Artificial Intelligence

CS 10

Ezinne Nwankwo
ezinne_nwankwo@berkeley.edu


The New York Times
SUBSCRIBE NOW
Account
A.I. Is Helping Scientists Understand an Ocean's Worth of Data
Machine-learning applications are proving to be especially useful to the scientific community studying the planet's largest bodies of water.

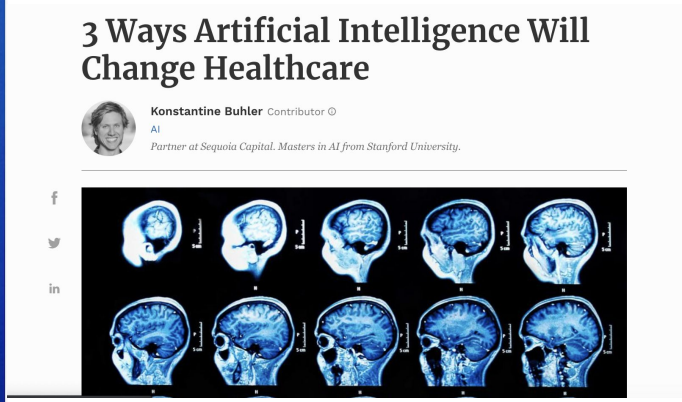
3 Ways Artificial Intelligence Will Change Healthcare
Konstantine Buhler Contributor
AI
Partner at Sequoia Capital. Masters in AI from Stanford University.


BUSINESS INSIDER
Log In
Subscribe
HOME > TECH
Farmers are using AI to spot pests and catch diseases — and many believe it's the future of agriculture
Daniel T. Allen Nov 8, 2019, 11:44 AM

Forbes

4,646 views | Mar 2, 2020, 04:19am EST

AI Bias Could Put Women's Lives At Risk – A Challenge For Regulators

MIT Technology Review

Artificial intelligence

Predictive policing algorithms are racist. They need to be dismantled.

Amazon Is Pushing Facial Technology That a Study Says Could Be Biased

In new tests, Amazon's system had more difficulty identifying the gender of female and darker-skinned faces than similar services from IBM and Microsoft.

WIRED

An Algorithm That 'Predicts' Criminality Based on a Face Sparks a Furor

Its creators said they could use facial analysis to determine if someone would become a criminal. Critics said the work recalled debunked "race science."

Goal of this lecture

I. Gain intuition for what AI can do
II. Understand some challenges in AI and introduce you to questions I think about all day
III. Pathways to AI @ Berkeley and In General

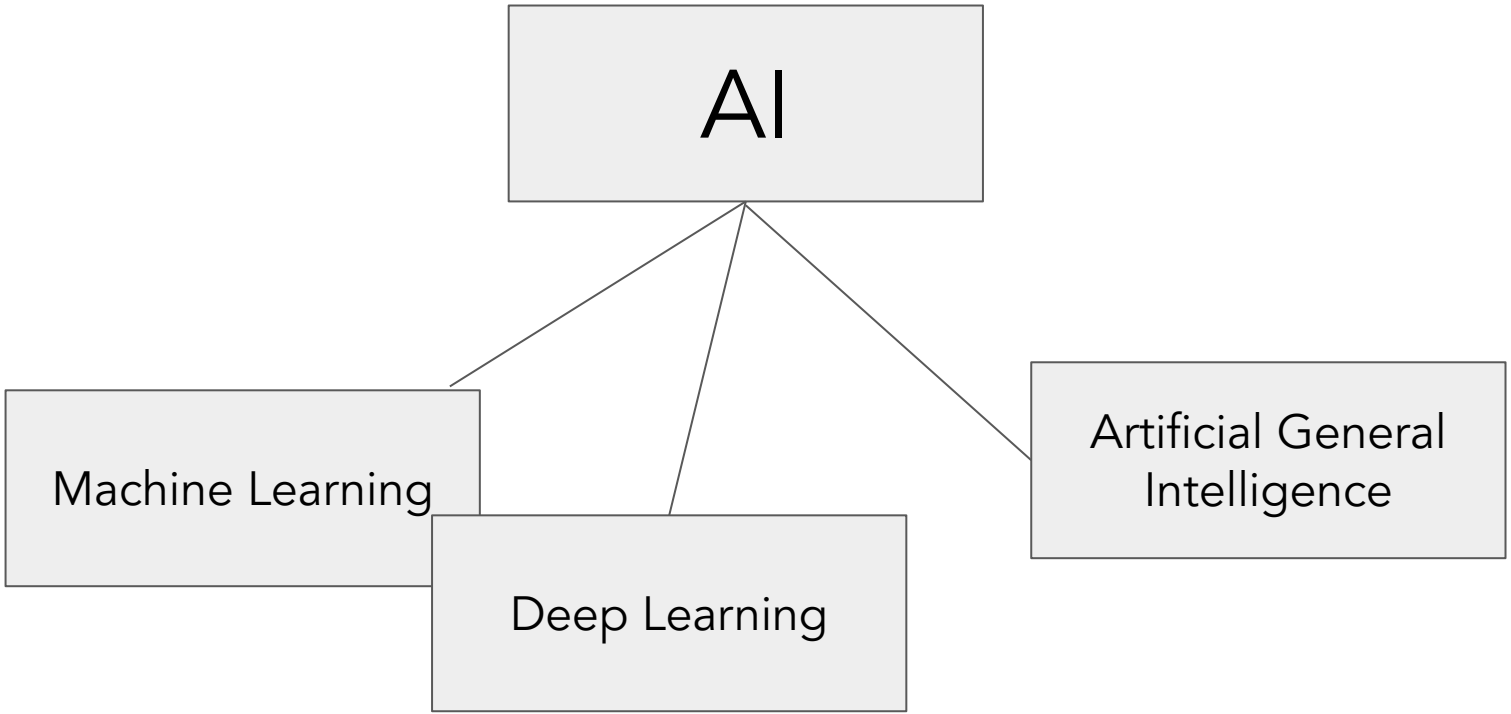

Probabilistic/Statistical Learning + Data + Algorithms

What is AI? What is Machine Learning?

The study of prediction from examples

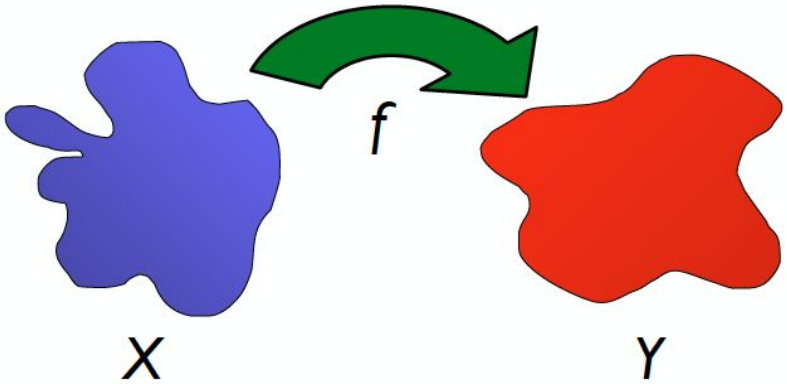

Estimate y = f(x) from observations $(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)$

Pattern Recognition

The study of prediction from examples

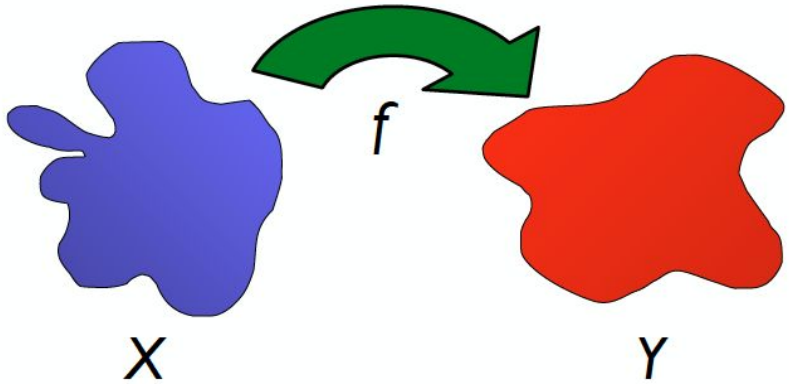

Estimate y = f(x) from observations $(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)$

Hope that this works well on new examples

Supervised Learning

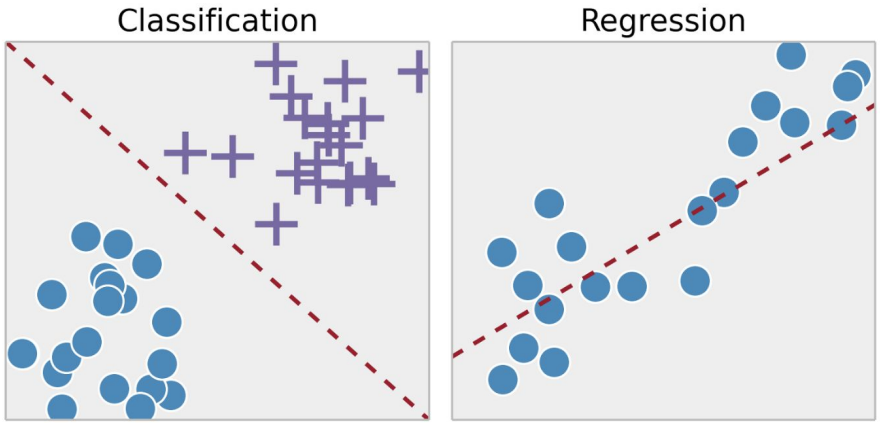

Unsupervised Learning

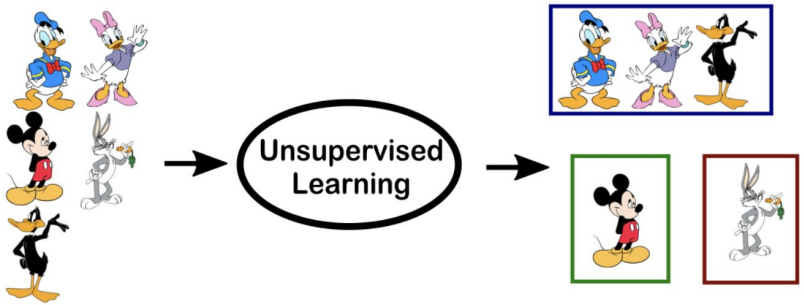

Images taken from https://towardsdatascience.com/supervised-vs-unsupervised-learning-14f68e32ea8d

The study of prediction from examples

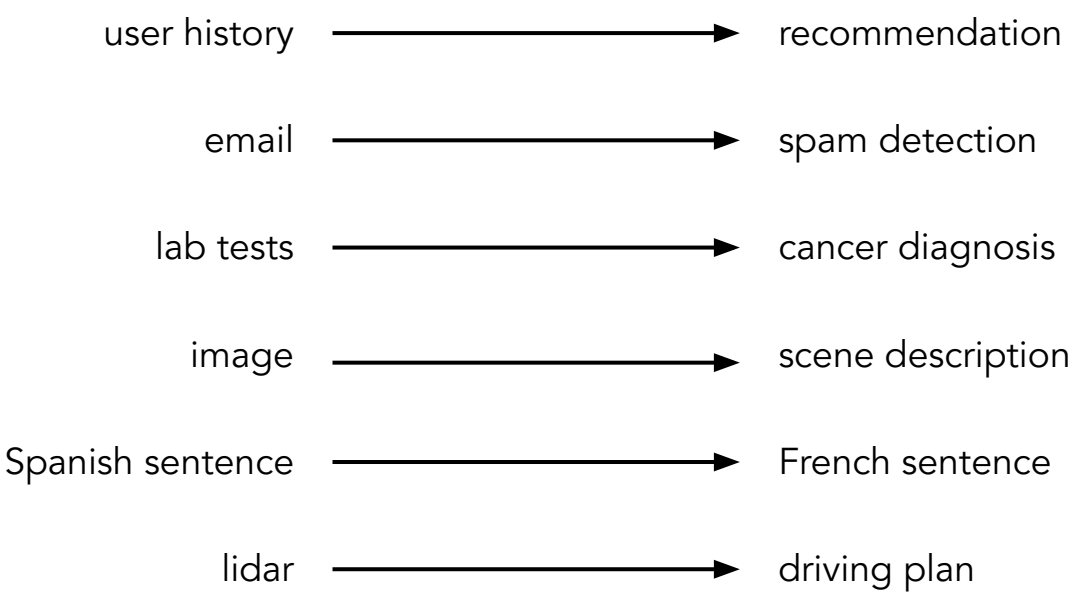

1958

NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo of Computer Designed to Read and Grow Wiser

WASHINGTON, July 7 (UPI) —The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The Perceptron Algorithm

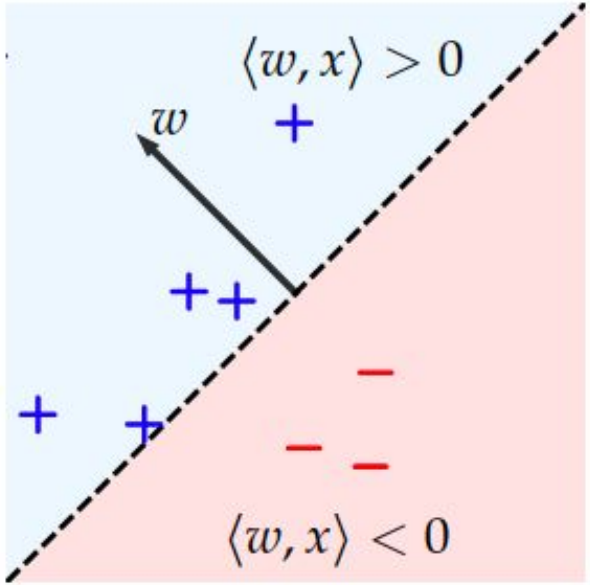

$$w_{t+1} = w_t - \eta \nabla_{w_t} loss(f_{w_t}(x_i), y_i)$$

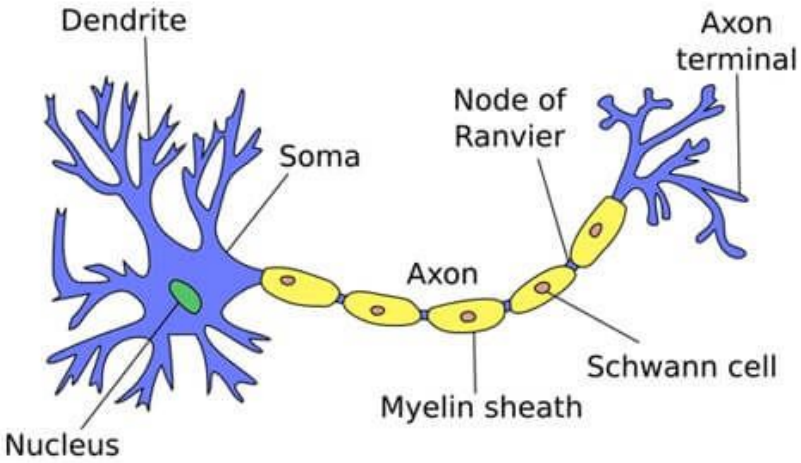

The Brain

Neural Networks

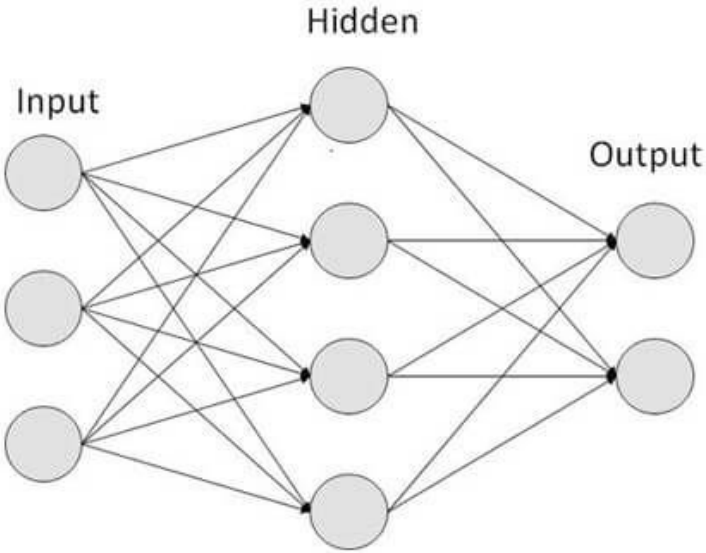

Computer Vision

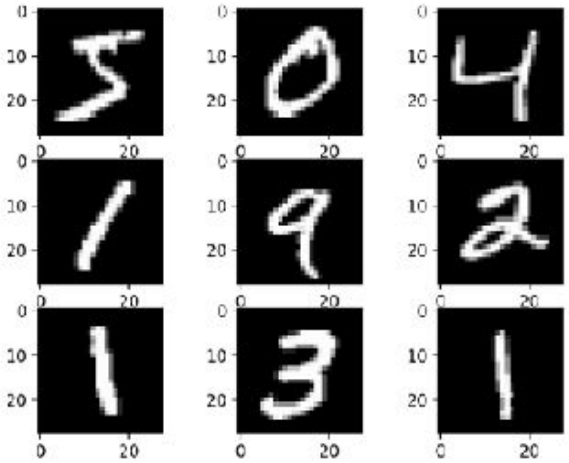

IMAGENET

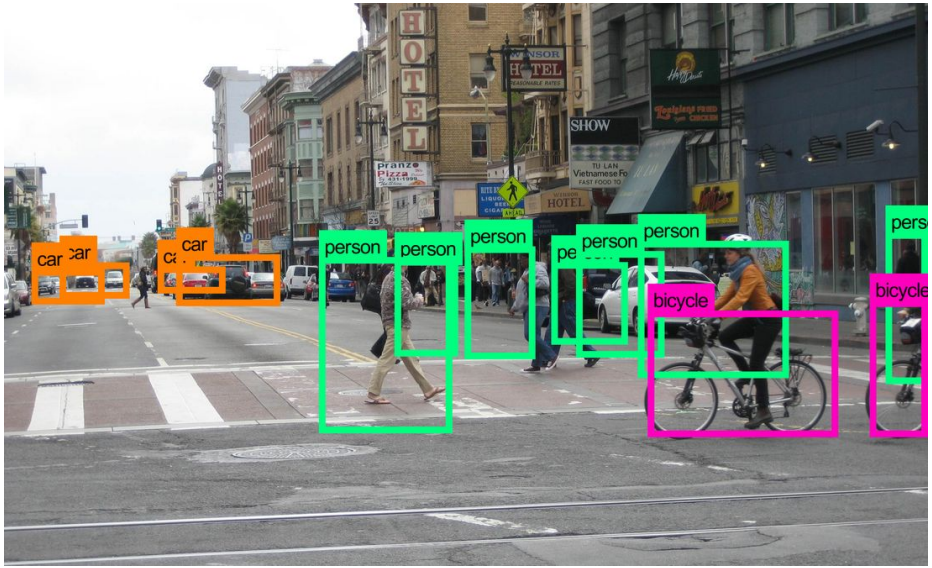

Object Detection/Identification/Segmentation

Computer Vision

(Chan et. al., 2018)

Computer Vision - Poultry Disease Detection

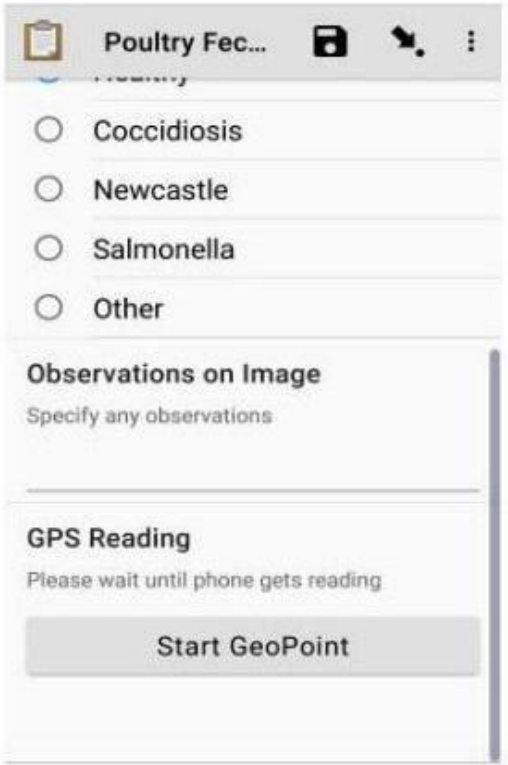

Natural Language Processing

Machine Translation

Language Understanding

NLP - Misinformation and Online Harm

Misinformation in Tweets

#Nigerian reports they have found a possible #vaccine #FightCovid19Africa #coronavirus #pandemic #stopcovid19 #Africa

Nigerian researchers announce COVID-19 vaccine

Nigerian Universities' Scientists, under the aegis of COVID-19 Research Group, on Friday announced the discovery of a vaccine for ...

citinewsroom.com

7 31

Pauline.
@kot_hacker

It's time for Africa to turn off their stupid Full Tunnel VPN because we going to need the bandwidth, there is no way they will allow Corona virus not to infect us. Hoping our Melanin is our immunity.
#FridayMotivation

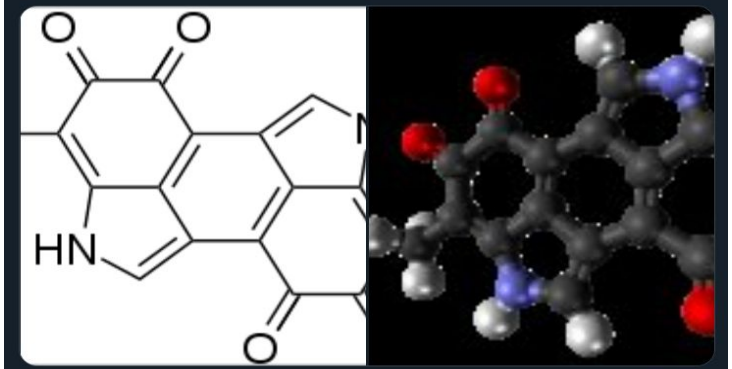

1:13 AM · Feb 28, 2020 · Twitter for Android

Masakhane

A grassroots NLP community for Africa, by Africans

Our Mission

Masakhane is a grassroots organisation whose mission is to strengthen and spur NLP research in African languages, for Africans, by Africans. Despite the fact that 2000 of the world's languages are African, African languages are barely represented in technology. The tragic past of colonialism has been devastating for African languages in terms of their support, preservation and integration. This has resulted in technological space that does not understand our names, our cultures, our places, our history.

Masakhane roughly translates to "We build together" in isiZulu. Our goal is for Africans to shape and own these technological advances towards human dignity, well-being and equity, through inclusive community building, open participatory research and multidisciplinarity.

Tweets from @MasakhaneNLP Follow

Masakhane Retweeted

David Ìfẹ́olúwa Adélaní @davlan... · Sep 30

Apply for Volunteer/Scholarship at AKBC 2022 even if you have no paper there @MasakhaneNLP @black_in_ai

6

Masakhane Retweeted

Salomey Osei @NanaYaaSally · Sep 29

Please don't forget to apply for the @black_in_ai funding for attending this year's workshop co-

Machine learning models are deciding or informing consequential decisions about people.

Data (X)

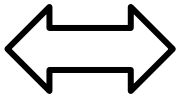

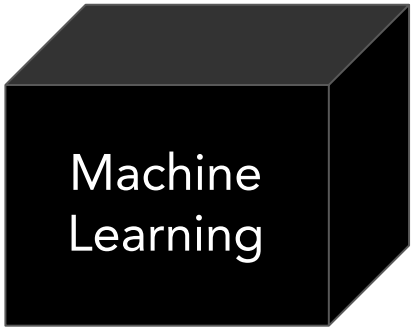

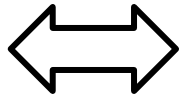

P[recidivism|X]

P[reamission|X]

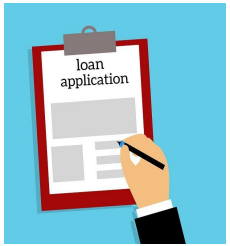

P[loan default|X]

Challenges in AI - Computer Vision

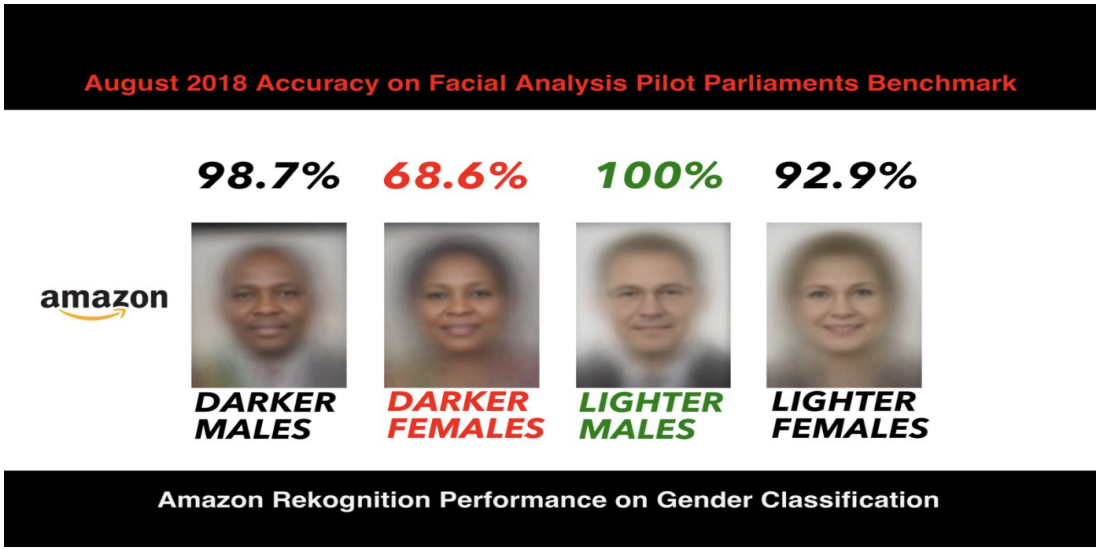

Facial Recognition

(Buolamwini et. al., 2019)

Challenges in AI - Crime

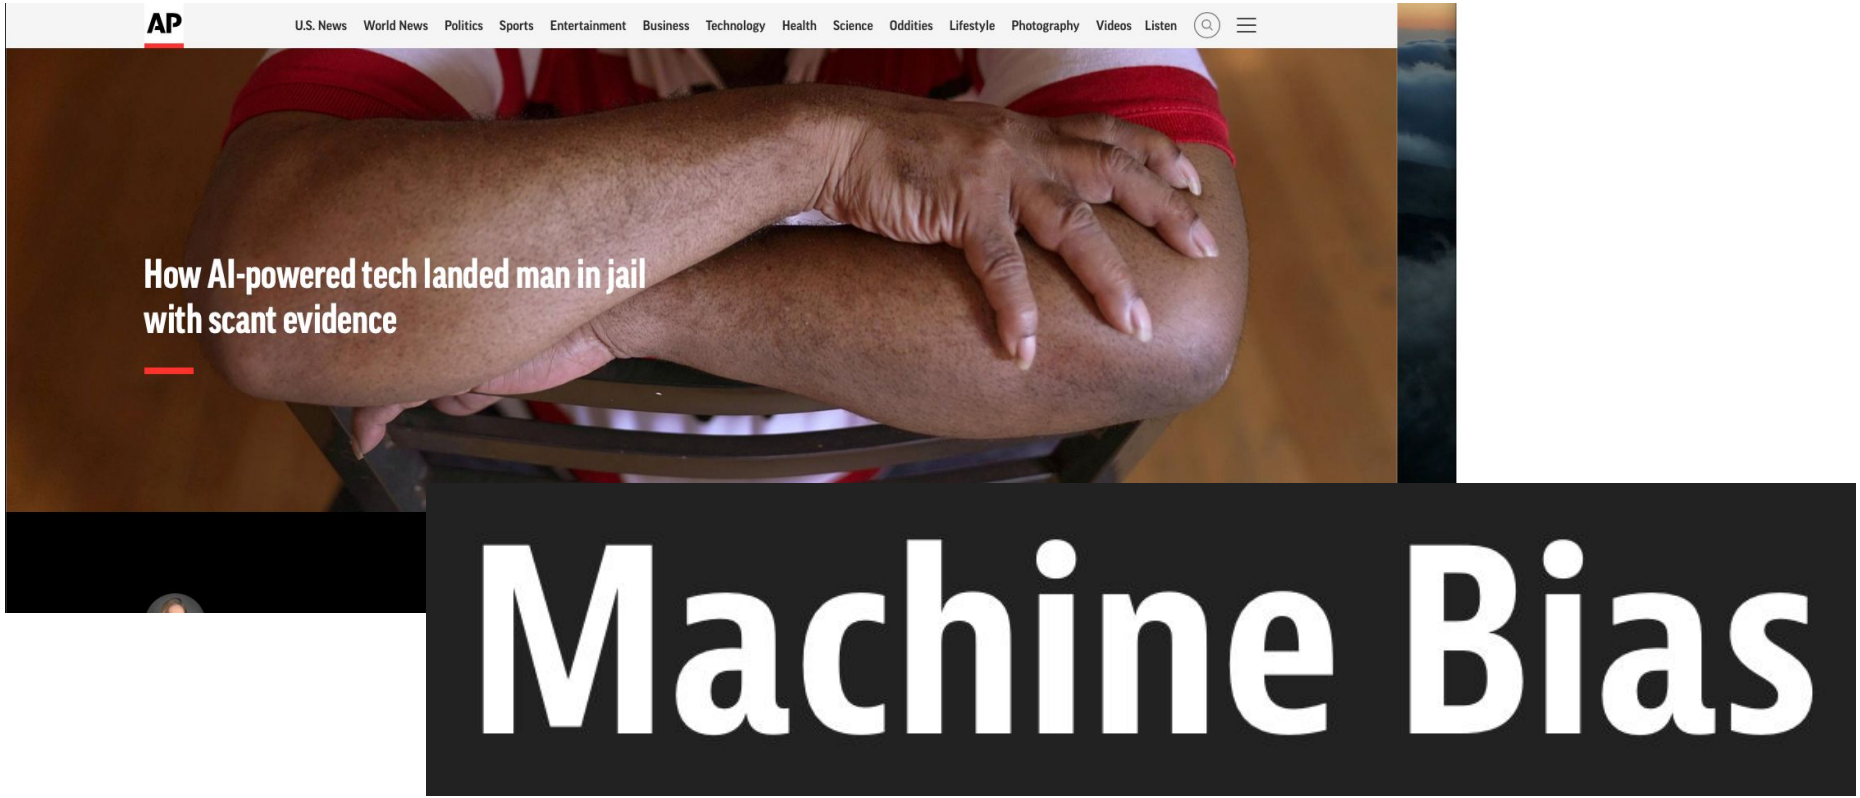

Challenges in AI - Social Work

A CHILD ABUSE PREDICTION MODEL FAILS POOR FAMILIES

"We definitely oversample the poor," says Erin Dalton, Director of Allegheny County's Office of Data Analysis, Research and Evaluation. "All of the data systems we have are biased. We still think this data can be helpful in protecting kids." NKBIMAGES/GETTY IMAGES

Wired (2017)

Challenges in AI - Health

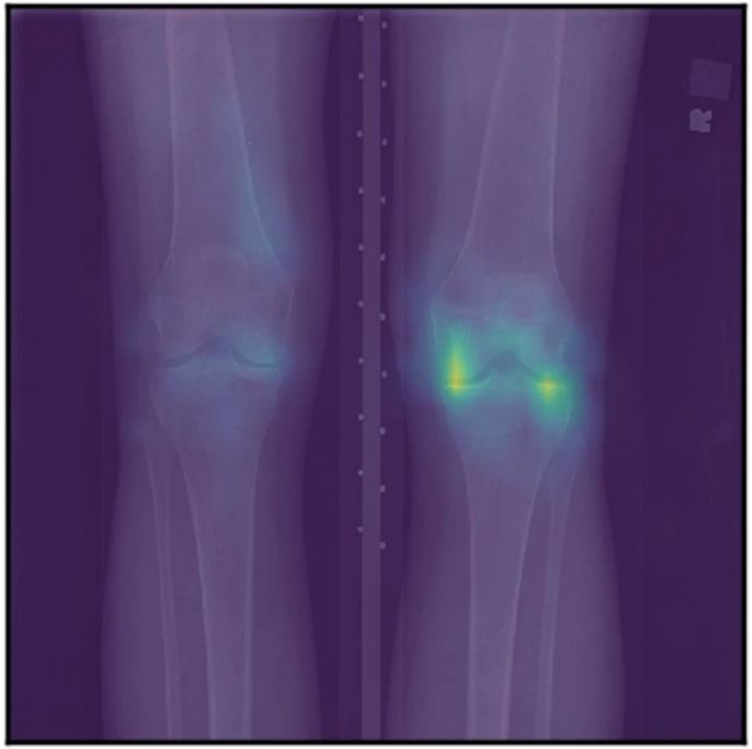

The model's prediction target is the pain score in the knee appearing on the right side of the image. Regions that influence the prediction more strongly are shown in brighter colors.

Pierson et al. An algorithmic approach to reducing unexplained pain disparities in underserved populations. Nature (2020).

Survey of Fairness in ML

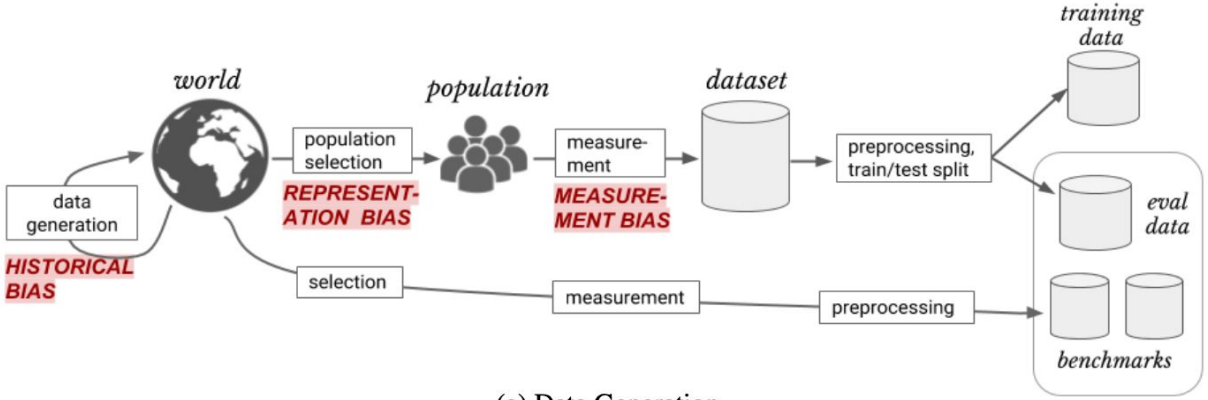

(a) Data Generation

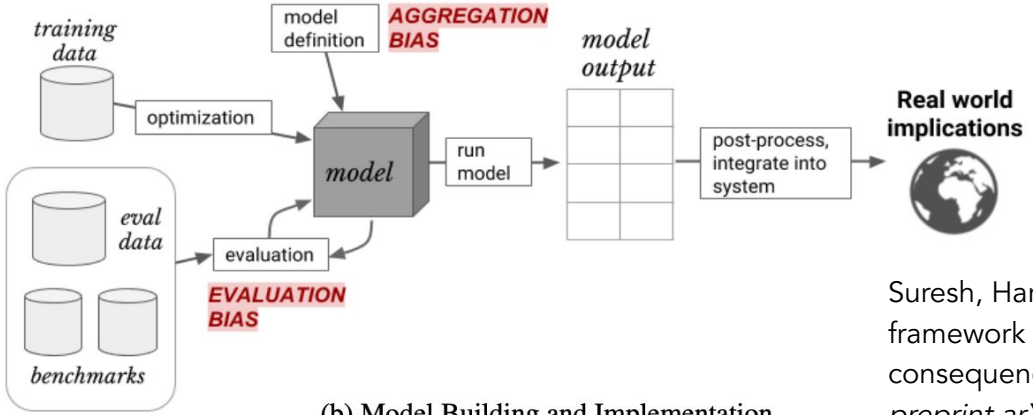

(b) Model Building and Implementation

Suresh, Harini, and John V. Guttag. "A framework for understanding unintended consequences of machine learning." *arXiv preprint arXiv:1901.10002* (2019).

For a long time, we believed...

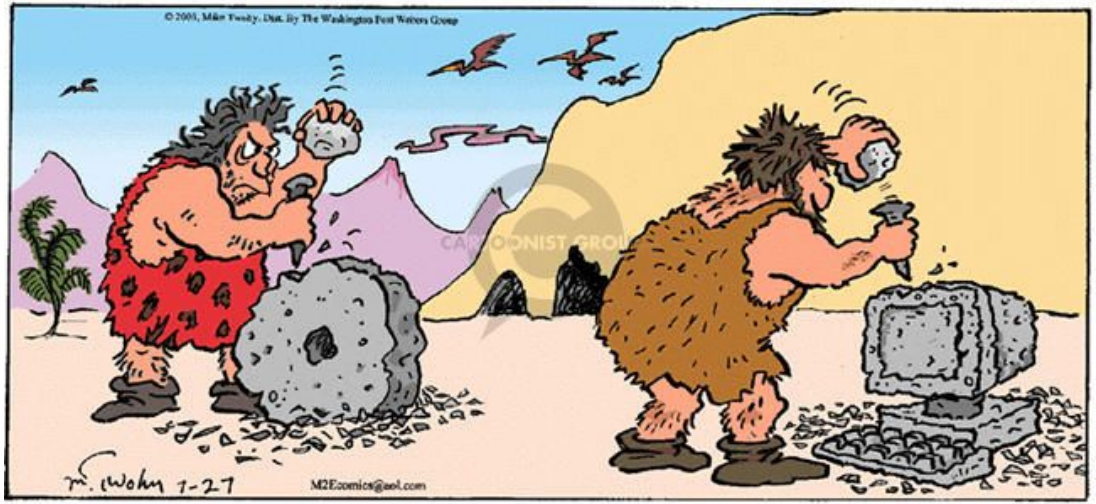

- Data was unbiased, objective, logical...
- So were models and algorithms.

(1) Sanmi Koyejo (2) https://www.thecomicstrips.com/subject/The-Stone+Age-Comic-Strips.php

The role of data in AI…

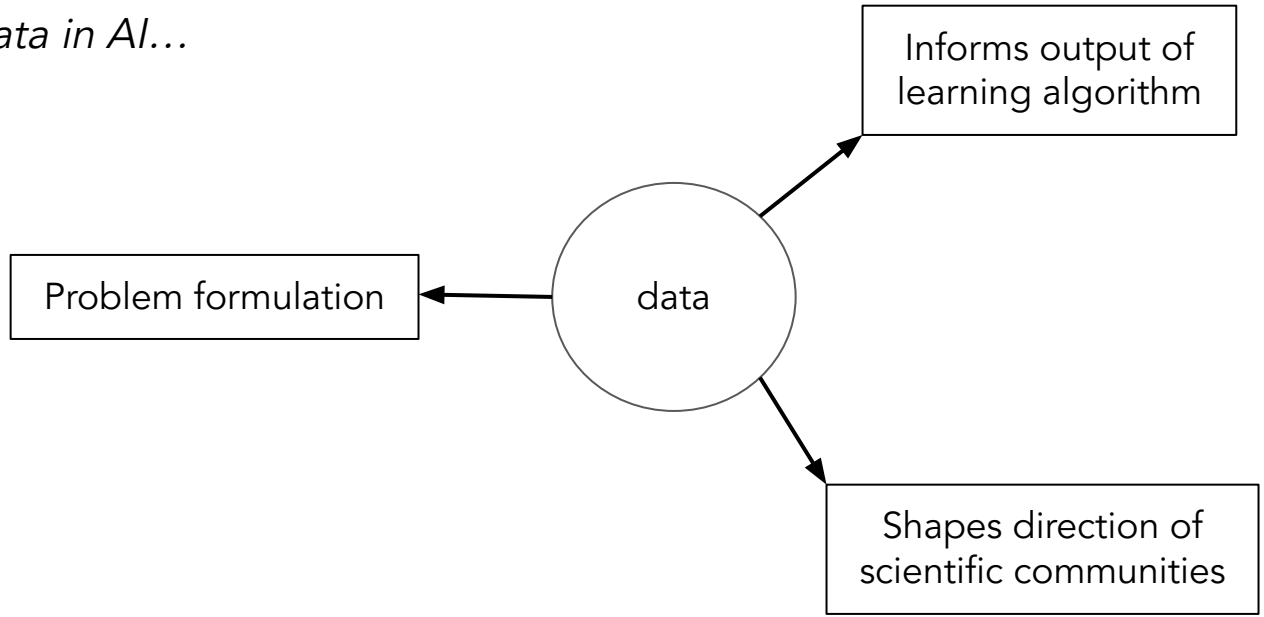

So what makes for good data and for what purpose (i.e. predicting life outcomes)?

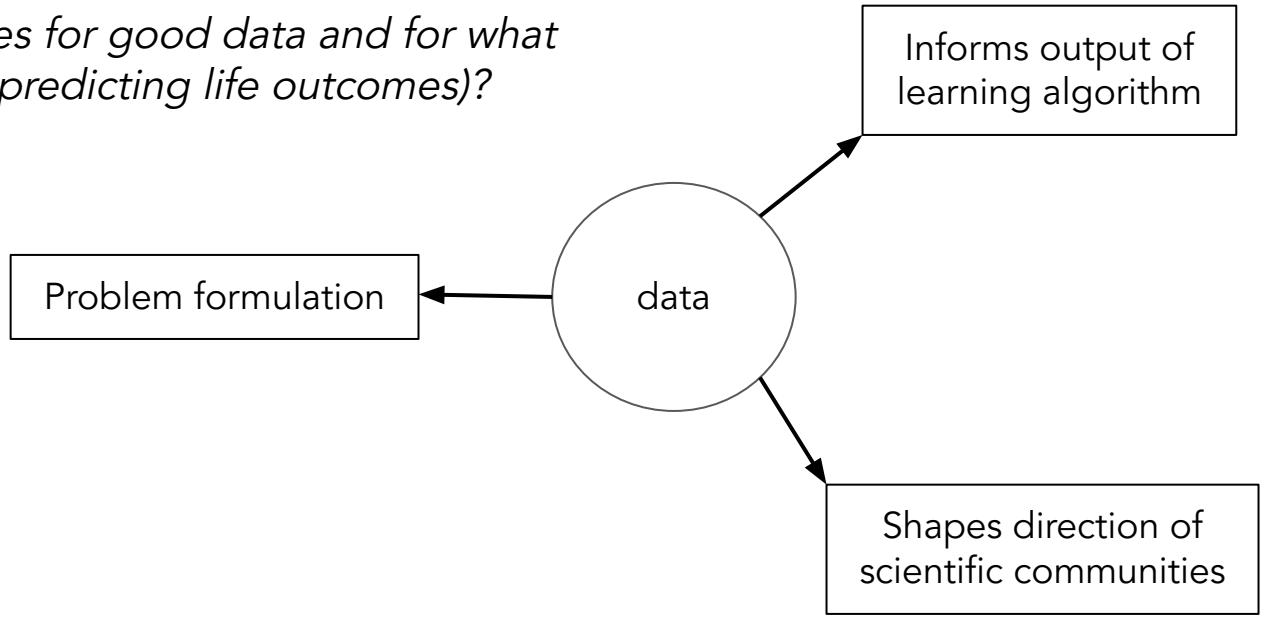

Limits of Prediction

Measuring the Predictability of Life Outcomes with A Scientific Mass Collaboration

Matthew Salganik, Ian Lundberg,
Alexander T. Kindle, et al.

Limits of Prediction

1.) Sensitive dependence on inputs
2.) Shocks
3.) Accumulation and amplification of advantage
4.) Unobserved and observed inputs
5.) The 8 Billion Problem
6.) The Drift
7.) Ill-conceived target variable
8.) Self-equilibration and strategic behavior
9.) Forward and inverse prediction

Narayanan and Salganik (2020). Limits of prediction: pre-read. Princeton University.

Forward and inverse prediction

Given the output from the data generating process, can we predict the input?

$$Y \leftarrow X$$

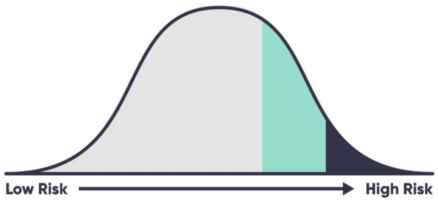

Race
Gender
Age

Self-equilibration and strategic behavior

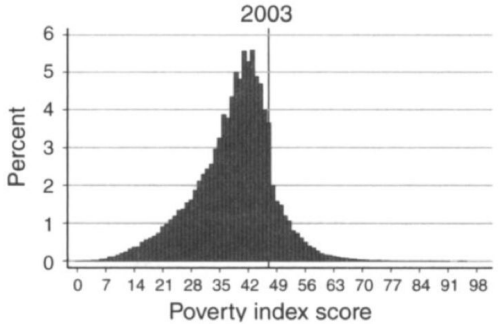

Camacho and Conover (2011). Manipulation of Social Program Eligibility.

Ill-conceived target variable

Some target variables are hard to measure or unstable.

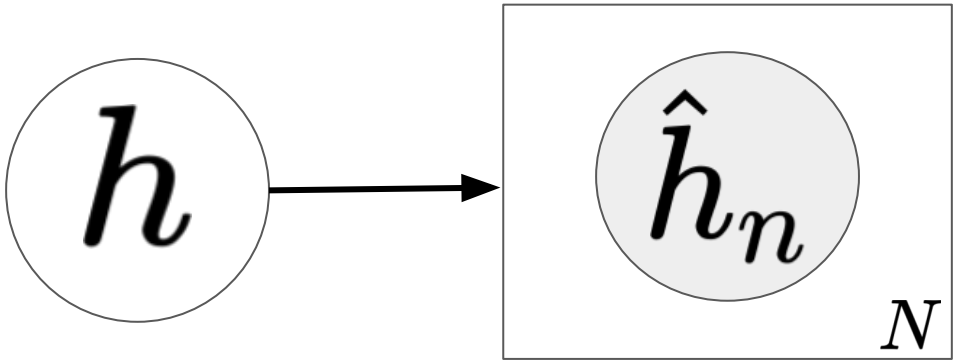

$$\hat{h}_n = h + \epsilon_n$$

Jacobs and Wallah (2020)
Measurement and Fairness

Modern systems bring new challenges

- Automated
- Dynamic
- Heavily dependent on multi-agent behavior
- Morally, politically, and mathematically challenging

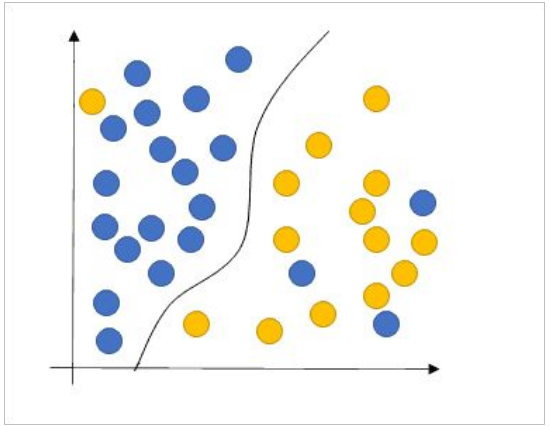

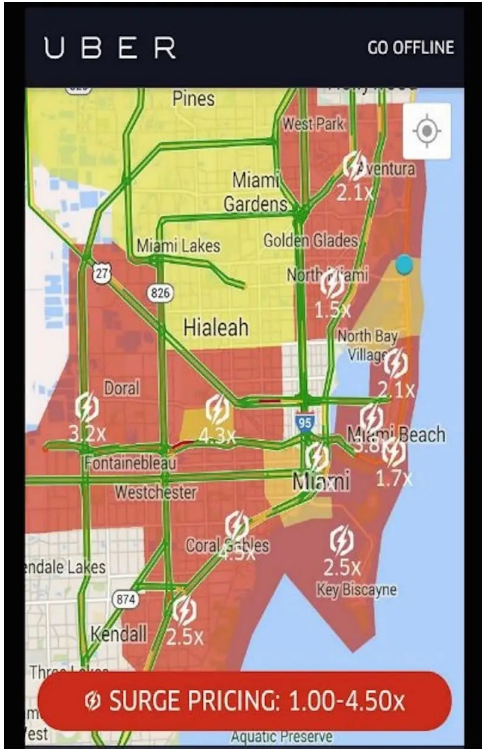

What Next?

- Introductory Courses to Consider
  - Stats: CS/Stat/Info C8, Stat 102
  - Programming: CS 61ABC
  - Probability: CS 70
  - Linear Algebra: EE 16A, EE 16B, Math 54
  - Optimization: EE 227A
  - Calculus: Math 1AB, Math 53
- Upper Divisions to Consider
  - AI-ML Essentials: CS 188, CS 189
  - Neural Networks: CS 182
  - CV: C280, RL: CS 285, NLP: 288
- Take Courses Outside CS
- Outside the classroom
  - Start personal projects!!
  - Get involved in research!!
  - Get creative with internships!!
  - Get involved in orgs: NSBE, Black in AI, Women in AI

My Path Thus Far

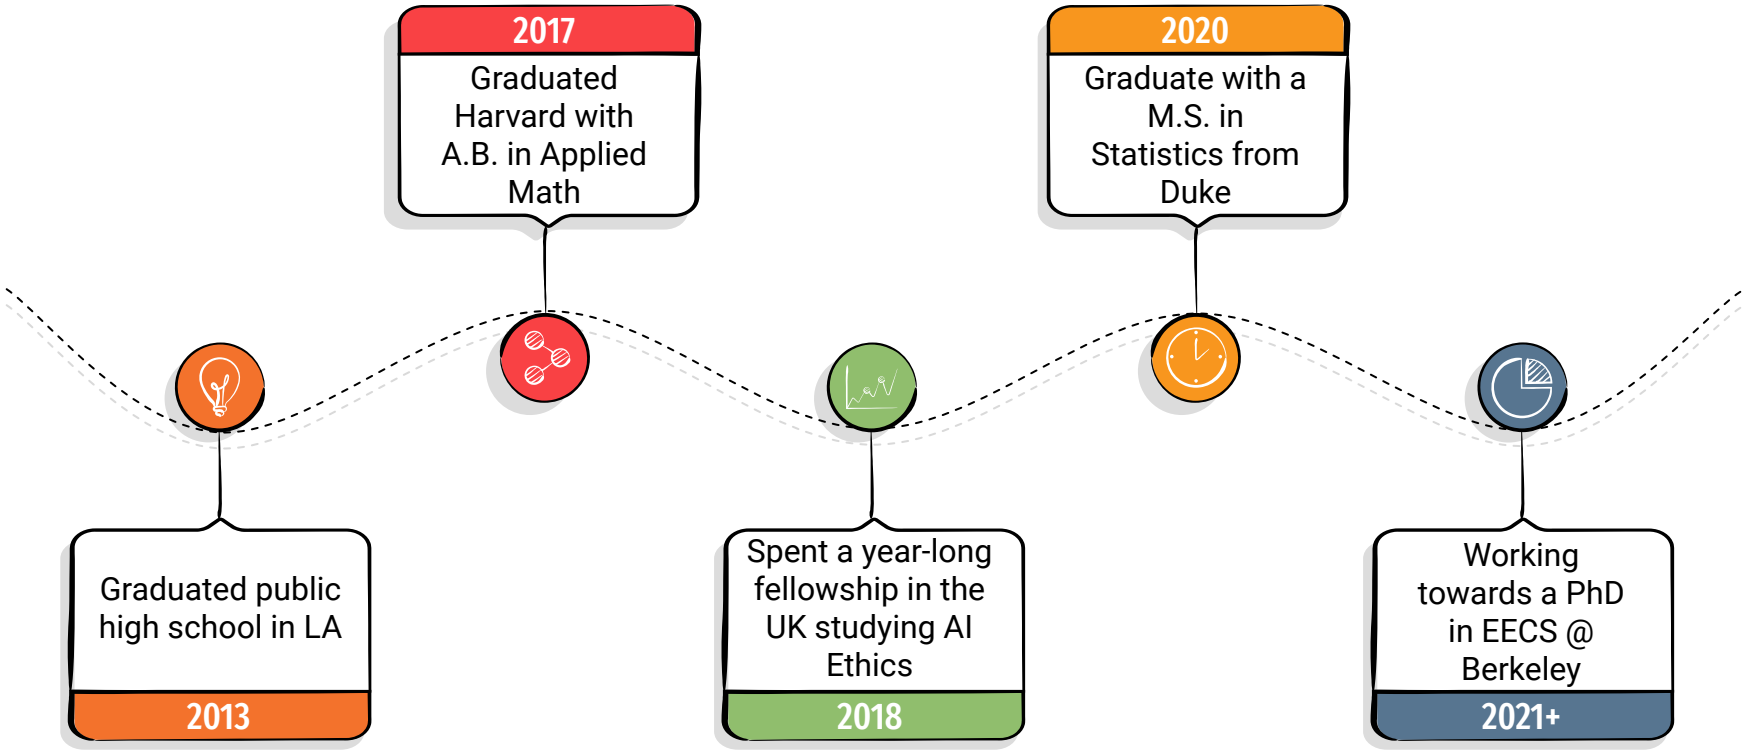

My Path Thus Far

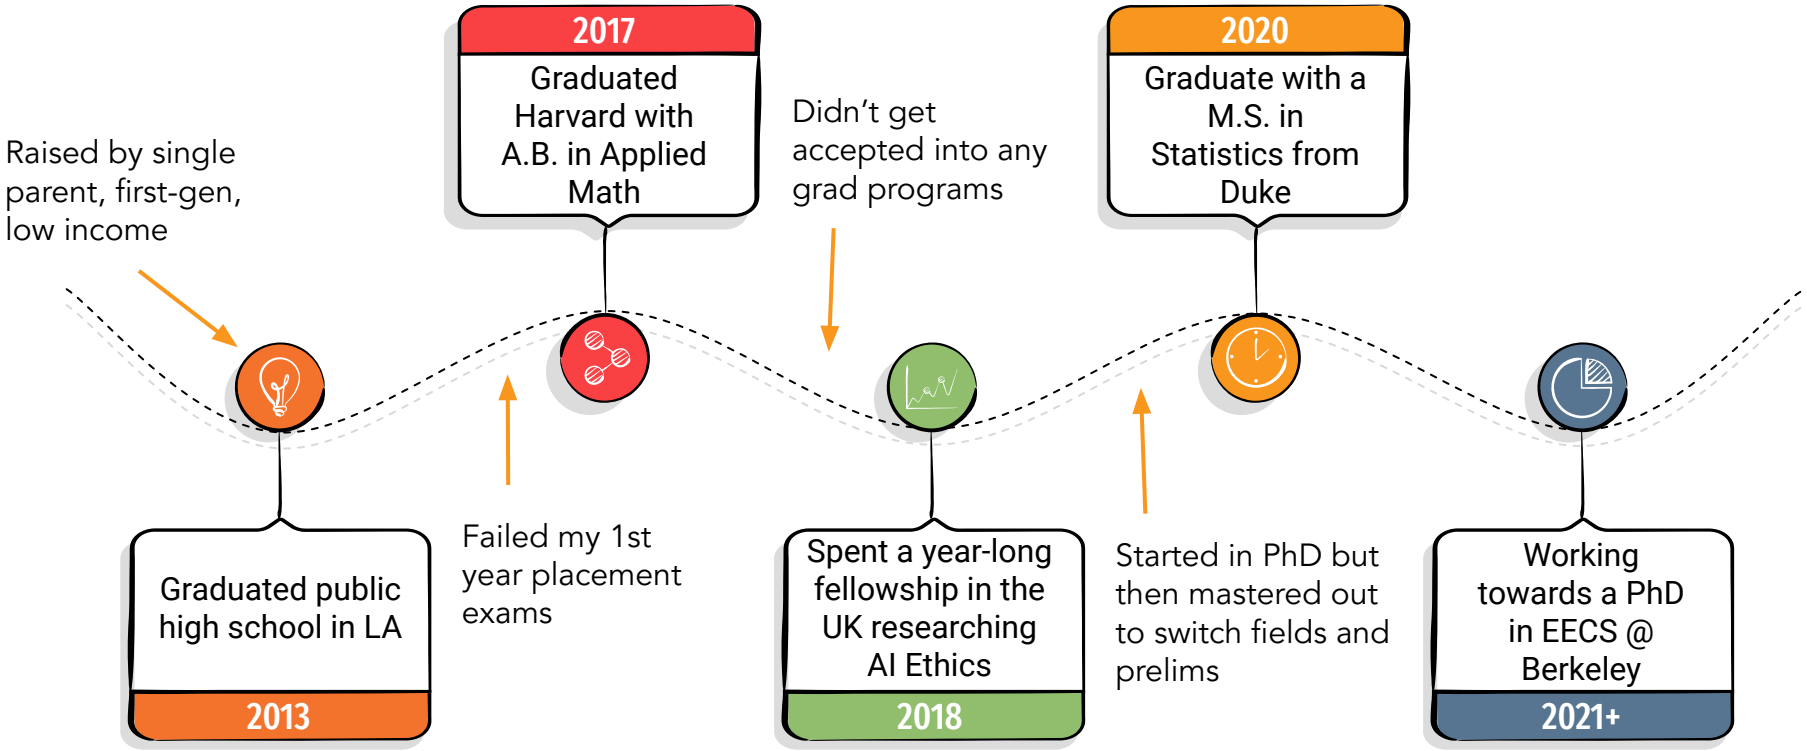

Thank you!

Any Questions?

Ezinne Nwankwo
ezinne_nwankwo@berkeley.edu