

CS 5/692: Research Ethics for CS & EE



4/4/2017: Intro & Course Logistics

Plan for today:

1. Overview of course
2. What's special about machine learning?
3. Responsible conduct of research
4. History of human subjects research
5. Key ethical principles
6. Models of ethical reasoning

Course overview:

Website/Syllabus:

<http://cslu.ohsu.edu/~bedricks/courses/cs692/>



Readings!

Discussions!

Writing assignments!

Grading!

Textbooks:



What's ethically special about machine learning?

... compared to:

... “traditional” research ethics?

... “traditional” engineering/computer ethics?

What's ethically special about machine learning?

1. Application areas

2. Scale

3. Sources of data

“Instead of relying on algorithms, which we can be accused of manipulating for our benefit, we have turned to machine learning, an ingenious way of disclaiming responsibility for anything. **Machine learning is like money laundering for bias.** It's a clean, mathematical apparatus that gives the status quo the aura of logical inevitability. The numbers don't lie.”

- *Maciej Cegłowski*



ML/AI/etc.

Traditional
research ethics



Emil OW Kirkegaard @KirkegaardEmil · May 8

The OKCupid paper has now been submitted. This means that the dataset is now public! Enjoy! :) [openpsych.net/forum/showthre...](https://openpsych.net/forum/showthread.php?p=11666116)



↻ 26



♥ 38



Ethan Jewett @esjewett · May 11

@KirkegaardEmil This data set is highly re-identifiable. Even includes usernames? Was any work at all done to anonymize it?



↻ 3

♥ 9



Emil OW Kirkegaard

@KirkegaardEmil



+ Follow

@esjewett No. Data is already public.

LIKE

1



1:30 PM - 11 May 2016



Reply to @KirkegaardEmil @esjewett



Ethan Jewett @esjewett · May 11

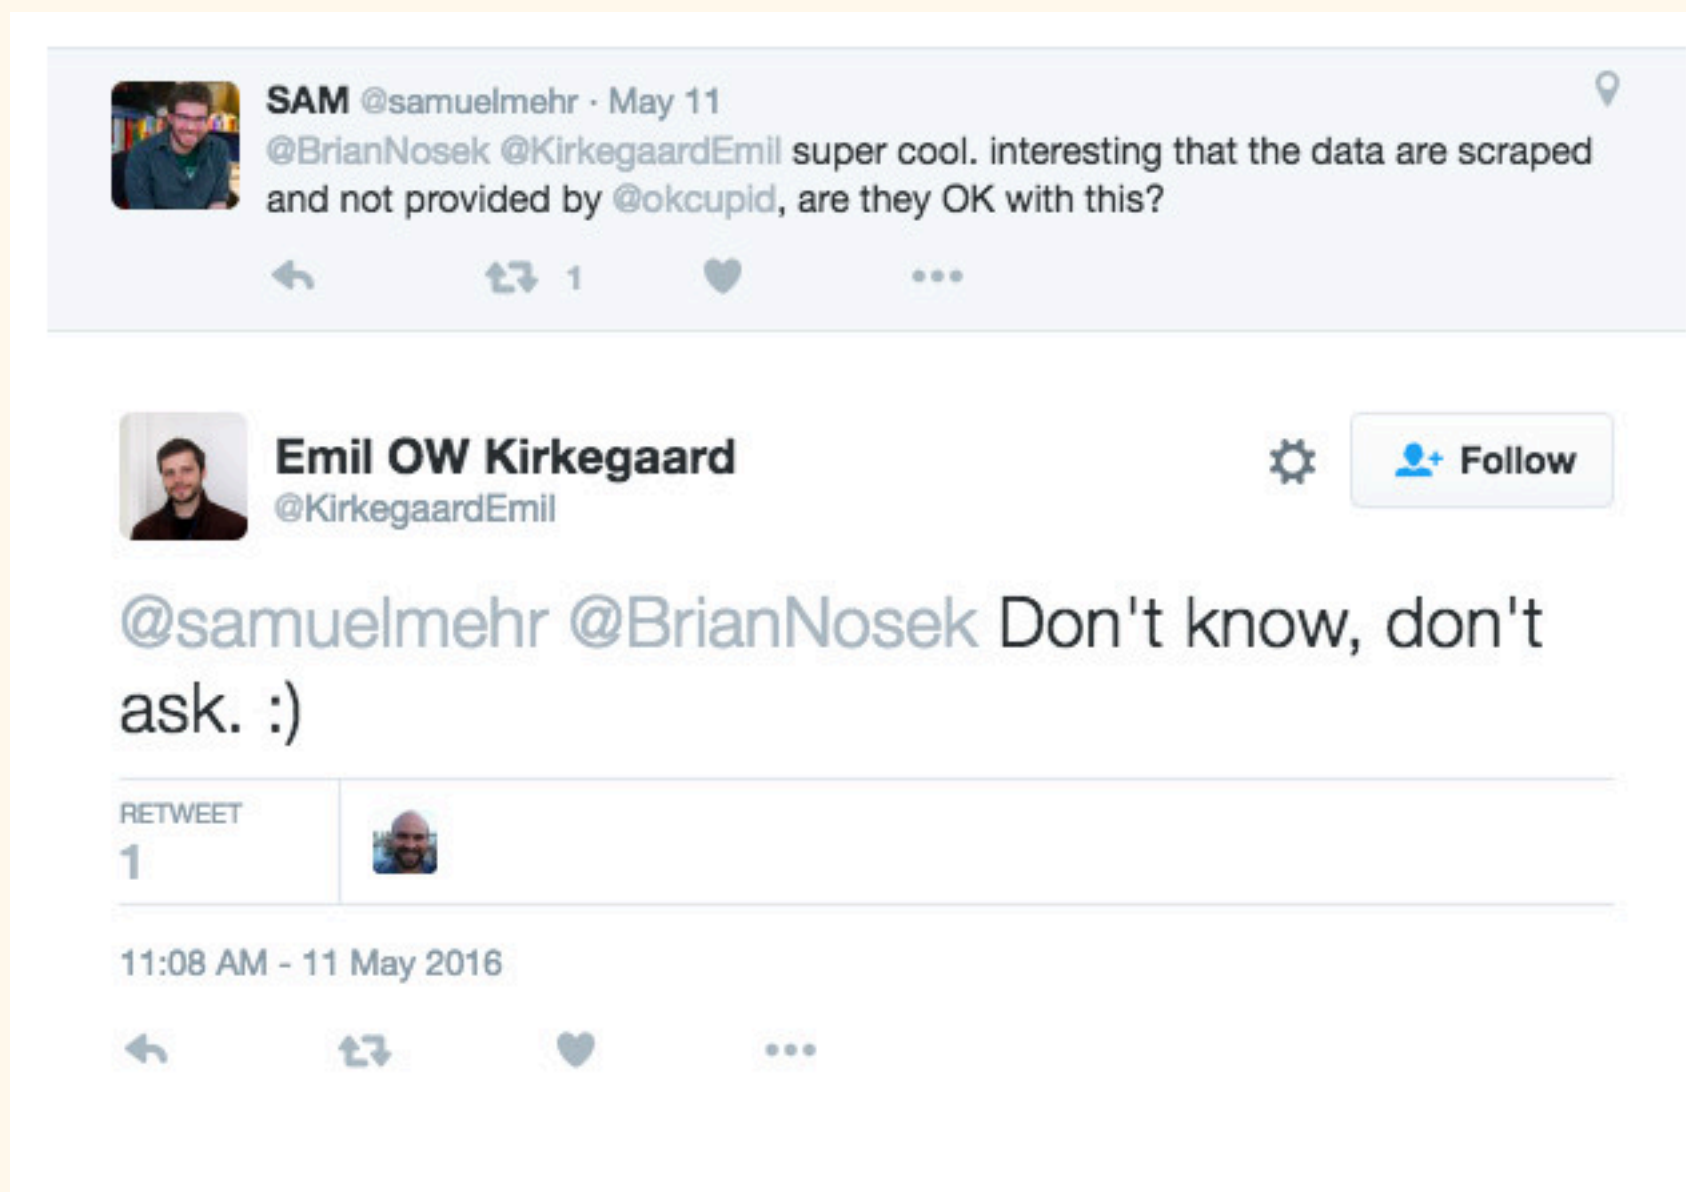
@KirkegaardEmil Differing degrees of "public". Also different ethical guidelines. IMO, you should speak with a research ethicist/IRB ASAP.



↻ 2

♥ 21





Note: Besides the privacy issues, this case had problems with other RCR areas including review standards, methodologies, etc.

Responsible Conduct of Research (RCR)

History of human subjects in research

Foundational principles of research ethics:

Respect for autonomy

Non-Maleficence

Beneficence

Justice

Respect for Autonomy:

Sometimes also phrased as “respect for persons”...

Key idea: individuals have the right to make informed and voluntary decisions about their participation & involvement.

Important corollary: researchers have a duty to protect individuals who are unable to fully exercise that right.

Non-Maleficence:

First, do no harm.

Research must be designed in such a way as to not harm subjects...

... or, if harm is necessary, it must be carefully controlled, minimized, and disclosed to the subject.

Investigators have a duty to actively consider possible harms, and ask “what could go wrong?”

Beneficence:

The goal of the research should be to benefit the participant in some way (inc. indirect benefit).

1. Do no evil/harm
2. Prevent evil/harm
3. Remove evil/harm
4. Practice good

Moral/ethical systems define beneficence in different ways...

Justice:



The benefits and burdens of research should be distributed fairly and equitably.

E.g.: specific populations should not bear disproportionate costs (i.e., subjects should not all be from one group)



Researchers bear responsibility for disseminating results equitably, etc.

Models of ethical reasoning:




Consequentialism:



“Nature has placed mankind under the governance of two sovereign masters, pain and pleasure. It is for them alone to point out what we ought to do, as well as to determine what we shall do.”

Jeremy Bentham , *Principles of Morals & Legislation*



Deontology:



“Act only according to that maxim whereby you can at the same time will that it should become a universal law without contradiction.”

Immanuel Kant, *Grounding for the Metaphysics of Morals*

Trite illustration:



IMAGINE YOU'RE IN AN OUT-OF-CONTROL TROLLEY. YOU'RE HEADED TOWARD THREE BUILDINGS, AND YOU CONTROL WHICH YOU SLAM INTO. TWO BUILDINGS CONTAIN FIVE PEOPLE ONE BUILDING CONTAINS ONLY ONE PERSON. YOU RANDOMLY SELECT A BUILDING TO SLAM INTO. THEN, ONE OF THE OTHER BUILDINGS IS REVEALED TO CONTAIN FIVE PEOPLE. SHOULD YOU SWITCH TRACKS?



So far, no ethicists are impressed with the Monty Hall Trolley Problem.

Assignment:

1. Pick one of the basic principles (autonomy, etc.)
2. Reflect (in writing) about how it might come into conflict with a machine learning / NLP/ AI research question