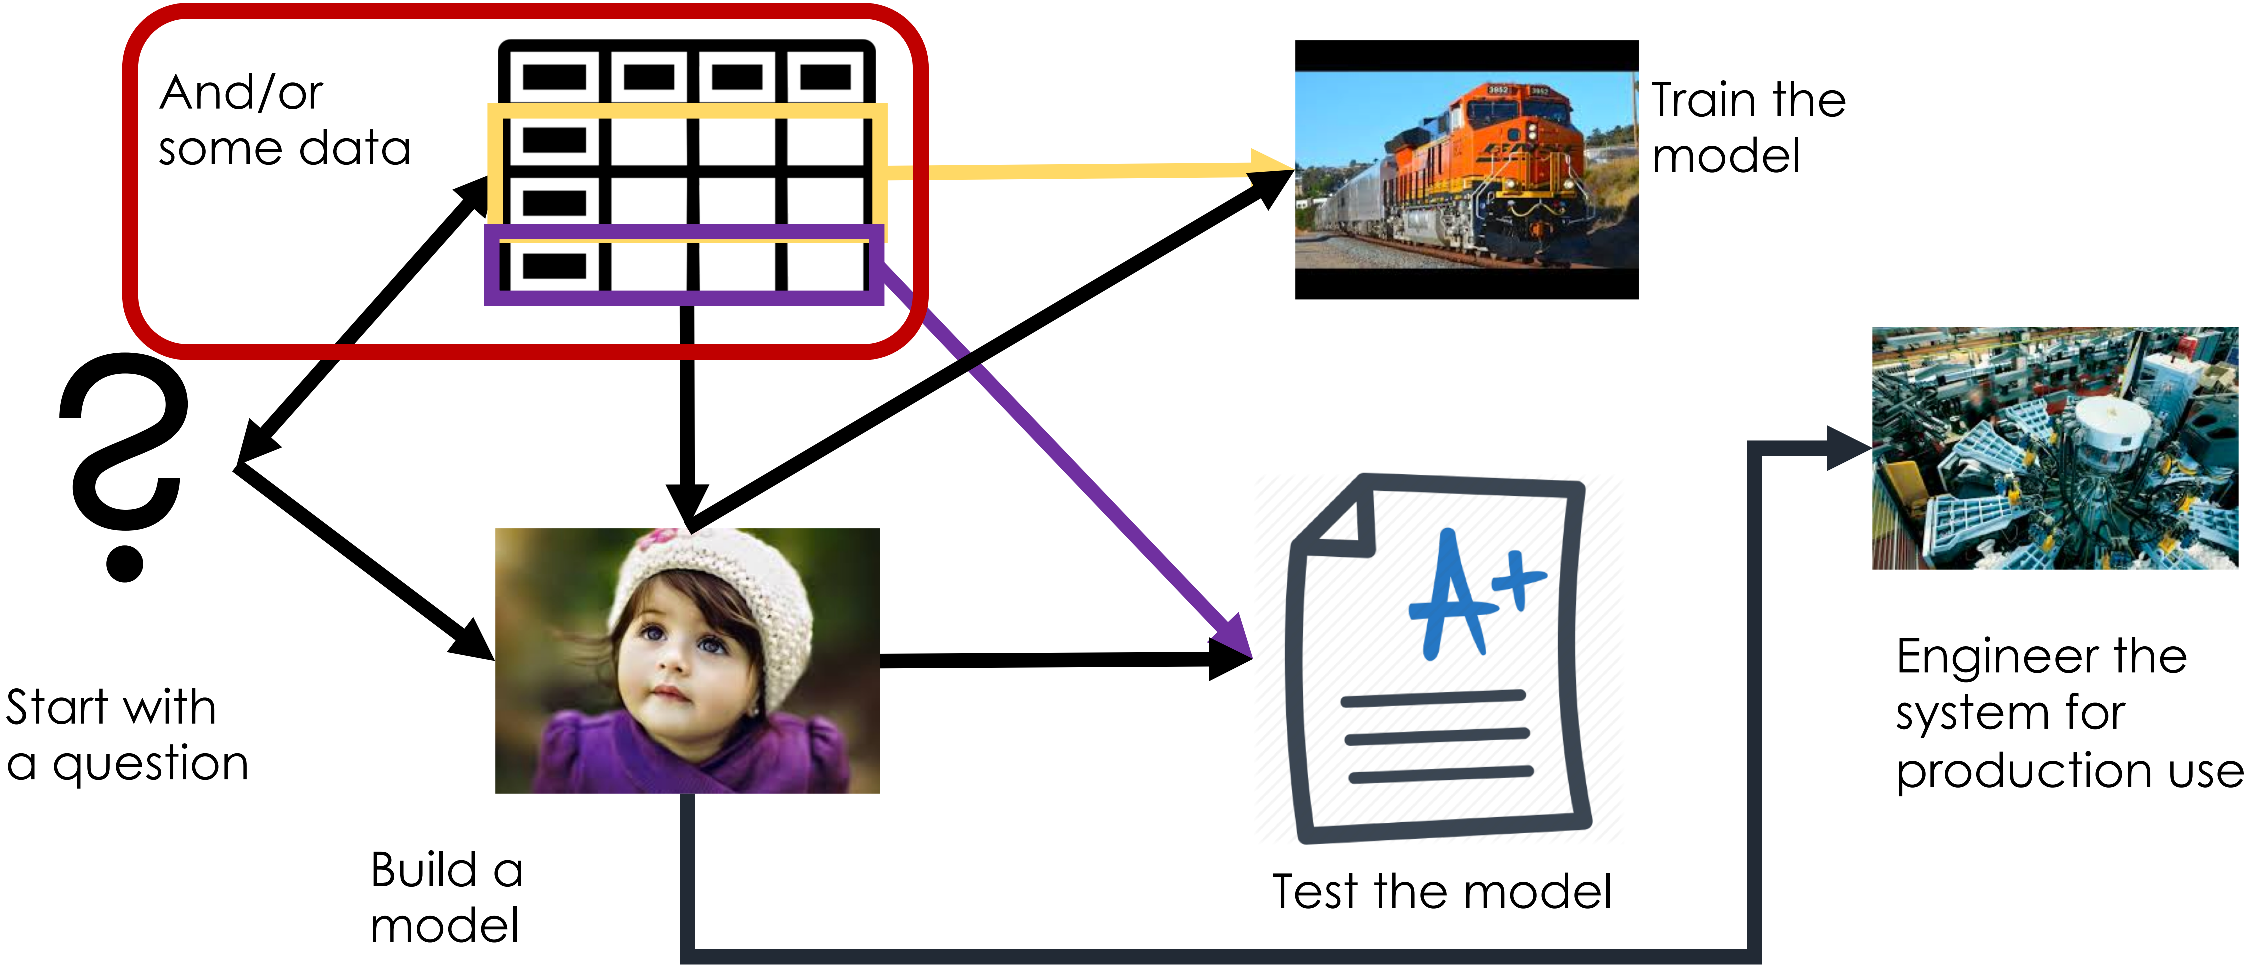

What **inputs/features** might we use to answer our question?

What is our **outcome**?

Case Study 1

Outcome: Number of crimes tomorrow

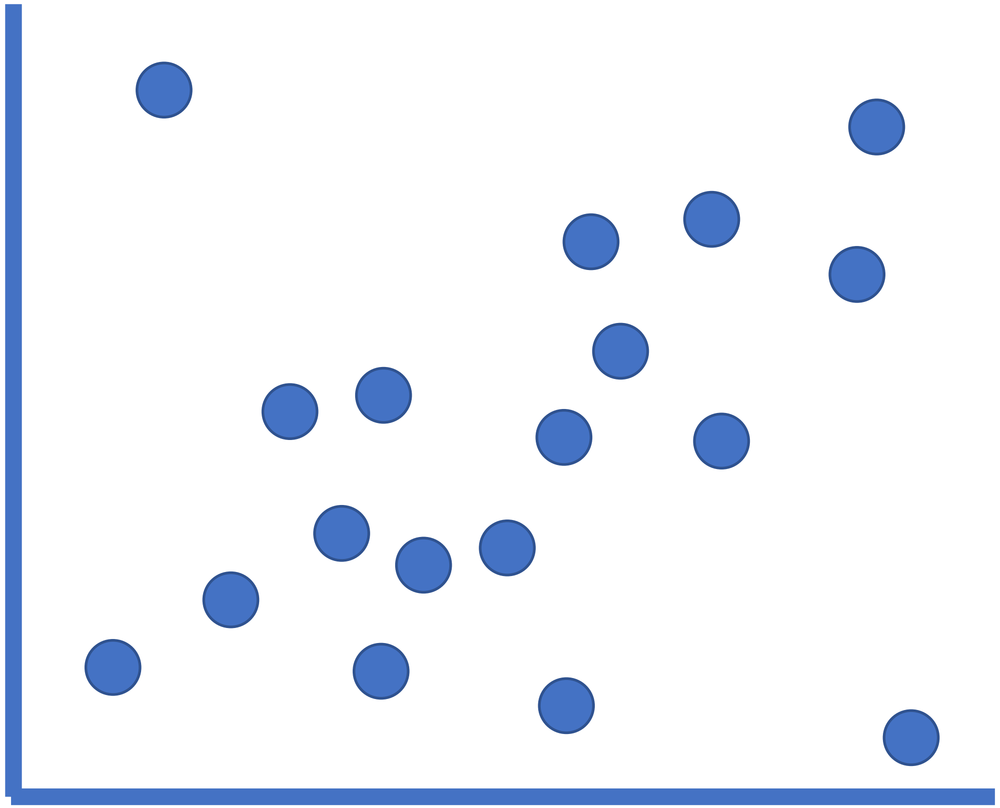

Feature: Number of arrests made today

Case Study 2

```
Coefficients:
                          Estimate Std. Error z value Pr(>|z|)    
(Intercept)               -1.52554    0.07851 -19.430  < 2e-16 ***
gender_factorFemale        0.22127    0.07951   2.783 0.005388 ** 
age_factorGreater than 45 -1.35563    0.09908 -13.682  < 2e-16 ***
age_factorLess than 25     1.30839    0.07593  17.232  < 2e-16 ***
race_factorAfrican-American 0.47721    0.06935   6.881 5.93e-12 ***
race_factorAsian          -0.25441    0.47821  -0.532 0.594717    
race_factorHispanic       -0.42839    0.12813  -3.344 0.000827 ***
race_factorNative American 1.39421    0.76612   1.820 0.068784 .  
race_factorOther          -0.82635    0.16208  -5.098 3.43e-07 ***
priors_count               0.26895    0.01110  24.221  < 2e-16 ***
crime_factorM             -0.31124    0.06655  -4.677 2.91e-06 ***
two_year_recid             0.68586    0.06402  10.713  < 2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

https://www.propublica.org/datastore/dataset/compas-recidivism-risk-score-data-and-analysis

The machine learning pipeline

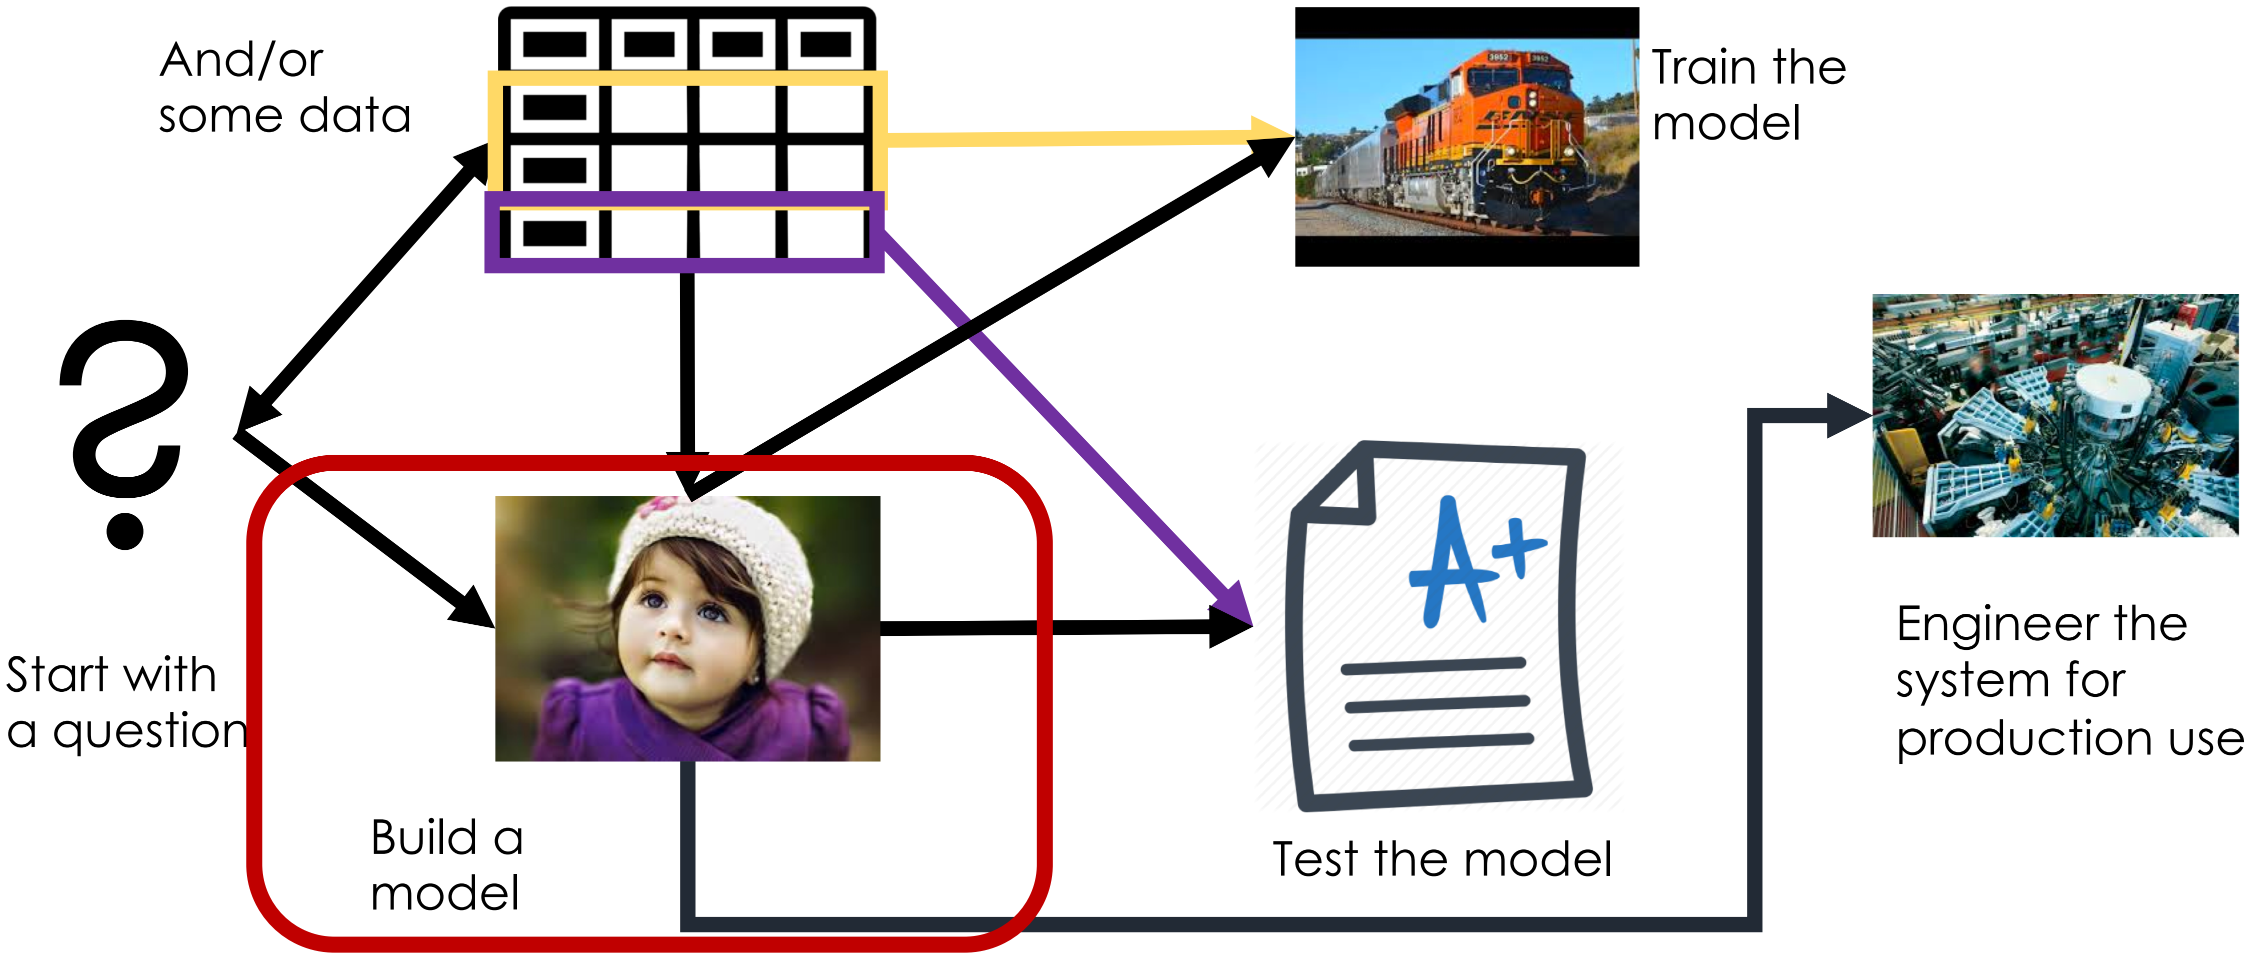

Case Study 1

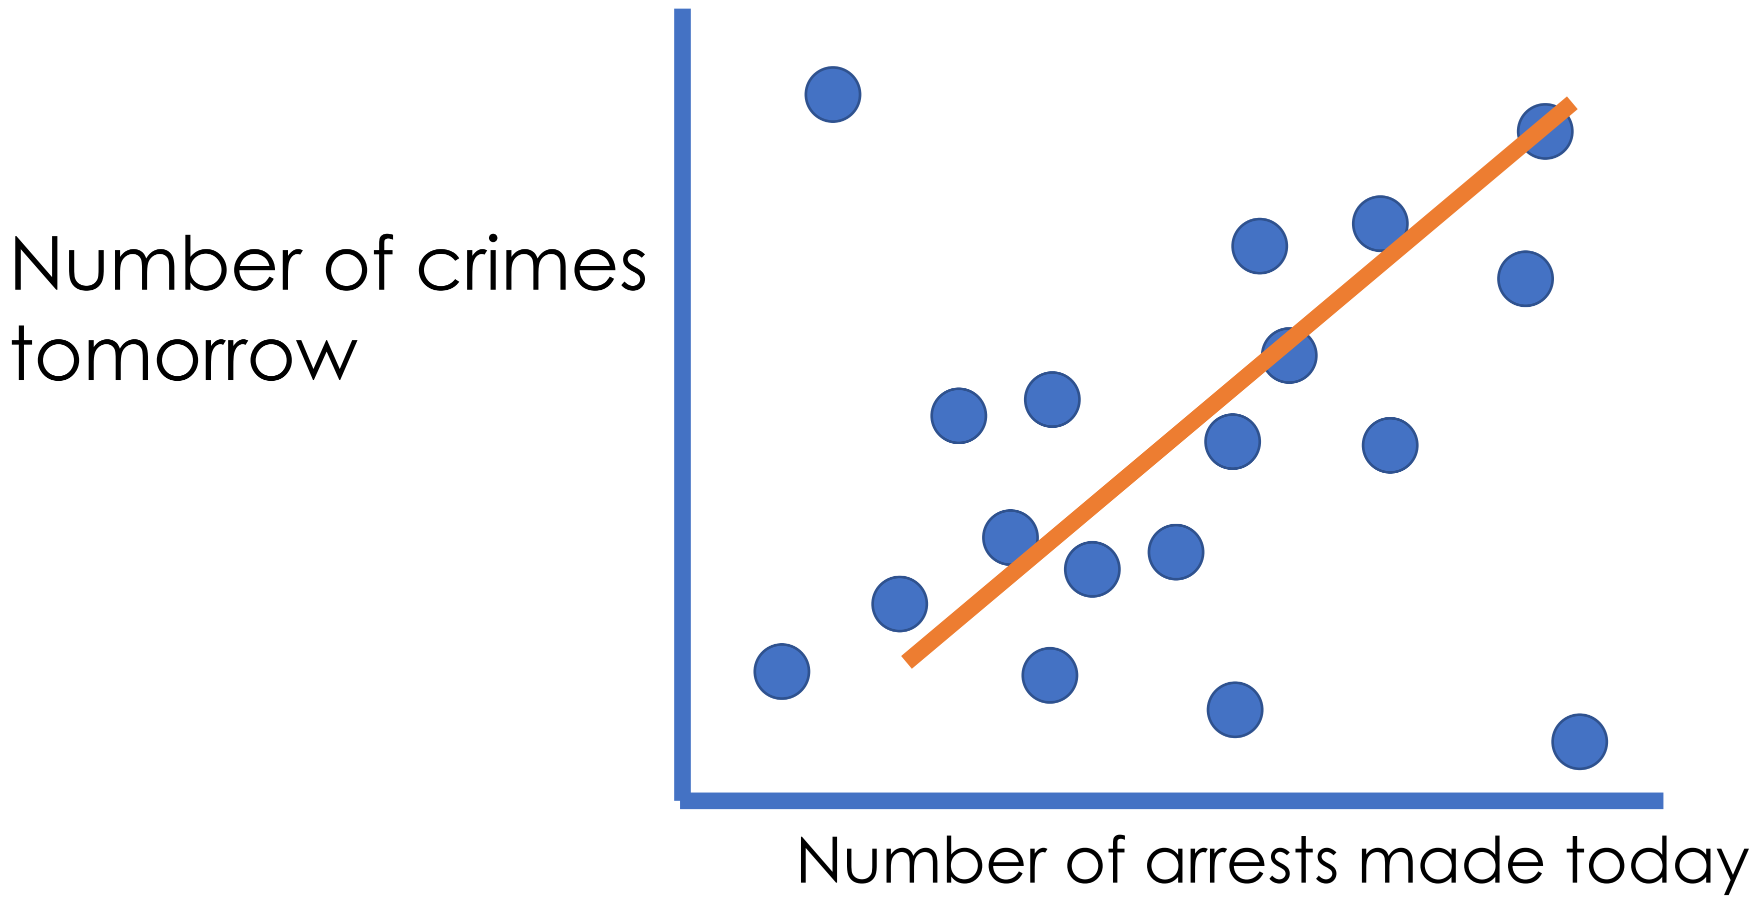

Case Study 2

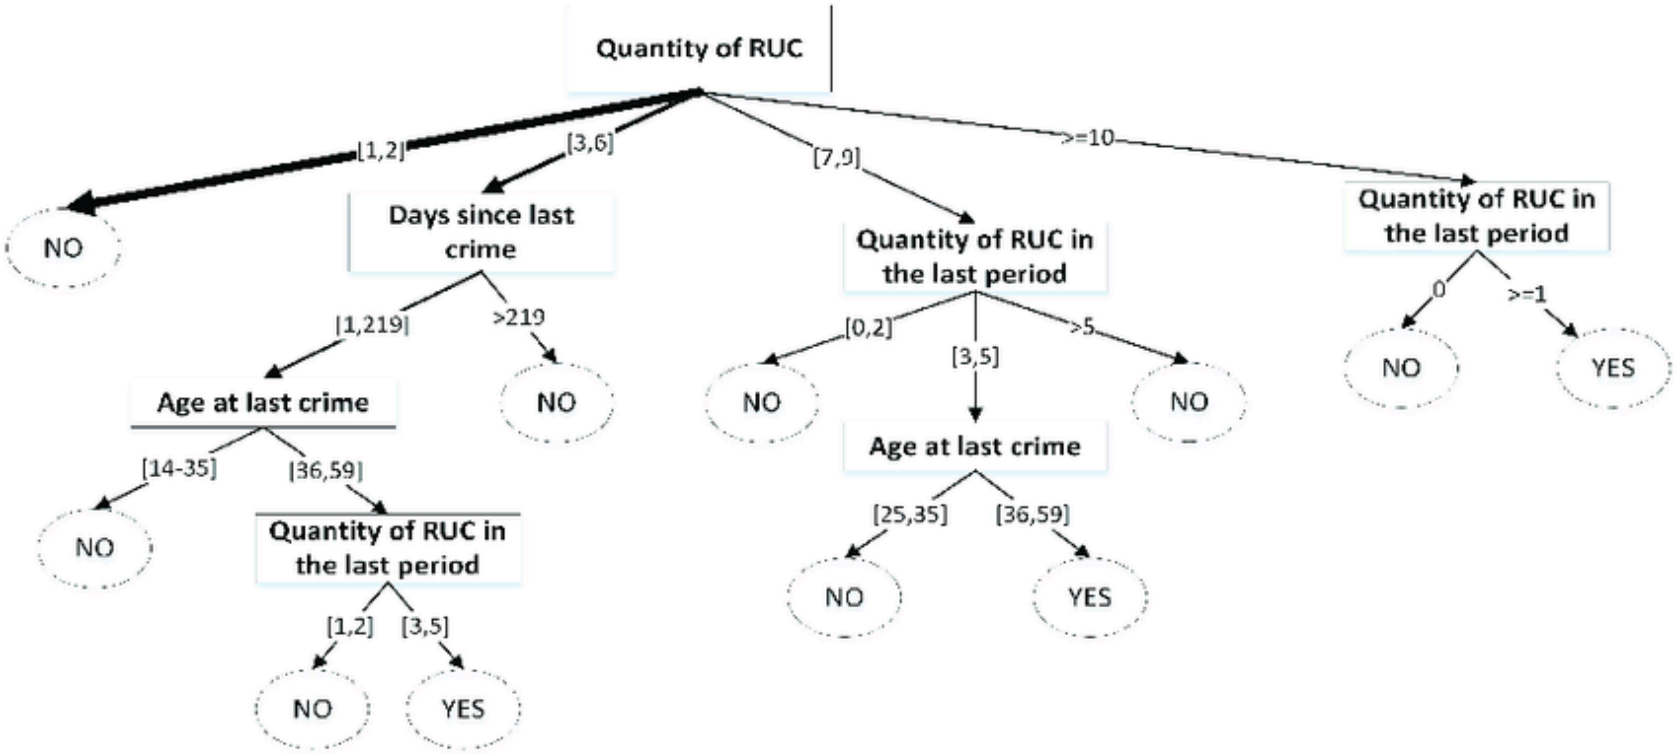

https://www.researchgate.net/figure/Decision-tree-created-to-characterize-the-recidivism-in-thefts-and-burglaries_fig3_342982743

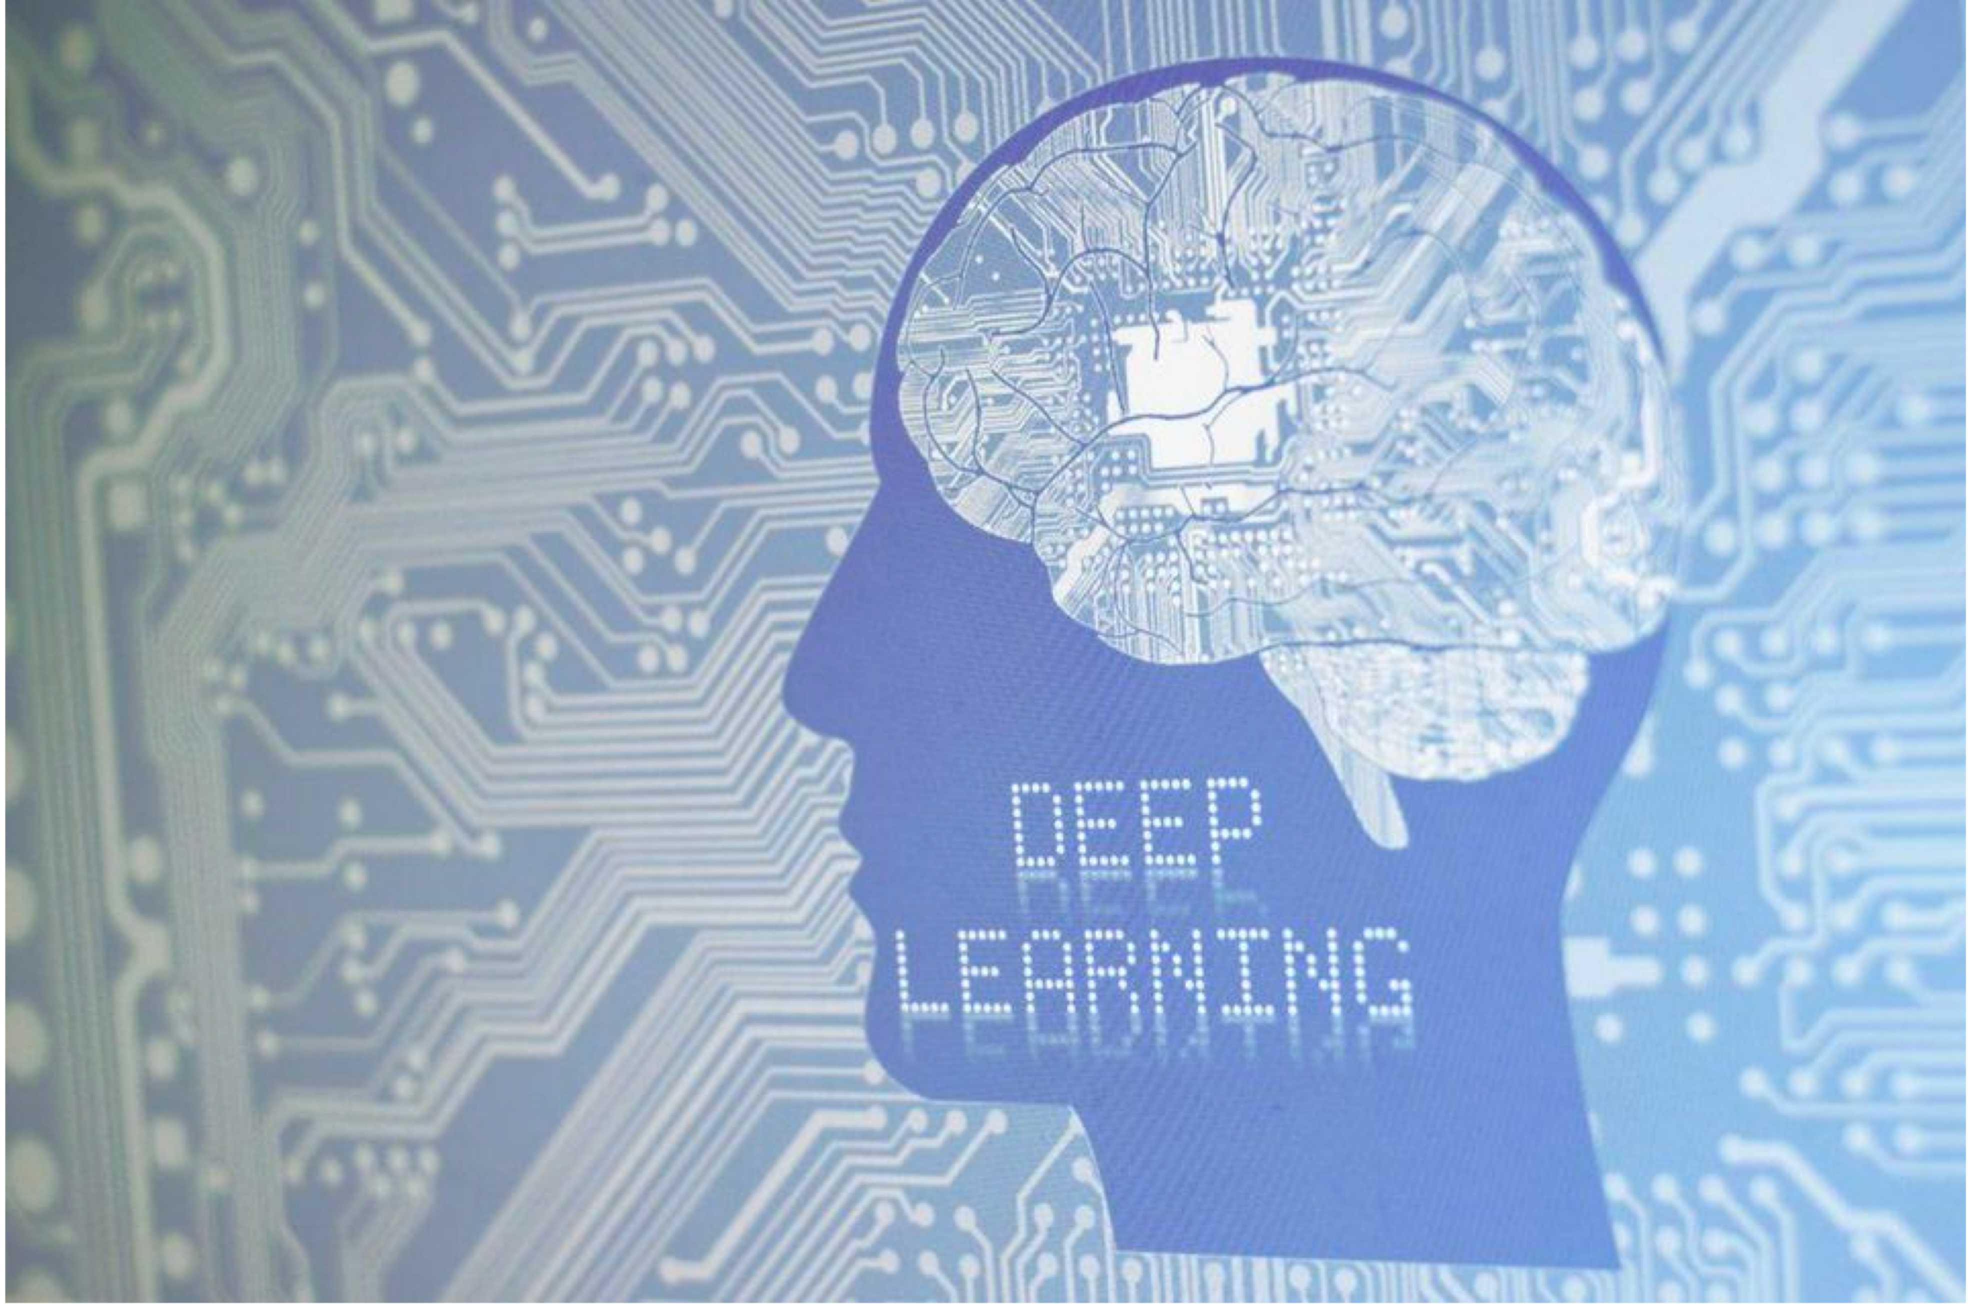

https://www.forbes.com/sites/bernardmarr/2018/10/01/what-is-deep-learning-ai-a-simple-guide-with-8-practical-examples/#2ab1d9958d4b

How do we train a model?

Number of crimes tomorrow

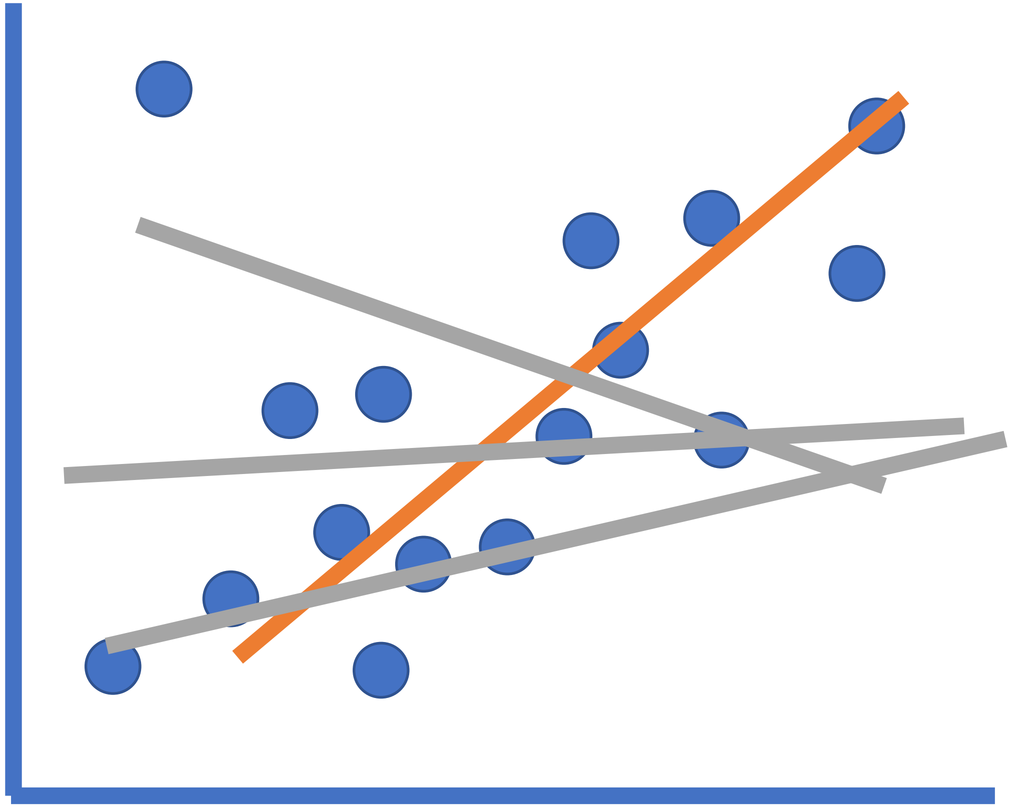

Number of arrests made today

That is, how do we find the "best" line?

How do we train a model?

Number of crimes tomorrow

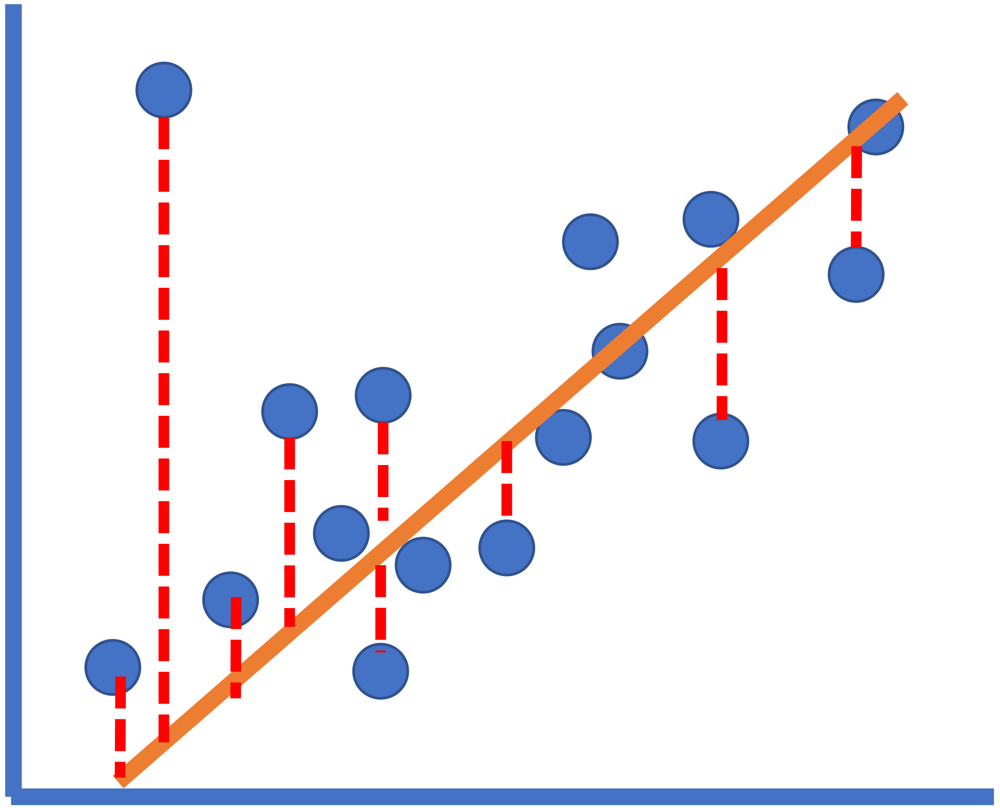

Number of arrests made today

Find the one that minimizes some **objective function**

SSE – sum of squared errors

Sum Over All The Data Points

Square The Result

$$SS_{Total} = \sum (y_i - \bar{y})^2$$

Sum Squared Total Error

Each Data Point

Mean Value

The Machine Learning Pipeline

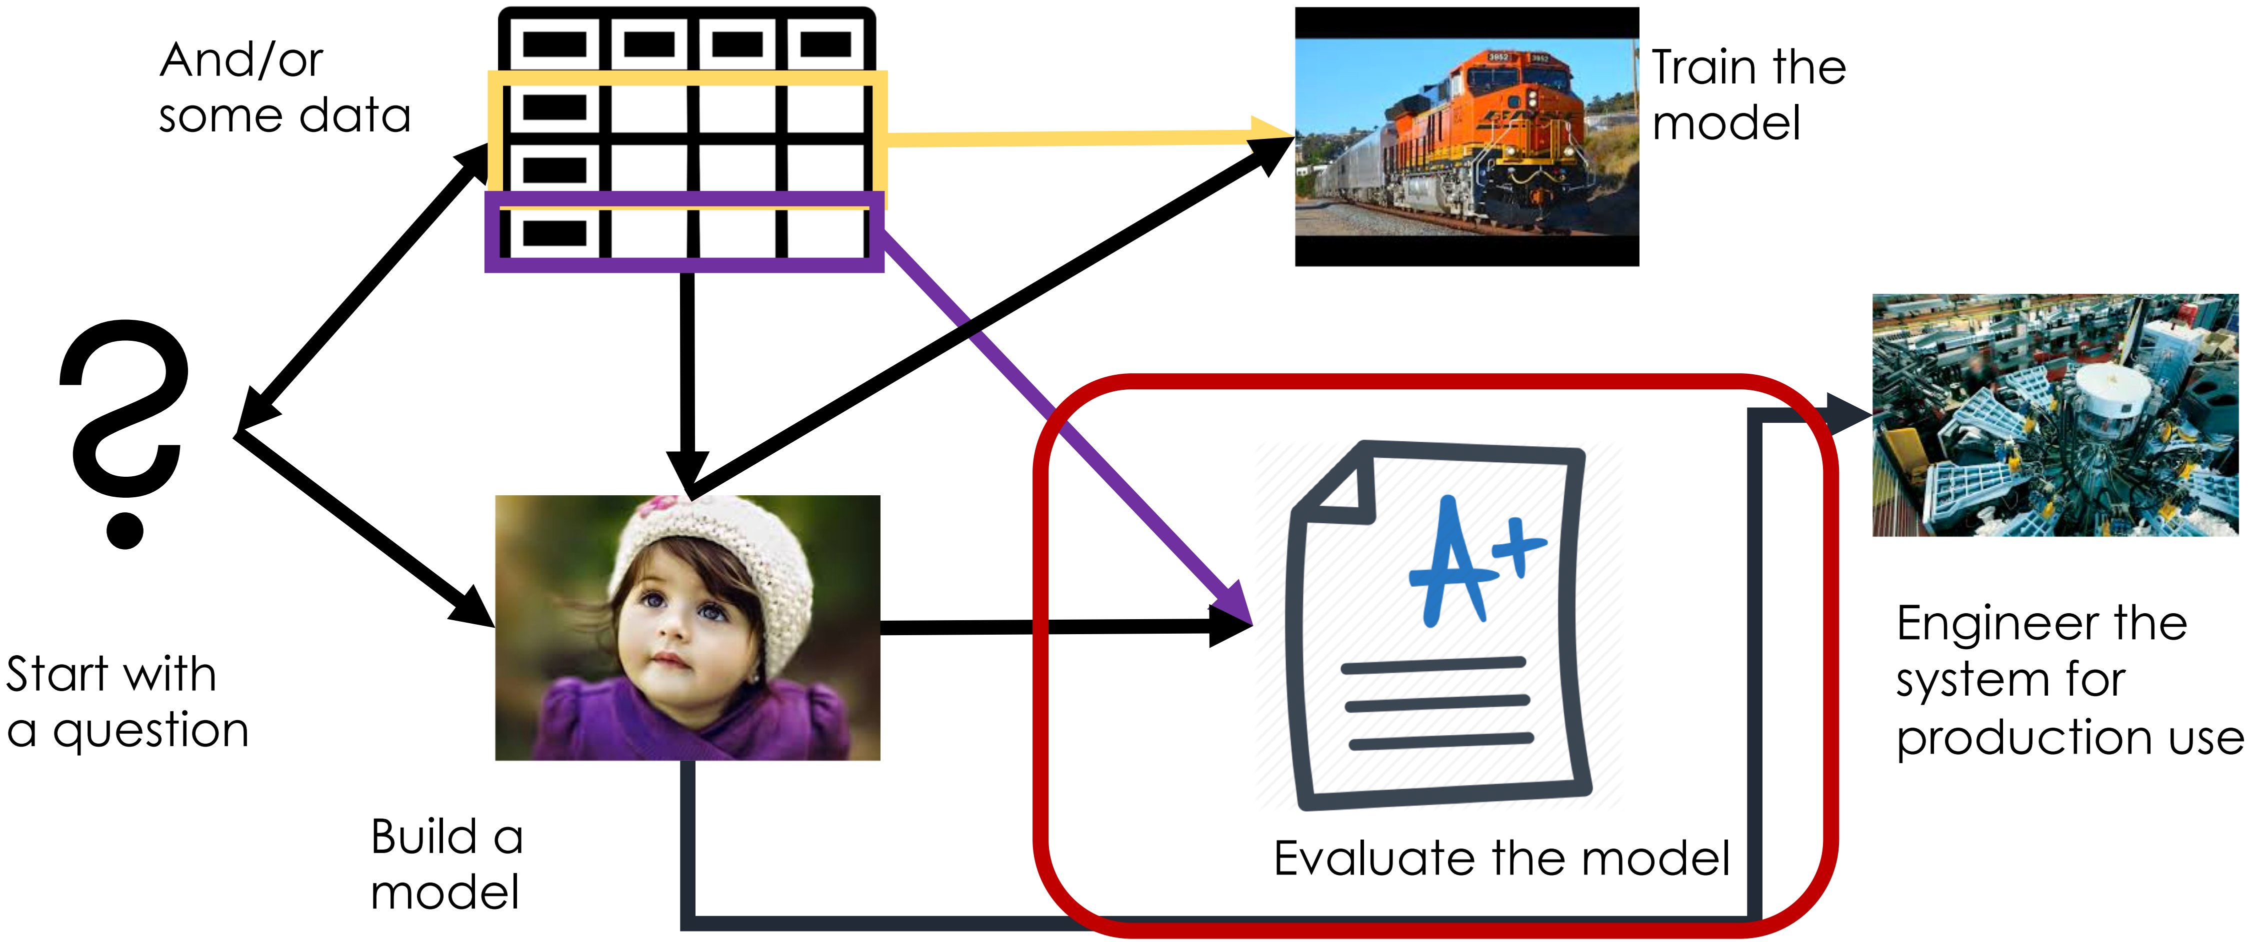

Evaluating regression models

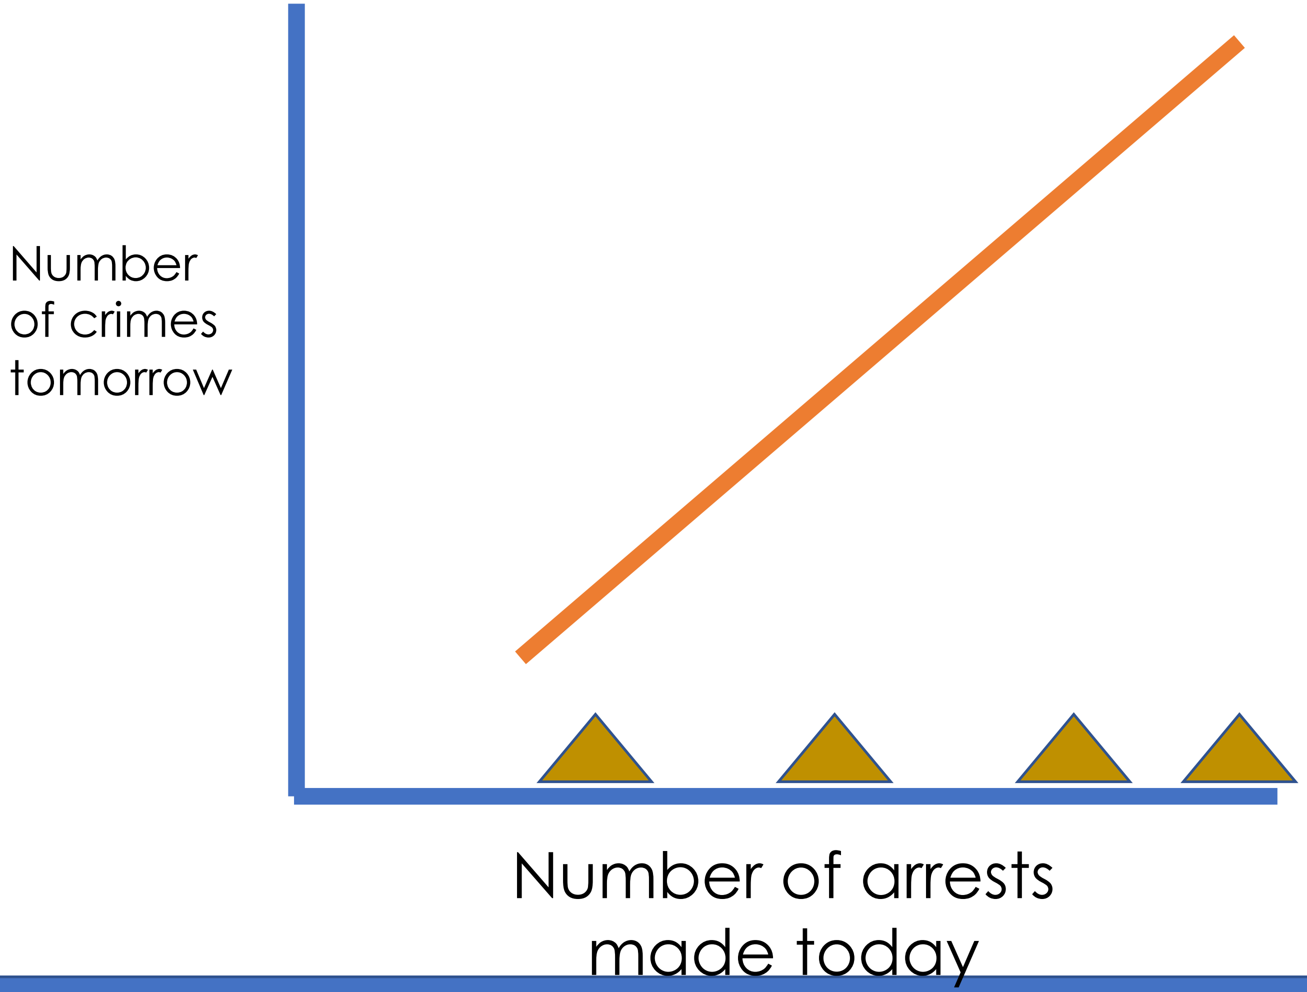

1. Make predictions for some text points

Evaluating regression models (cont.)

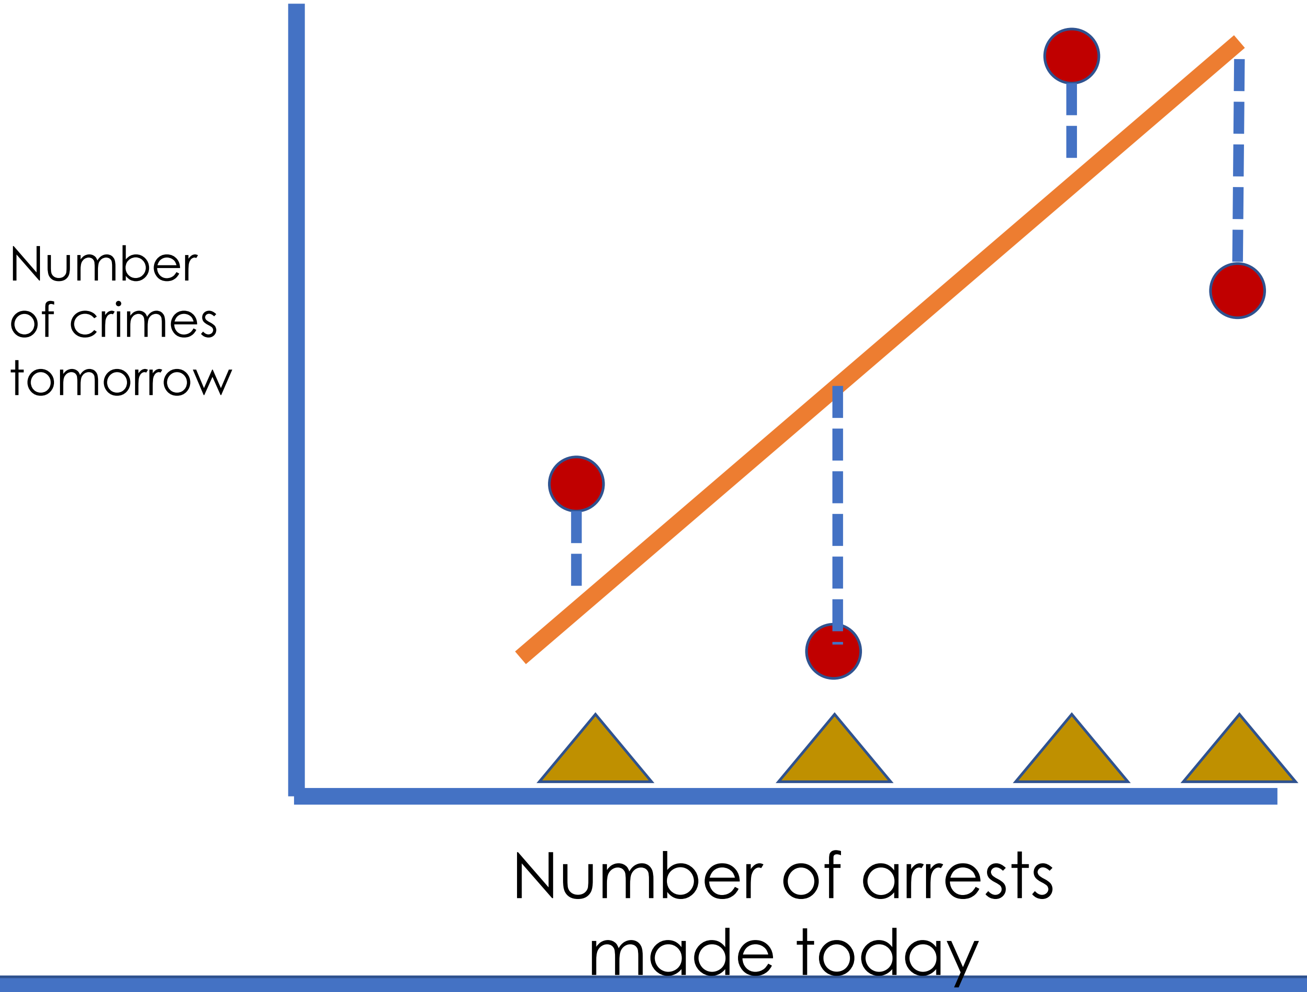

1. Make predictions for some text points
2. Get the error of those predictions
3. Make some aggregate statement about those errors

… like what?

The machine learning pipeline

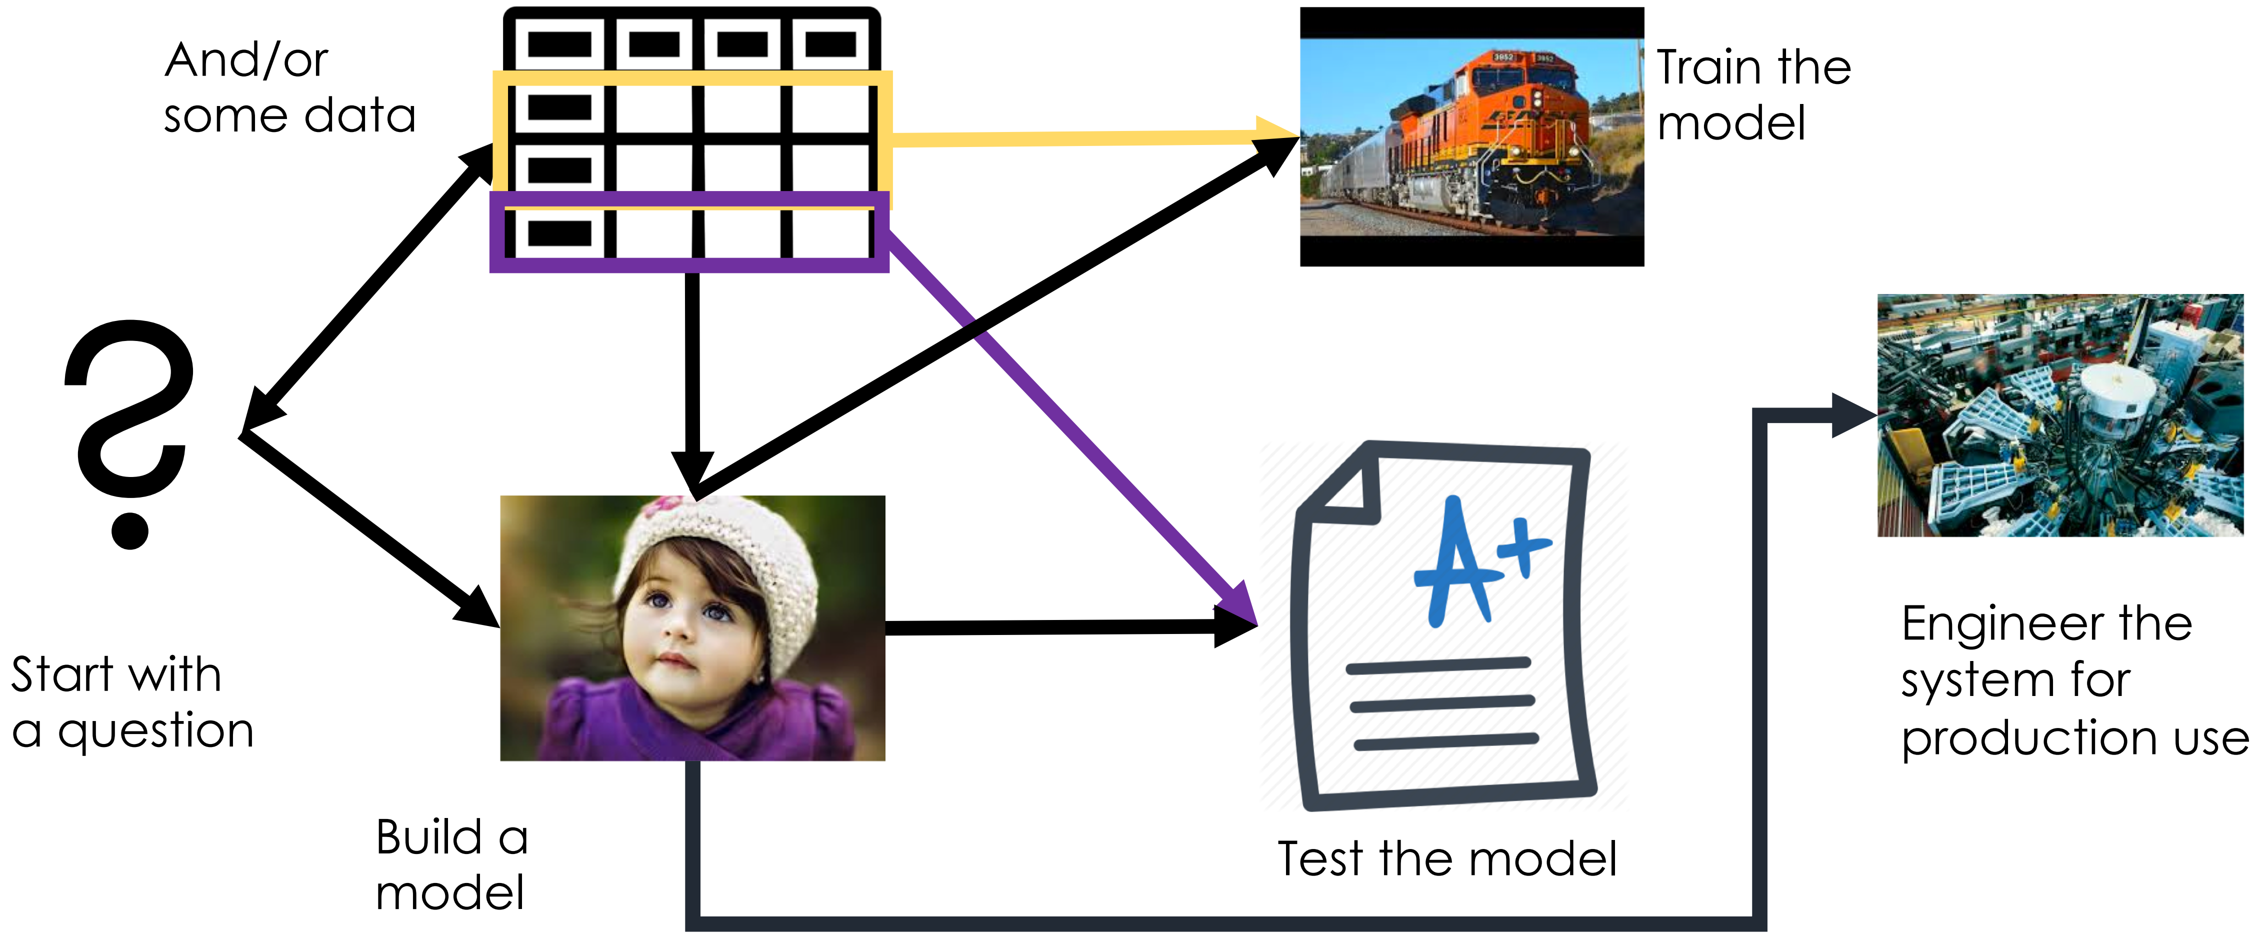

A real view of the ML Pipeline

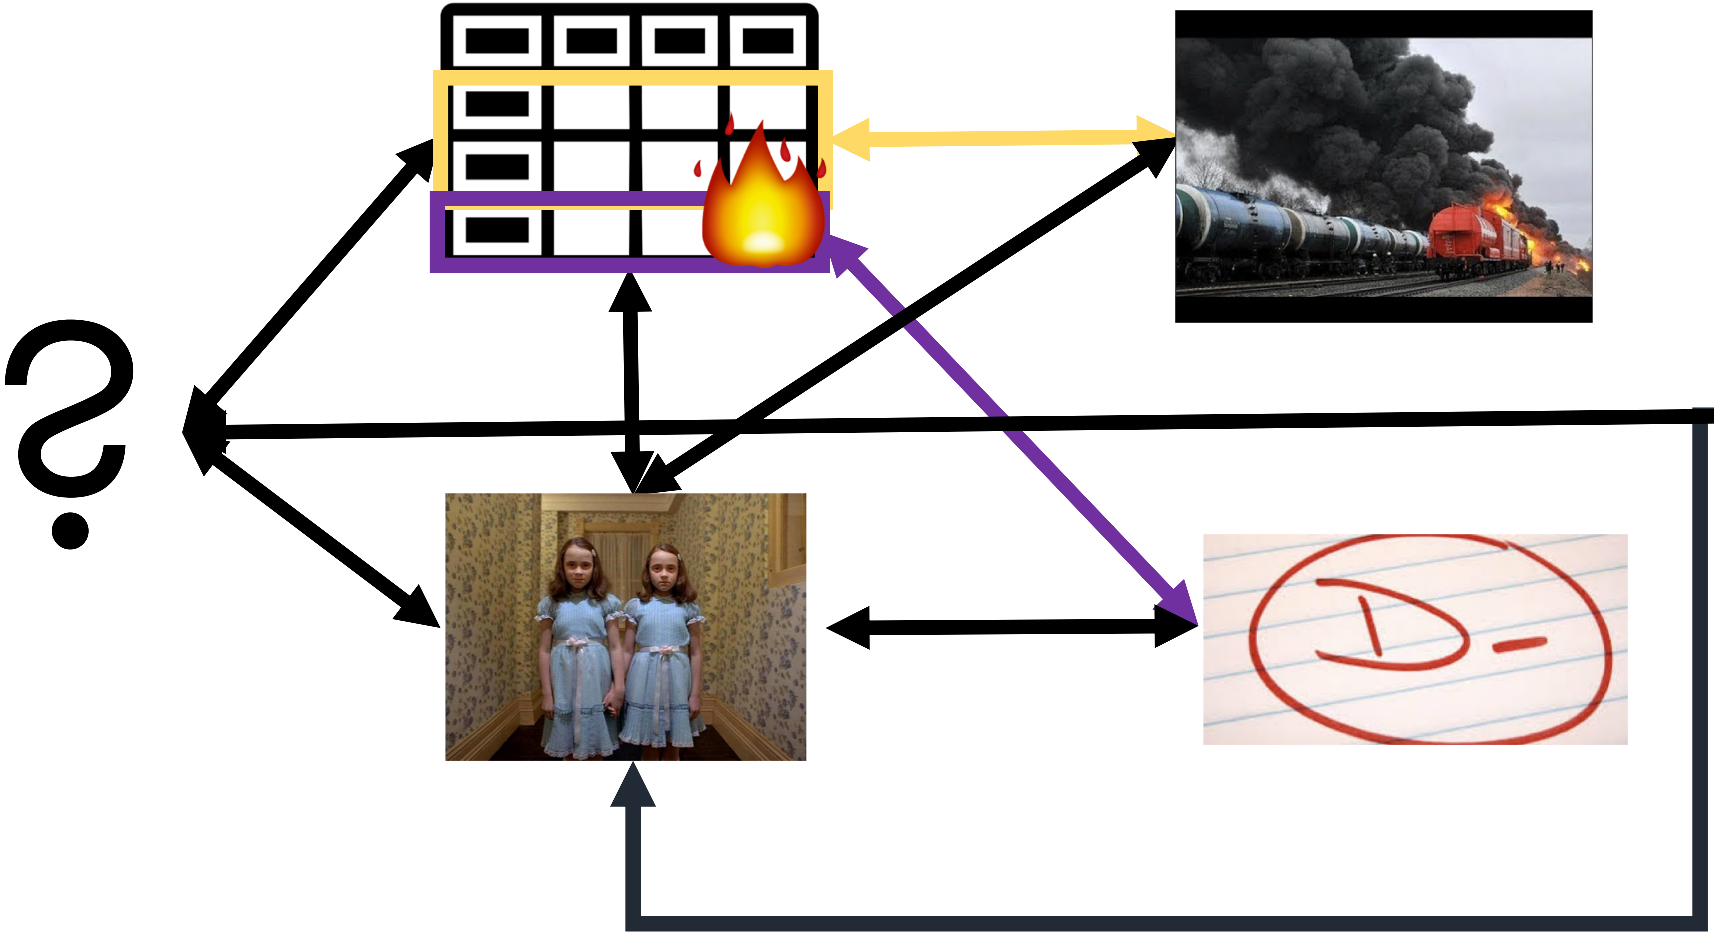

Real world machine learning can be awesome, but there are a lot of gotchas and

Some important considerations

- Question
  - **Who** is asking?
  - What are they seeking to **optimize?**
  - **Why** are they trying to optimize it?
- Data
  - **How** was it collected?
    - Was this influenced by **the algorithm**?
    - By the person who **asked the question?**
  - Does it really **measure** what it claims to?
- Evaluation
  - Do I **believe** the evaluation (e.g. precision/recall)
  - Are they **checking** for the right things?

Case study 1

- Question
  - **Who** is asking?
  - What are they seeking to **optimize?**
- Data
  - **How** was it collected?
    - Was this influenced by **the algorithm**?
    - By the person who **asked the question?**
  - Does it really **measure** what it claims to?

Case study 2

- Data
  - **How** was it collected?
    - Was this influenced by **the algorithm**?
    - By the person who **asked the question?**
  - Does it really **measure** what it claims to?
- Evaluation
  - Are they **checking** for the right things?

What can we do about it?

What **NOT** to do

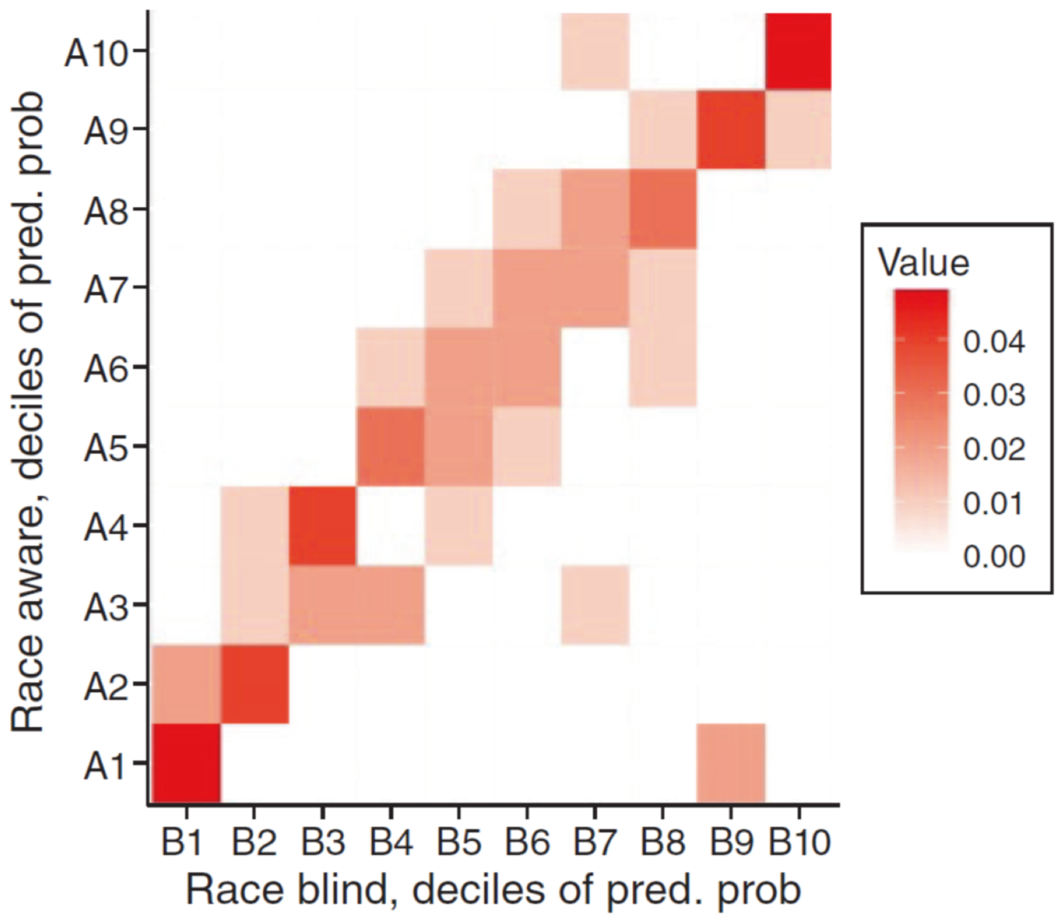

Source: Kleinberg, Ludwig, Mullainathan and Rambachan (2018). Using data from the NELS:88 dataset, we first predict for each observation their predicted college performance (measured as Y=1 if GPA<2.75) using an algorithm that is blinded to applicant race, and then again using an algorithm that has access to each applicant's race. We then take the sample of black students in the NELS:88 and use their predicted values to bin them into deciles based on the race-blind predictions (x-axis) and race-aware predictions (y-axis). If the two models rank-ordered everyone the same way, all the data would be along the 45-degree line. The "off diagonals" in the figure show mis-ranking.

What can we do? Smarter Data

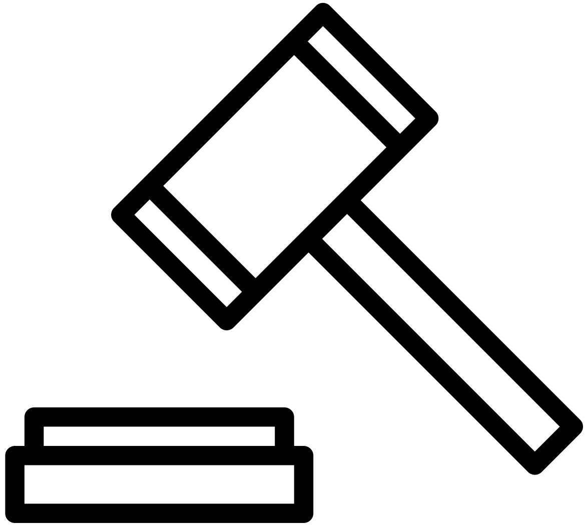

Past decisions from judges

Whether or not people reoffend

What can we do? Smarter Data

may care specifically about violent crimes or about racial inequities. We deal with these problems using different econometric strategies, such as quasi-random assignment of cases to judges. Even accounting for these concerns, our results suggest potentially large welfare gains: one policy simulation shows crime reductions up to 24.7% with no change in jailing rates, or jailing rate reductions up to 41.9% with no increase in crime rates. Moreover, all categories of crime, including violent crimes, show reductions; these gains can be achieved while simultaneously reducing racial disparities. These results suggest that while machine learning can be valuable,

Is that enough?

What can we do? Better Evaluation

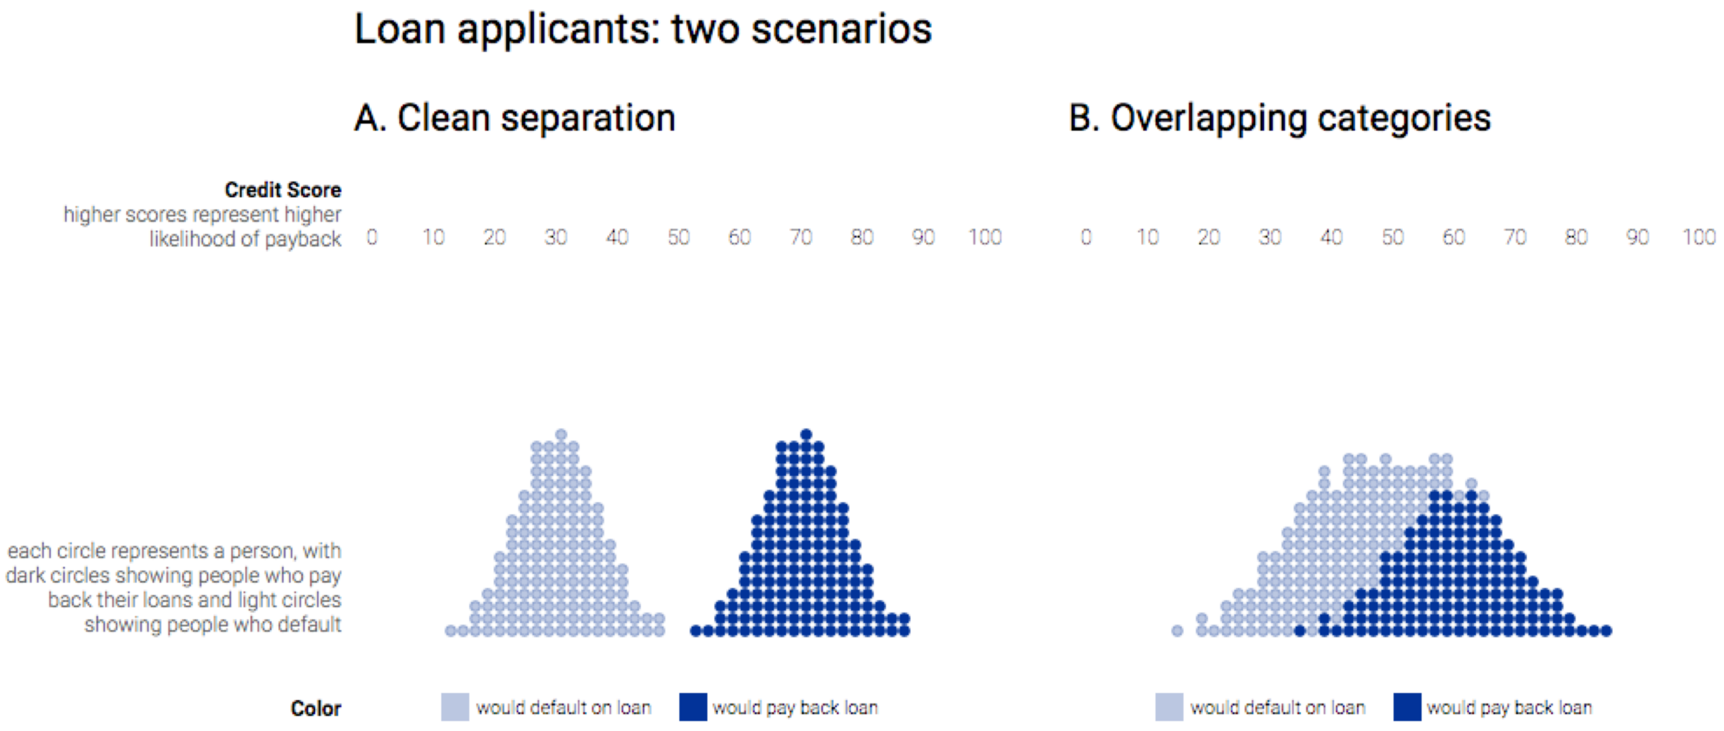

http://research.google.com/bigpicture/attacking-discrimination-in-ml/

What can we do?

Be aware and ask questions

Compute by and with society