TEACHING DATA SCIENCE ETHICS

SURESH VENKATASUBRAMANIAN
UNIVERSITY OF UTAH

THE PROMISE OF DATA SCIENCE...

Spotify

NETFLIX

amazon

THE UNIVERSITY OF UTAH

MIRADOR

HireVue

banjo

Palantir

… AND THE PERIL

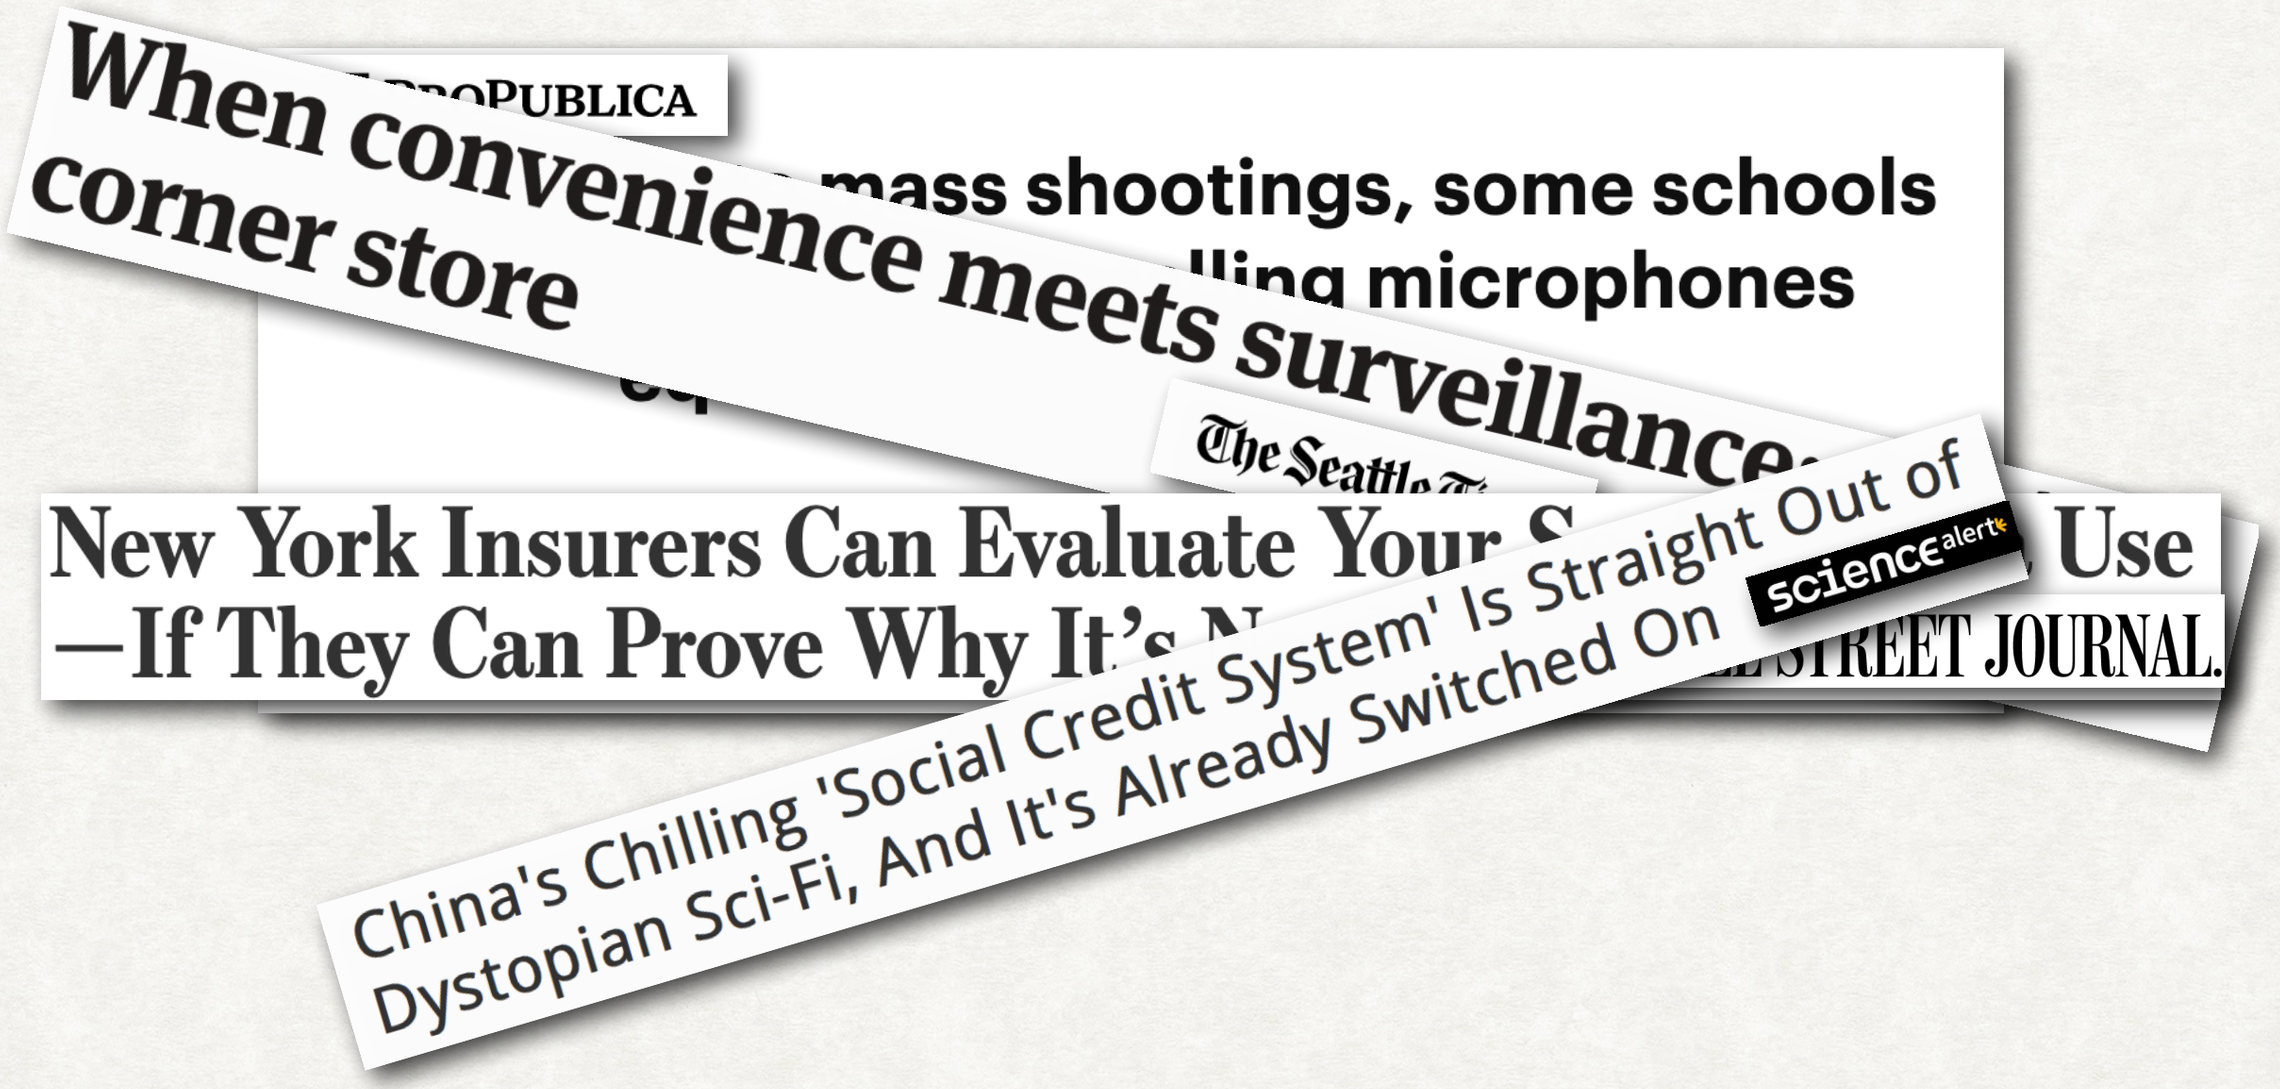

MISTRUST & INEQUALITY REFRACTED THROUGH COVID19

Contact-tracing apps have a trust problem, even if they do protect your privacy

Experts believe tha
contact-tracing ap
convincing the pub
big tech.

THE FIRST 100

COVID-19 Took Black Lives First. It Didn't Have To.

ProPublica, May 9, 2020

THE CORONAVIRUS CRISIS

The Coronavirus Doesn't Discriminate, But U.S. Health Care Showing Familiar Biases

April 2, 2020 · 12:37 PM ET

NPR, Apr 2, 2020

DATA SCIENTISTS NEED ETHICS TRAINING

Do We Need To Teach Ethics And Empathy To Data Scientists?

Kalev Leetaru Contributor ⓘ

AI & Big Data

I write about the broad intersection of data and society.

Forbes, Oct 18, 2018

Ethics I

APRIL 17, 2018

Data scientists
curricula. Rec
discuss the pr
classroom. Th

Responsible Computer Science Challenge

With Great Code Comes Great Responsibility

a partnership of

OMIDYAR NETWORK moz://a SCHMIDT FUTURES Craig Newmark Philanthropies

WHAT SHOULD WE TEACH?

ETHICS IN DATA SCIENCE (S19,F19)

In this course, we will explore the moral, social, and ethical ramifications of the choices we make at the different stages of the data analysis pipeline, from data collection and storage to understanding feedback loops in analysis. Through class discussions, case studies and exercises, students will learn the basics of ethical thinking in science, understand the history of ethical dilemmas in scientific work, and study the distinct challenges associated with ethics in modern data science.

and 200+ more courses like this!
https://medium.com/@cfiesler/tech-ethics-curricula-a-collection-of-syllabi-3eedfb76be18

GOAL: MORAL RESPONSIBILITY

Broadly, understand the difference between responsibility and blame and how that influences ethical thinking. Emphasize that as data scientists we have a role to play.

Responsibility comes from agency, which comes from power and choice:

Learn to be aware of the (ethical) choices at our disposal when we design systems for data analysis

And we should be consistent:

Be able to formulate and articulate actionable outcomes that are consistent with an ethical analysis of a task.

GOAL: EPISTEMIC VALIDITY

"Your ethics are only as good as your facts"

- Data science is an attempt to derive knowledge from data.
- What kind of knowledge do we derive?
- Are we sure we've derived it in ways that are accurate and valid?
- Do we use it in ways consistent with its nature?

GOAL: EMPOWERMENT

"All of this has happened before"

- Understand how science has dealt with ethical issues before
- Understand the mechanisms — individual and collective — by which we can place checks and safeguards on how we use data in society.

*Learn to face and deal with **moral distress**: helplessness in the face of hierarchy and power.*

HOW SHOULD WE TEACH?

“IN THE BEGINNING, ARISTOTLE…”

- Most ethics 101 classes build a foundation on ethical frameworks — virtue ethics, deontological ethics, consequentialism, contract theory, and so on….

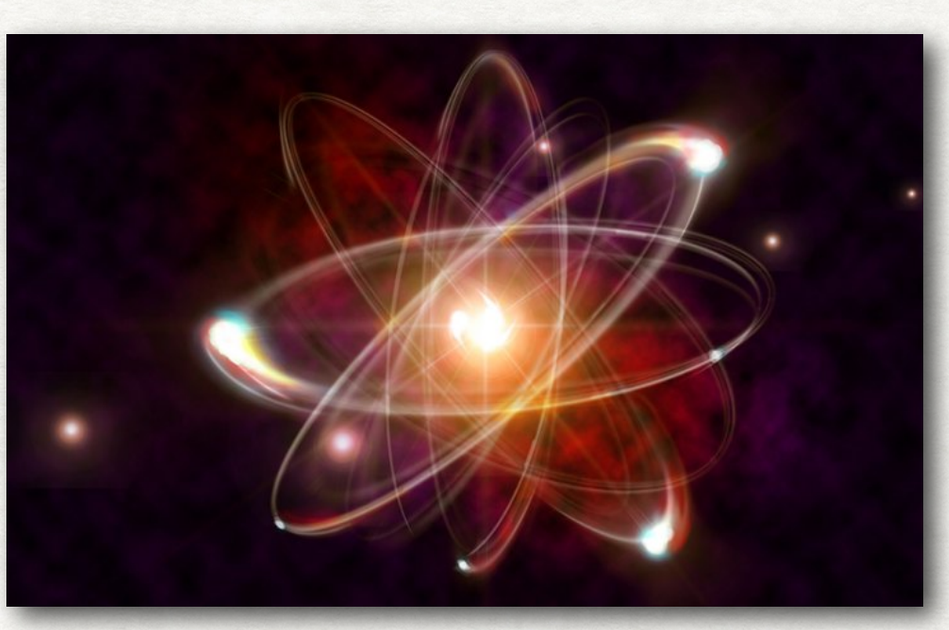

https://www.wonderopolis.org/_img?img=/wp-content/uploads/2017/03/Quantum_Physicsdreamstime_xxl_60222747_1.jpg&transform=resizeCrop,720,450

Can’t use quantum physics to build a bridge!

APPLIED ETHICS FOR DATA SCIENCE

- Can we ground the discussion in the processes of data science (like with business ethics, public health ethics, bioethics and other applied ethics frameworks)?

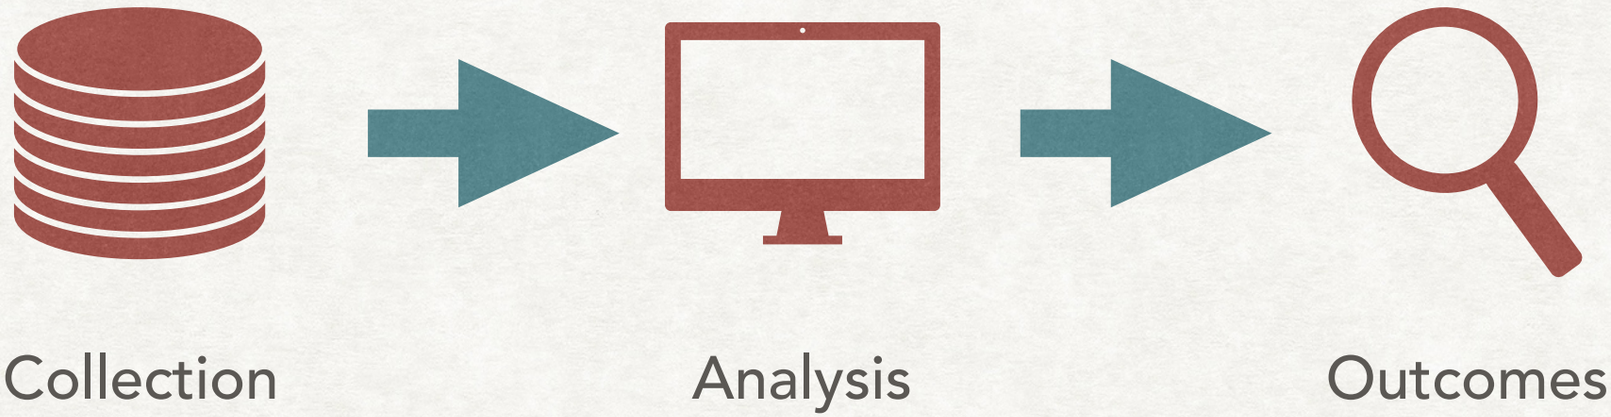

The data analysis pipeline

Choices made in the design of the pipeline have moral consequences

COLLECTION: THE PROBLEM OF ABSTRACTION

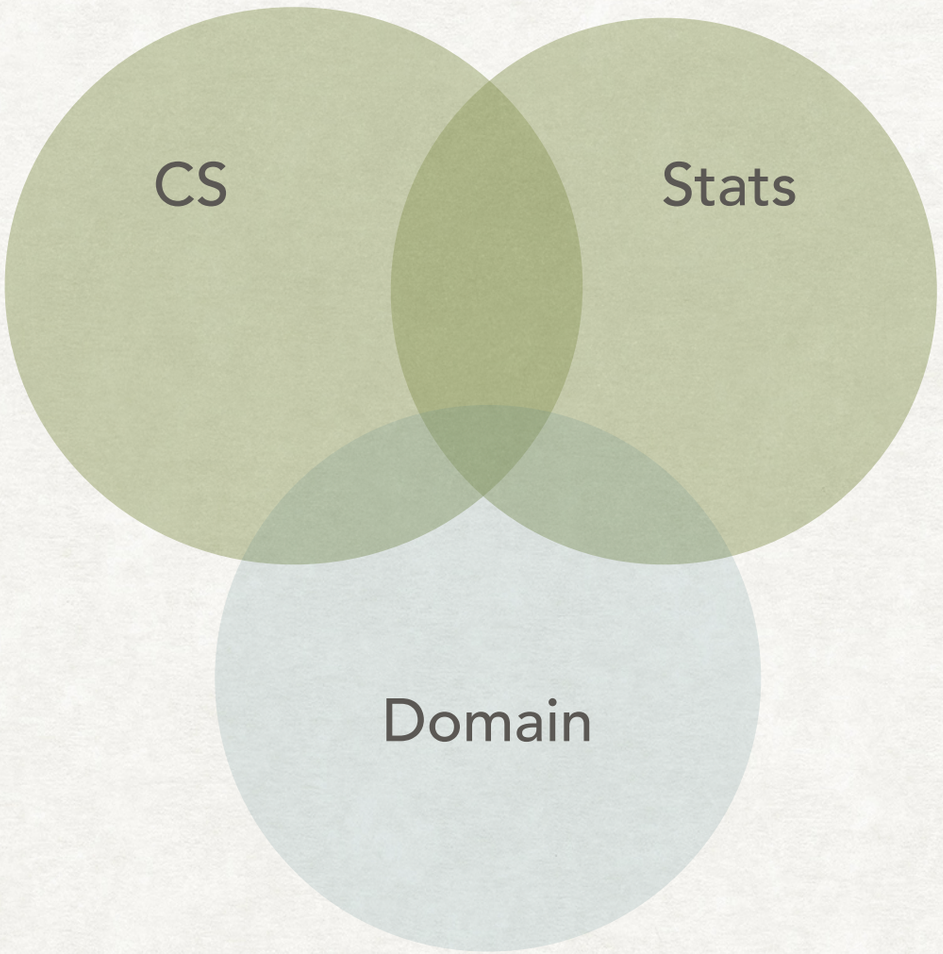

Data processing is disconnected from the domain

Data collected for one purpose might be used for another

We cannot reason about whether data should be used or not

COLLECTION: DATA METAPHORS

How we think about data and our relationship to it changes our ethical compass

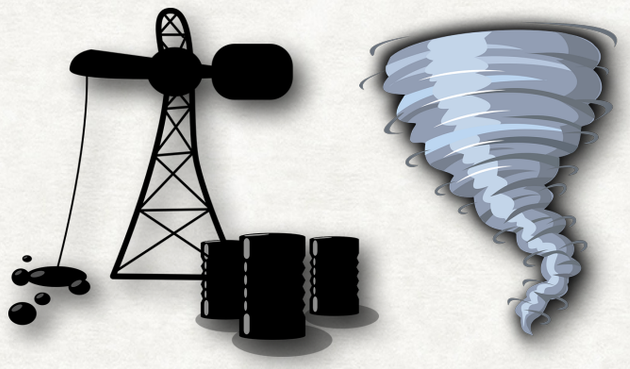

Data as natural resource

Data as personal property

Data as public good

Data as hazardous material [Schneier/Doctorow]

Data as gift [Ferryman]

COLLECTION: POWER RELATIONSHIPS

Who controls the data

Who's asking the questions

Who's data is being used?

Who benefits?

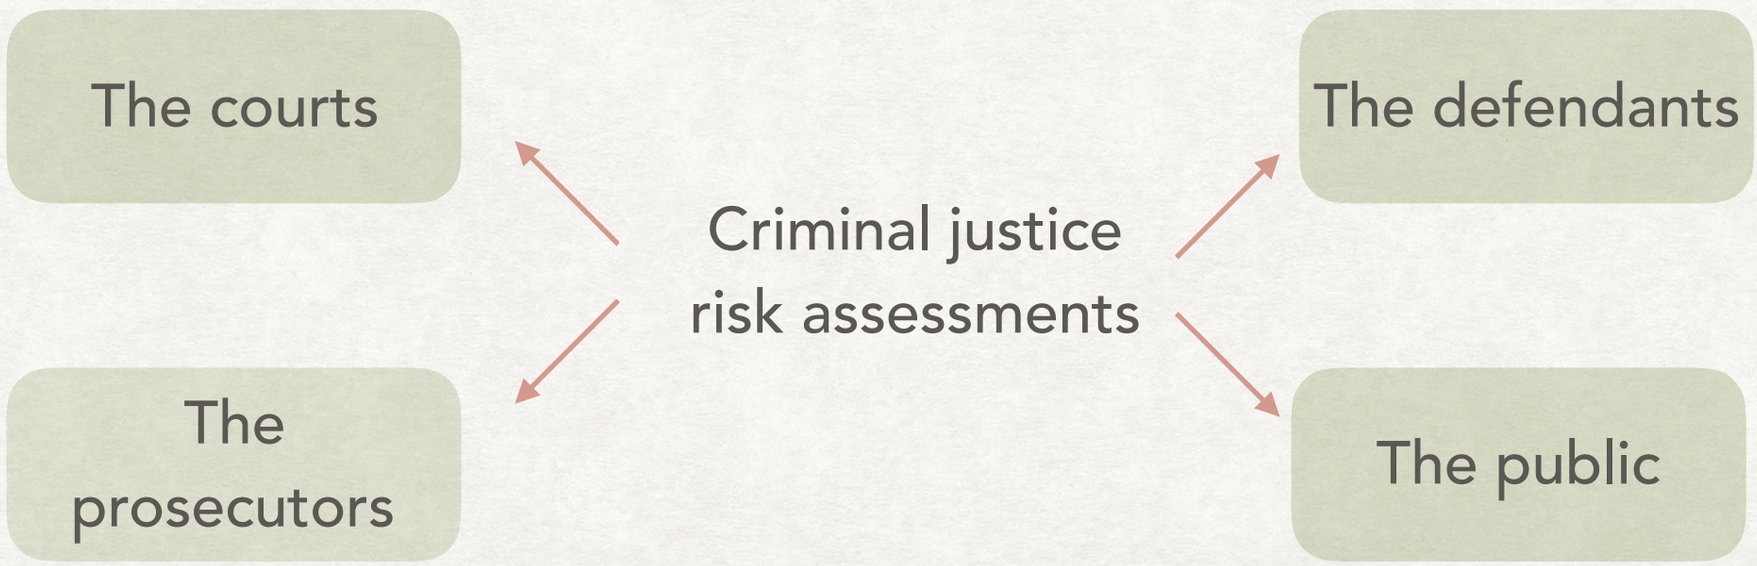

DATA ANALYSIS: MODELING THE DATA

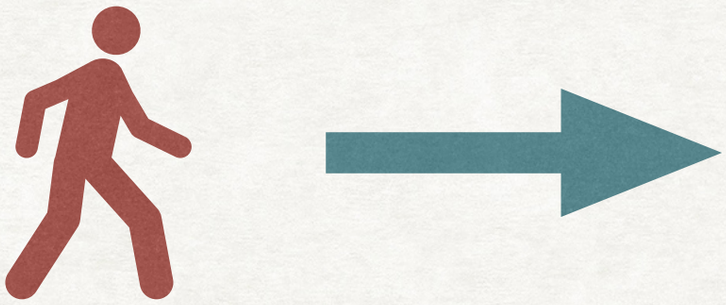

$$(x_1, x_2, x_3, \ldots, x_n)$$

- How do we decide which features to include?
- What encoding do we use for the features that we include?
- How do we normalize the features?
- What happens if we “throw in everything we can get”?

DATA ANALYSIS: MAKING A MODEL

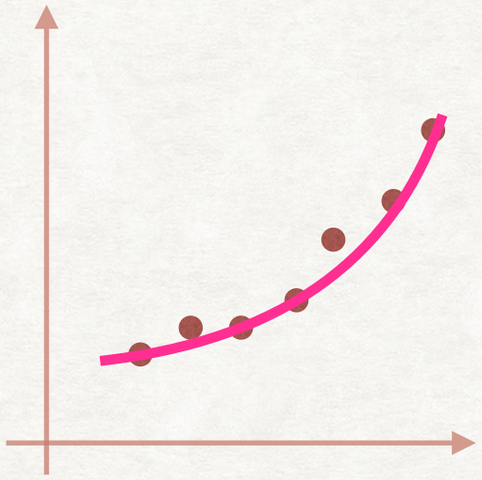

Choosing a hypothesis space

$$\sum_i (f(x_i) - y_i)^2$$

Choosing an error function

Choosing a method of learning

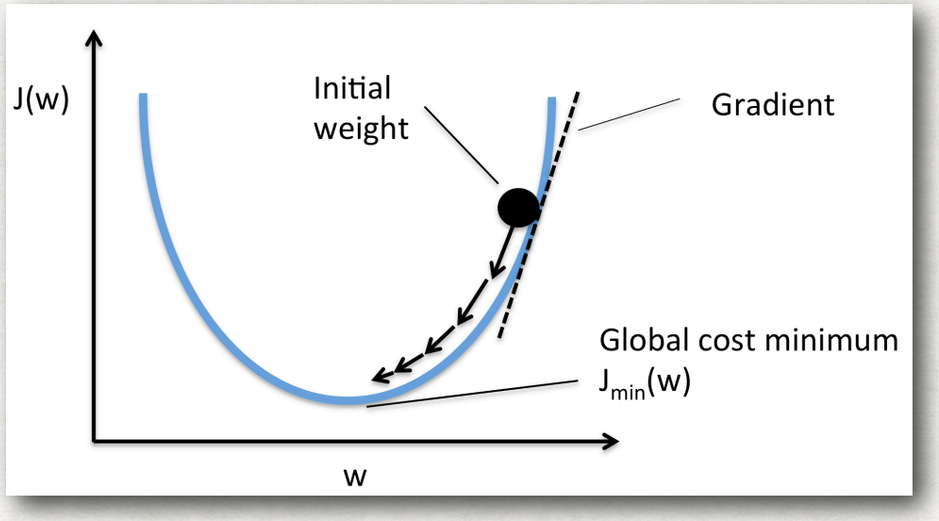

https://rasbt.github.io/mlxtend/user_guide/general_concepts/gradient-optimization/

DATA ANALYSIS: VALIDATING A MODEL

If you torture the data enough, they will confess.

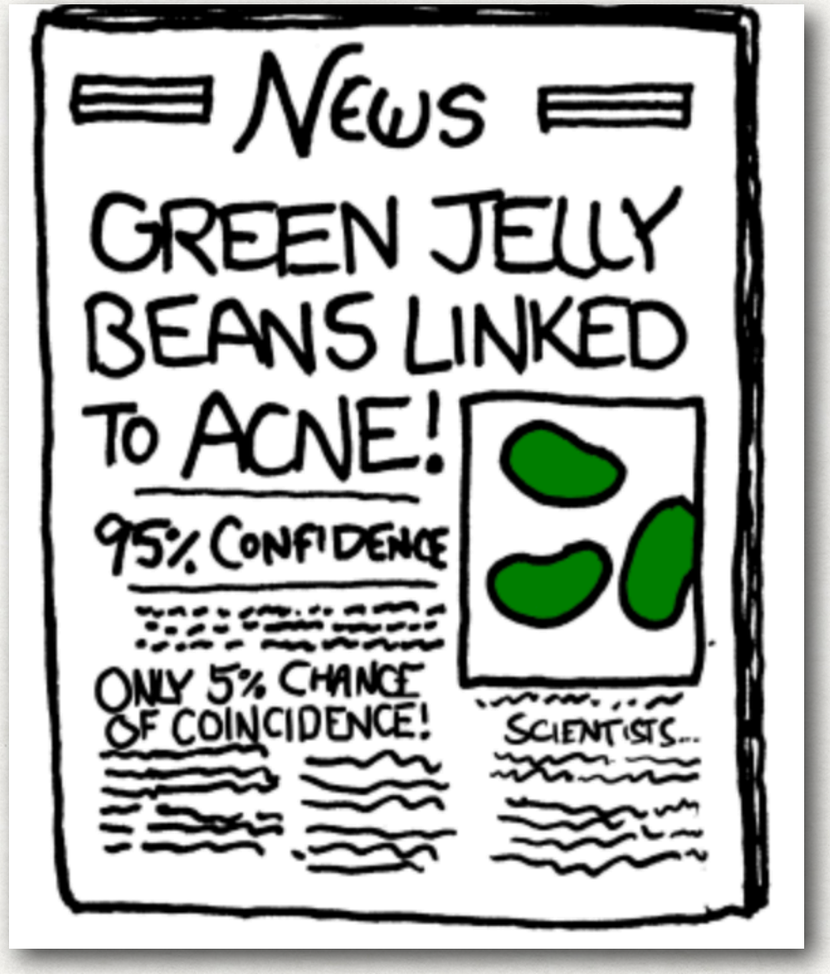

xkcd.com/882/

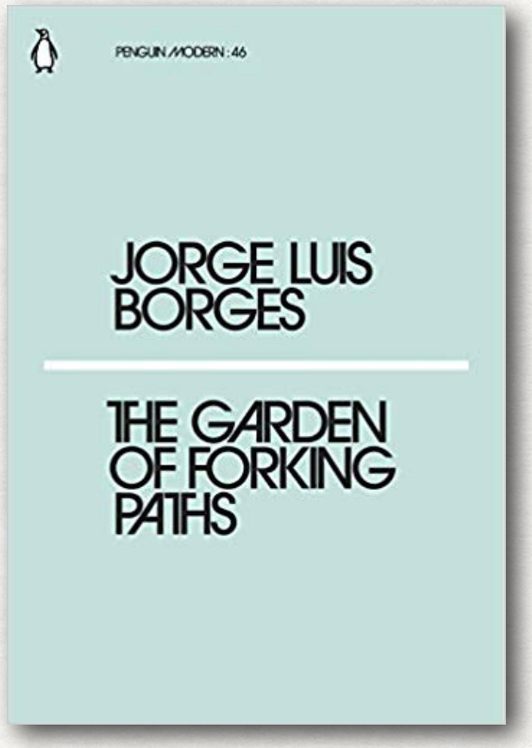

The garden of forking paths: Why multiple comparisons can be a problem, even when there is no "fishing expedition" or "p-hacking" and the research hypothesis was posited ahead of time*

Andrew Gelman† and Eric Loken‡

14 Nov 2013

DATA ANALYSIS: CORRELATION AND CAUSATION

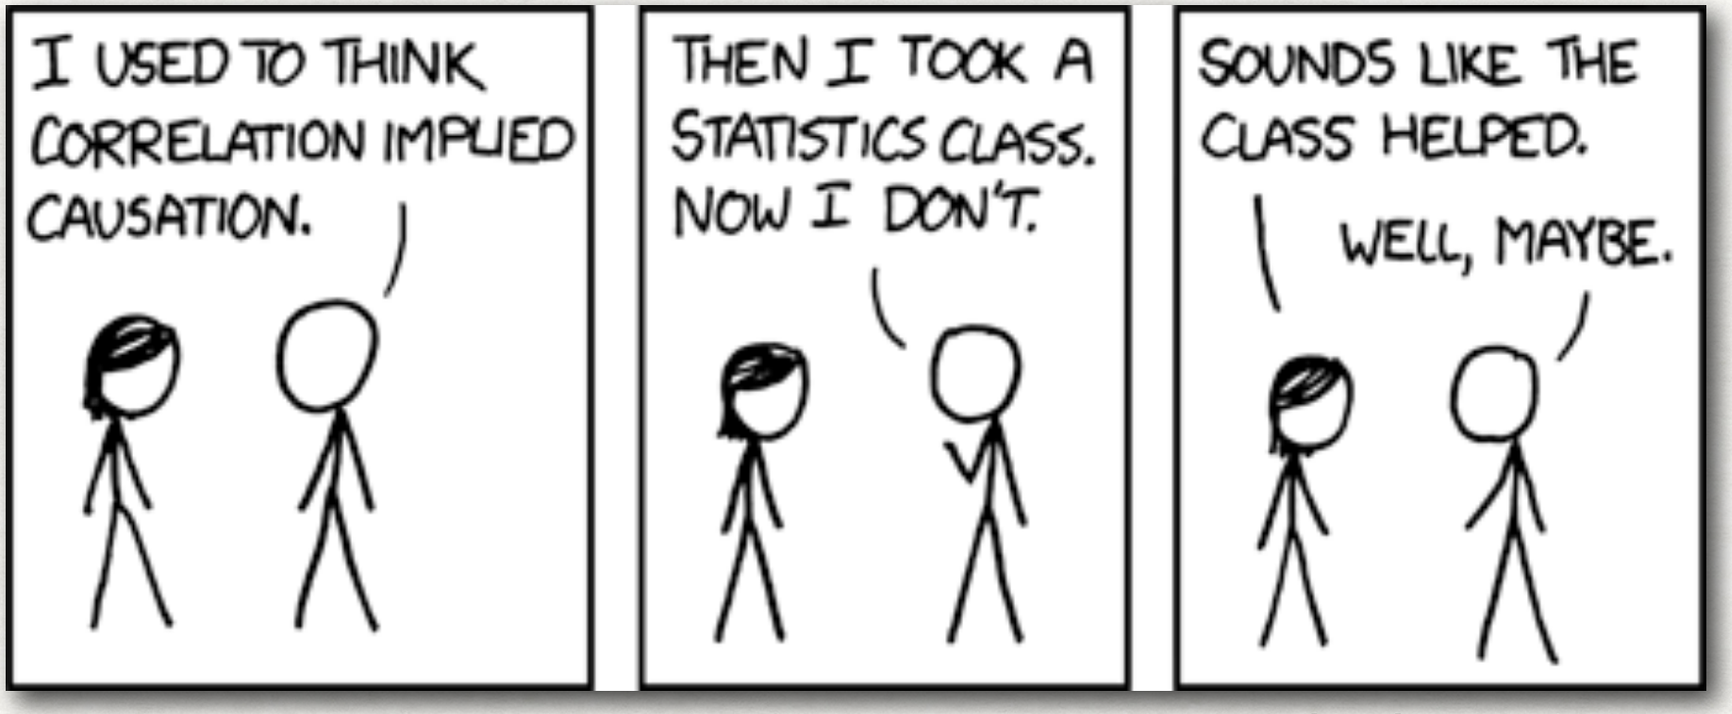

xkcd.com/552

Most ML will tell us that X and Y are correlated. This could mean that

- X is caused by Y
- Y is caused by X
- X and Y are caused by a common factor
- X and Y both cause a common factor Z that is conditioned on (Berkson's paradox)
- X and Y cause each other
- X causes W and W causes Y
- There is no connection between X and Y (bald/hairy Russian leaders)

DATA ANALYSIS: WHAT IS A SCIENTIFIC EXPLANATION?

- Explanations matter because the result of an analysis will affect **people**.
- Not all explanations are the same (Aristotle's 4 types, …)
- Machine learning systems are often inscrutable
  - can we make them more *explainable*
  - can we only use systems that are explainable by design?
  - Should we?

OUTCOMES: DEPLOYING TOOLS

- How do we understand what harms might come from deploying a tool?
- Dimensions of harm:
  - Justice (temporal and static)
  - Distributed vs individualized
  - Scope
  - feedback loops
  - alternate uses

OUTCOMES: RESPONSIBILITY AND OUR ROLE

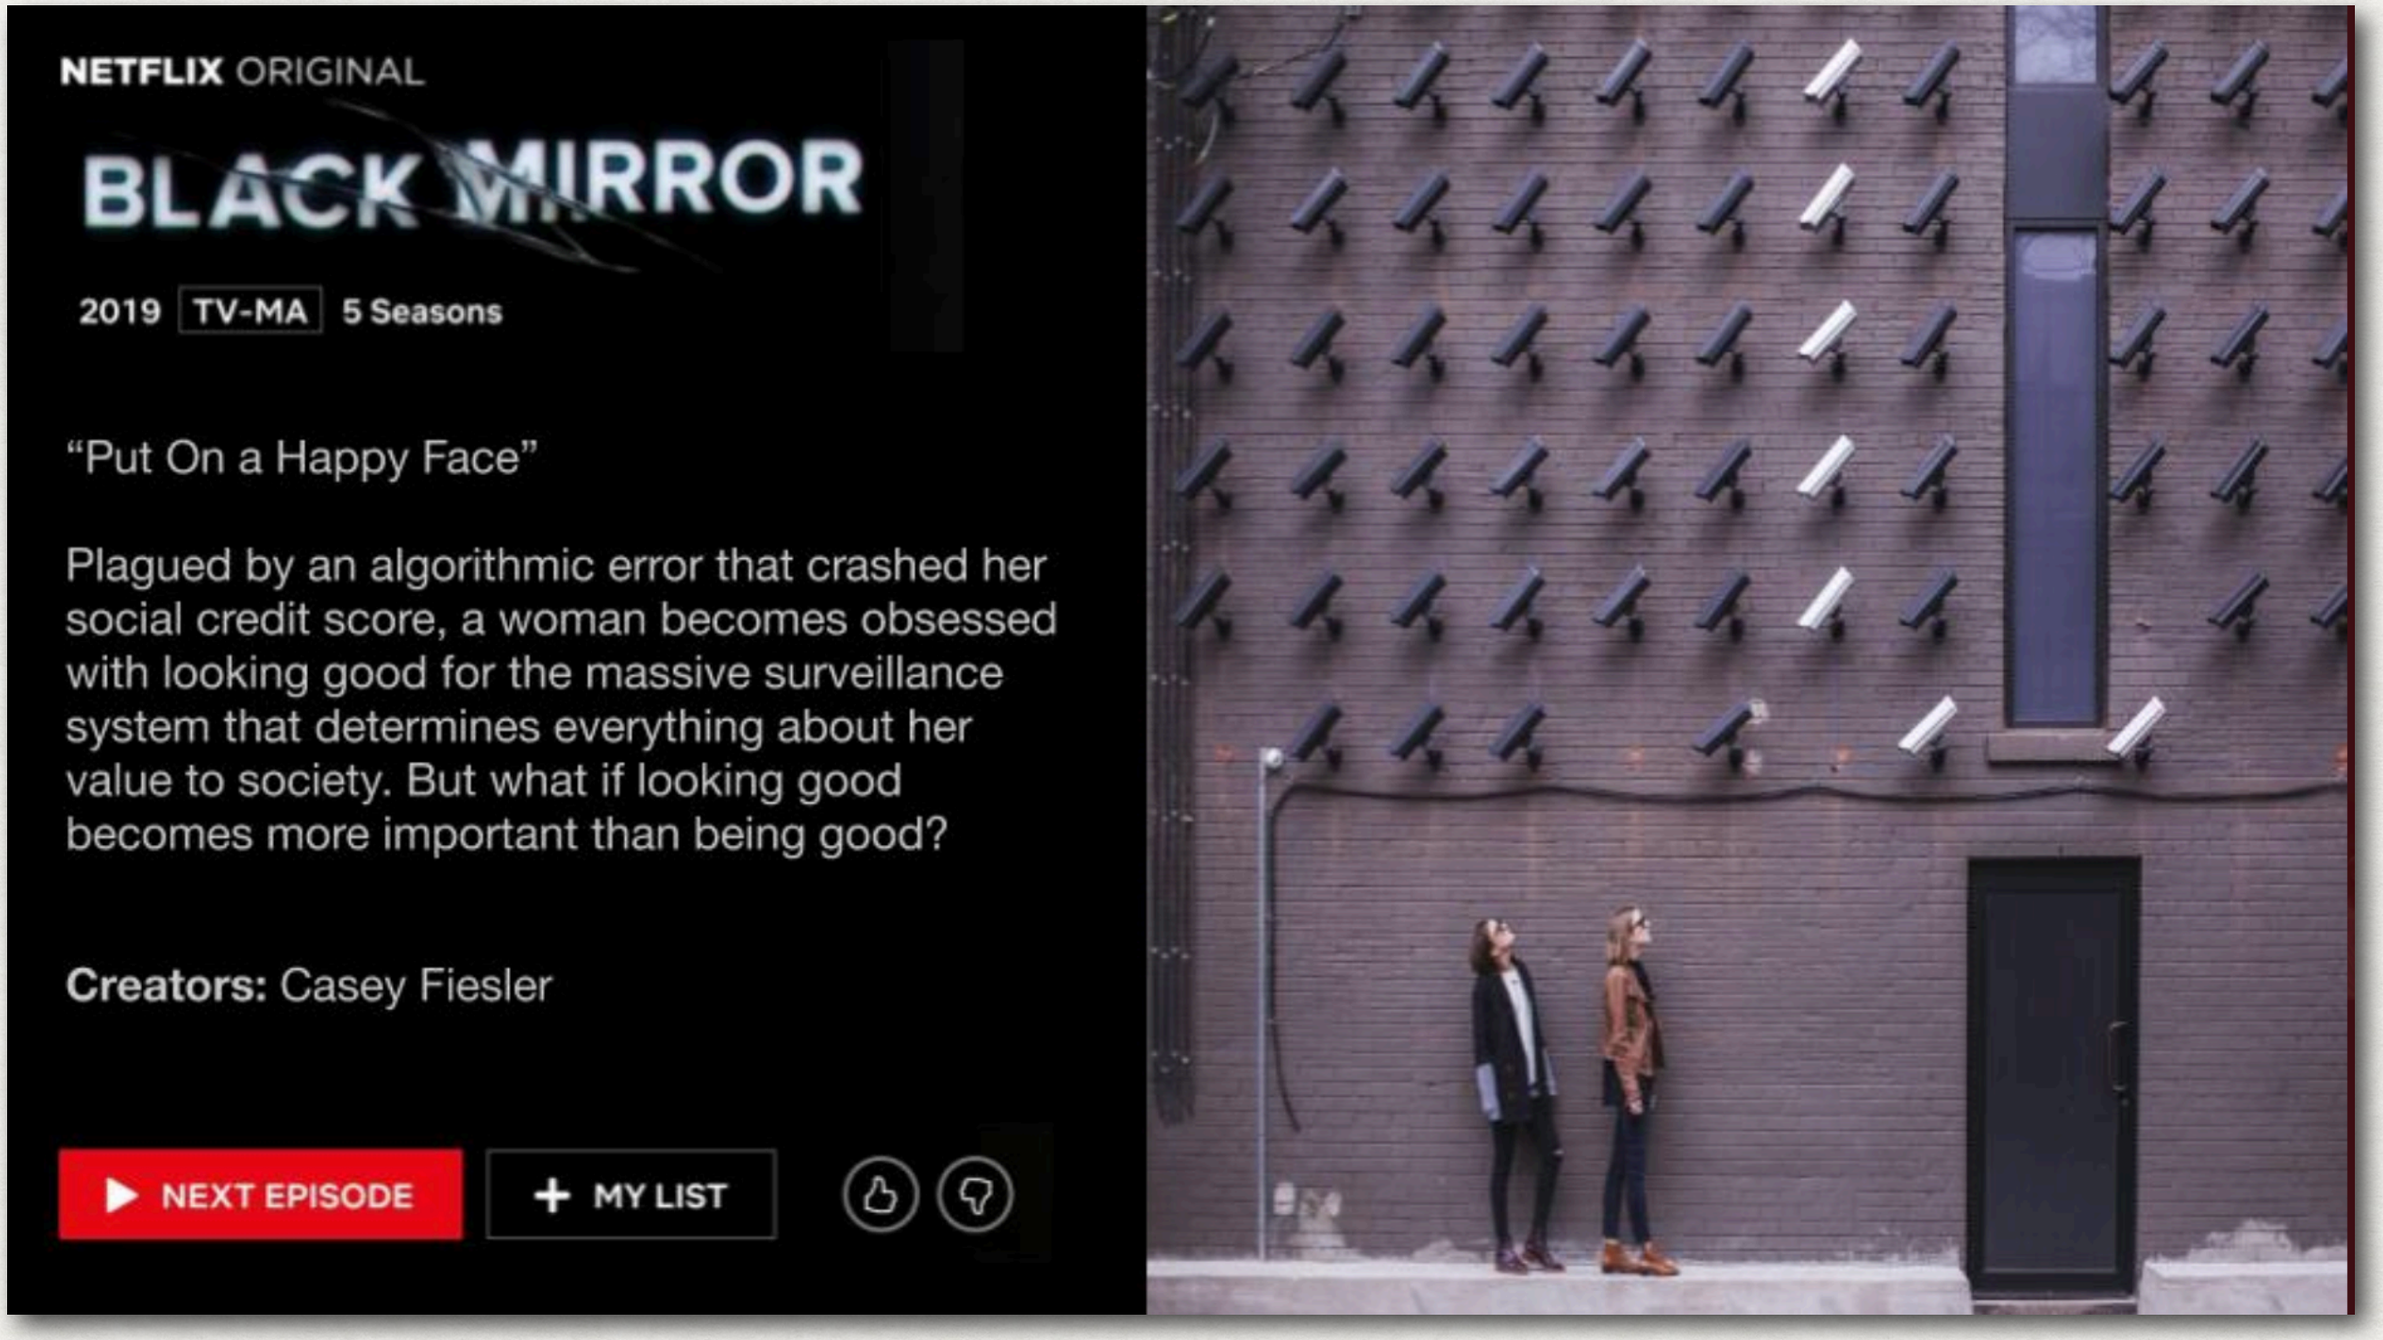

Activity idea by Casey Fiesler @ U Colorado

THE ROLE OF SOCIETY AT LARGE

POLITICS | JAN. 2, 2020

Will the 2020s Be the Decade of Eugenics?

By Sarah Jones

New York

IRB
INSTITUTIONAL REVIEW BOARD SERVICES

AXON AI AND POLICING TECHNOLOGY ETHICS BOARD

ACM Code of Ethics and Professional Conduct

OTHER FRAMES FOR TEACHING

Teach algorithm design through lens of responsibility in modeling.
— with Sorelle Friedler (Haverford), Seny Kamara, Kathy Fisler (Brown)

Certificate in Computing, Ethics and Society (in progress)

- Jointly with humanities and social sciences
- Targeting undergraduates in technical and social sciences.

Award Abstract #2041960

EAGER: SaTC-EDU: Teaching Security in Undergraduate Artificial Intelligence Courses Using Transparency and Contextualization

Use ideas from security such as transparency and contextualization to help students understand sociotechnical vulnerabilities in AI. (Led by Eliane Wiese, together with Mu Zhang — all at Utah)

PEDAGOGY

LECTURES VS INTERACTION

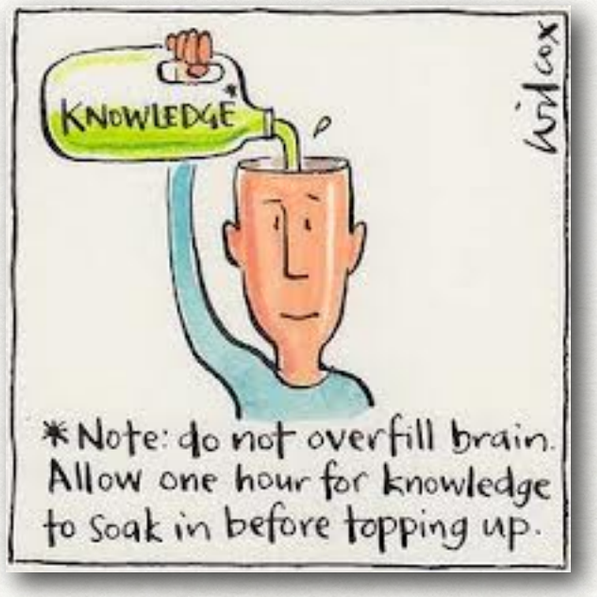

http://capsunm.tumblr.com/post/66989758659/5-tips-for-exam-preparation

Students are not used to discussion-oriented classes in CS.
Used a lot of ‘think-pair-share’

CONSTRUCTING REASONED ARGUMENTS

Sample writing activity

1. Post a constructive question about the reading.
    1. DUE Friday, November 15th
2. Provide a *response* to **your own** question.
    1. DUE Friday, November 15th
3. Provide a *critique* of **your own** response.
    1. DUE Sunday, November 17th
4. Provide a *rebuttal* to the critique of **someone else's** position.
    1. DUE Tuesday, November 19th

See Writing Tips.docx for tips on how to pose a constructive question to stimulate discussion and for more information on posting responses and formulating arguments. Make sure to view the rubric (click the three vertical dots at the top right corner of this box) for this assignment before posting.

THEORY AND CASE STUDIES

Week

Lecture 1	Lecture 2	Lecture 3
Theory (philosophical/ ethical/ mathematical)	Group discussion, case studies, activity	More discussion, reflection on themes

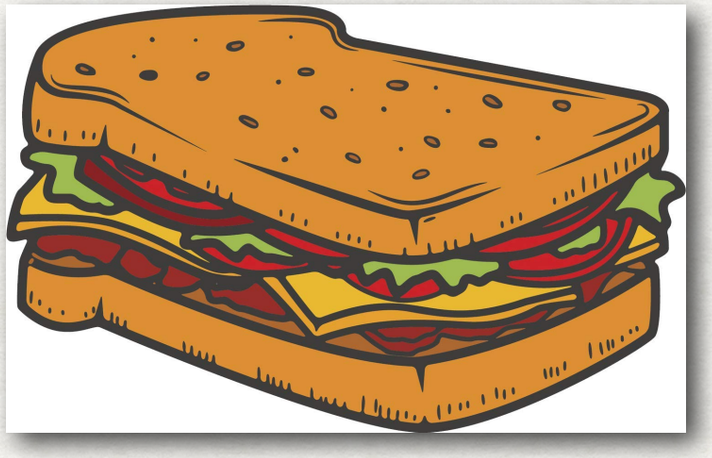

CHALLENGES

- **SCALE:**
  - Discussions are crucial but are hard to do with 100+ person classes (our classes were small).
- **CULTURE:**
  - students (and faculty more so) are hesitant to embrace such classes (viewed as “lightweight”). Not easily transferrable without faculty expertise
- **CURRICULUM**
  - hard to make space in the curriculum for mandatory ethics courses or modules (ABET requirements help)

ACKNOWLEDGEMENTS

- Eleanor Gilmore-Szott
- Kaitlin Pettit
- Katie Shelef

THANK YOU!

suresh@cs.utah.edu