



UNIVERSITY OF NOTRE DAME
MENDOZA COLLEGE OF BUSINESS

ITAO 40510 Ethics of Data Analytics

Professor: Kirsten Martin, Ph.D.
Kmarti33@nd.edu

Class Meets: 2X per week (75-minute class periods)

Course Overview and Objectives

Course Overview

Every day, in a business close to you, someone furrows their eyebrows and thinks, “This just seems wrong.” Sometimes that moment turns into a larger discussion. Sometimes that moment gets buried and ignored – at least until the problem manifests into a full-blown dilemma. With data analytics, increasingly that problem finds its way to the press.

In this course, we will explore the ethical implications of data analytics. We will marry old ideas – privacy, surveillance, power, justice, accountability, corporate responsibility, stakeholder theory – with new technologies and cases, such as the use of machine learning to predict crime. The goal is to tie our discussions around data analytics with ethical concepts that transcend our current issues. This way, you will be equipped to deal with newer, ambiguous capabilities once you graduate. Technology changes, but issues of power, privacy, justice, surveillance, and marginalization withstand the test of time.

The course is designed so that each day, we analyze a current ‘case’ (in the form of an article or two) with a theory or concept. Our readings will include the original works of philosophers, economists, law professors, and computer scientists. Tackling the ethical issues of data science requires a wholistic approach, and no one discipline has the answer.

We are one of a few business schools to offer a course in the ethics of data analytics. While the computer scientists you may work with in the future study these topics, most business schools do not. And because we are in a business school, the course is also designed to explicitly link the ethics of data analytics to theories in business ethics (corporate responsibility, stakeholder theory, economics of trust, etc). The goal is to provide you the tools to be able to make strategic arguments about the ethics of data analytics using the language and concepts of business.

Course Objectives

In this course, you will:

1. Develop the capacity to have constructive conversations about ethical dilemmas of data analytics.
2. Identify types of ethical issues within each stage of data analytics.
3. Examine a particular ethical issue in data analytics and apply multiple theories to understand possible solutions.
4. Explain the current issues in AI/ML/Predictive Analytics in the language of business and business ethics.
5. Develop a series of questions to ask of any data analytics program in order to identify the possible moral quandaries and ethical implications.

Required Material

Textbook: *None*

Articles: All required readings are provided on Canvas within the ‘Module’ labeled for the day the reading is due. Each day includes readings labeled “Case” and readings labeled “Theory.” You are expected to have read the required readings prior to class.

Optional

For each lesson, I have provided extra readings in case you are interested in (a) the particular case, (b) the philosophy behind the concepts of that day, or (c) the relevant computer science papers. These are labeled as “Extra” on the lesson for that day with a button. These are *optional* and we will not go over them in class. For many of you, the similar cases may provide ideas for your final report. I am not assuming any of you want to read philosophy or an engineer’s conference paper but the option is there.

Course Format, Assignments, Evaluation and Grading

Format

We will run the classroom as a case discussion. We will rely upon you and your peers to discuss and question the cases and examples in order to identify lessons, develop moral awareness and empathy, and practice arguing about ethical issues. In order to get the most out of a case discussion, *everyone must be ready, attentive, and participating*. Please be sure to read and study the cases and readings that are assigned. Students are expected to attend each class and be prepared and ready to contribute.

Performance Measurements

- Participation 30%
- Pre-Class Qs 20%
- Memos 20%
- Final Report 30%

Participation (30%).

Your participation will make this class what it is and will help your colleagues think through the ethical issues. Actions that hinder your ability to be attentive, respectful, and participating hurt your grade. For example, if you are texting in class, I simply write “text” in for your grade for that particular day. If you have questions about participating, please ask and we can work on ways to make sure you participate. You will receive feedback (from me and anonymously from your peers) on your participation in order to ensure there are no surprises. The cases and assignments are designed to reinforce the importance of participation.

If you are not comfortable participating – please come see me or reach out to me ahead of time. The time we have together is short and case discussion start on the first day. I have ideas as to how to make your participation easier.

Each class will assume that you have done all the readings, are prepared to discuss with your peers, and apply the materials during class learning activities. At a minimum, I expect you to attend class and prepare readings in advance; that would constitute an entry-level “C.” In order to receive a better grade, you must participate. Participation is graded based on attendance, your willingness to participate, demonstration of having prepared for class, and thoughtfulness of class participation.

I provide a midpoint grade and feedback for participation. You will provide feedback to me in an assignment – where your contribution was particularly helpful in moving the class forward and who you look forward to hearing from in class. This is 1-2 sentences and should take 30 seconds. My grade of that assignment is your first half grade in participation.

Participation Rubric.

	<u>No credit</u>	<u>Partial Credit</u>	<u>Full Credit</u>
Spoke	Nothing		Contributed to case discussion
Use of theory	No theory	Some theory	Actual theory assigned that day
Link to other examples			
Other	Use of phone		Give voice to alternative view

Pre-Class Qs (20%) – due before each class.

Prior to each class, you will submit a question you would ask of any future data analytics project given the case/concepts read for that day with 1-2 sentences explaining why. It should be clear form your question and explanation that you read the material. This is lightly graded (0, 1, 2) with most everyone getting a 2. This is due BEFORE class and no late submissions will be accepted. I only keep the top 10. The goal of this assignment is to have you thinking about more generalized yet practical implications of the specific case you read for class.

The link to this ‘quiz’ is in the module for a given day.

Memos (20% total)

Each student is expected to submit two written case study memos – one before and one after the midpoint. The memos are 1-2 pgs and will analyze the case assigned for that day using the theory also assigned for that day. You are to give me your assessment of the issue using reasoning from the course. I will grade you on your ability to understand and apply the ethical framework. The write-ups are due *before* the class discussion that day – in other words you do not get the benefit of the class discussion before completing the write-up. A question designated ‘memo’ below cases in the reading section provides the prompt for your case memo.

Memo Rubric – The emphasis is on the theory being applied.

	<u>No credit</u>	<u>Partial Credit</u>	<u>Full Credit</u>
Structure 20%	No structure	Structure but not around theory	Structure based on theory
Theory 80%	No theory	Parts of theory but not all	All of the theory correctly applied
Case	No case mentioned	Case generally mentioned with no specifics	Case facts used in analysis

Final Report (30%)

You are assessed on a final case report (there is no final exam). Each student will write a report on an issue in data analytics of their choice. *This assignment will allow you to connect classic ethical cases, problems, and theories to current dilemmas and offer solutions.* The 5-6 page report must have headings and use of at least two theories from the class to analyze an issue of your choice (but not a case we covered in class). All papers will have explicit use of theory applied to a real case. We will talk more about the paper closer to the due date. The goal of this assignment is to use theories from the class to analyze a real data analytics dilemma. I have provided time in class for peer feedback on the proposal and outline. The proposal and outline will also be submitted for feedback from me. The project is due during the last week of class.

20% of this assignment is your ability to link your analysis to four examples from the class discussion.

The paper or project must include the following (% corresponds to portion of the final report grade).

1. 10%. Explain an ethical issue/case in data analytics. This can be from a few news articles. 1-2 pages maximum. This should read similar to a newspaper article without any ethical analysis.
2. 60%. Apply a minimum of 2 concepts from class to analyze what (if anything) is wrong with the case. The concepts must be clearly labeled (headings, bolded, underlined, etc). 4 pages minimum.
3. 10%. Provide a possible alternative or solution to any problem identified in #2 above. This should include a technical alternative such as a different set of training data used, different algorithm or approach in data analytics, or

different type of governance approach (just as examples). Half page approximately.

4. 20%. The report should include four references to examples or cases from class. These can be either assigned cases or examples that come up in the class discussion. This would be as short as “this company violated the privacy expectations of their users by using the commercial data set similar to the case of ad tech in class, both.....”

There are three steps to this process. You will submit a proposal and outline as attempts in Canvas for the final report. I will provide feedback on both. You will be given time to work on the proposal and outline in class to get feedback from your peers.

Grading

Grading criteria for assignment of final course grade is based on:

<u>Percentage of Course Points Earned</u>	<u>Grade Earned</u>
93.0 % and above	A
from 90.0 % but less than 93.0 %	A-
from 87.0 % but less than 90.0 %	B+
from 83.0 % but less than 87.0%	B
from 80.0 % but less than 83.0%	B-
from 77.0 % but less than 80.0 %	C+
from 73.0 % but less than 77.0 %	C
from 70.0 % but less than 73.0 %	C-
from 60.0 % but less than 70.0 %	D
less than 60.0%	F

*The professor reserves the right to ‘curve’ grades to comply with the Mendoza College of Business’ grade point average range of 3.0 to 3.4.

The students are expected to adhere to the University Honor Code (<https://honorcode.nd.edu/>). Any violations of the Honor Code will be referred to the appropriate committee.

Schedule.

Class	Topic	Case/Readings	Extra
1	Introduction to Algorithmic Bias	Bias: Stanford Vaccine Algorithm Nissenbaum, Winner	
2	Privacy	Finding Consumers Nissenbaum on Privacy	
3	Biased Datasets	Facial Recognition: Wrongfully Accused	
4	Purpose of Corporation/AI.	Inferences: Predicting Ethnicity Friedman on Corp Resp & Profits Freeman et al on Stakeholder	
5	Fairness, Predictive Analytics, & Mistakes	Predictive Analytics: Compas. Rawls, Nozick, Walzer on Fairness	Time to work on proposal
6	Surveillance & Power	Secretive Company. Lyons on Surveillance ** last day to submit memo	
7	Disparate Treatment/Impact	Hiring Algorithms: Amazon Design college admissions in class. Selbst and Barocas on DT/I	Time to work on proposal.
8	Econ of Trust & Stakeholders	Content Moderation: Facebook Robert Frank on Econ of Trust	
9	Transparency & Accountability	Houston Teachers Crawford et al on Contestability	
10	Creating and Measuring Accuracy	Algorithmic Traps: Criminal Students	Time to work on Final
11	Summary Class.	People Analytics: Cheating Detection	
12	OFFICE HOURS	Schedule 15 min	
13	Role of AI Ethics in Corp/ Role of Corp in AI Ethics	Google Research CASE or Student PPT	
14	FINAL REPORT DUE		

Class 1 Introduction to Algorithmic Bias Stanford Vaccine

The idea that computer programs or technology has bias - preferences or a forcing function to guide users - has long been understood. The question for us is how do biases 'look' in the area of AI or data analytics? We will read two articles on technological biases and discuss the possible biases in the Stanford Vaccine algorithm.

The goal of the class is be able to identify types of bias in analytics, sources of those biases, and how those biases could be ethical or unethical.

Case:

1. Guo, Eileen and Karen Hao 2020. "This is the Stanford vaccine algorithm that left out frontline doctors" MIT Technology Review
<https://www.technologyreview.com/2020/12/21/1015303/stanford-vaccine-algorithm/>

Questions:

1. *Is there anything wrong with using a computer program to determine who gets a vaccine? Why or why not?*
2. *If the program is wrong, who is at fault?*
3. *What type of bias do you see in the Stanford algorithm?*

Concepts:

- Pg. 330—336 only. Friedman and Nissenbaum. Bias in Computer Systems.
<https://nissenbaum.tech.cornell.edu/papers/Bias%20in%20Computer%20Systems.pdf>
- Alex Miller, Want Less-Biased Decisions? Use Algorithms, Harvard Business Review, July 26, 2018. <https://hbr.org/2018/07/want-less-biased-decisions-use-algorithms>
- Pg. 121-128 only. Winner, Langdon. 1980. "Do Artifacts Have Politics?" Daedalus 109(1): 121-136.

Class 2 Privacy & Shared Responsibility Ad Tech

Privacy has been defined in three ways: (1) as the degree you control your information, (2) as whether others have access to you and your information, and (3) the negotiated norms around how information should be gathered, shared, used (and by whom). We will focus on the third definition and discuss the limitations of the first two. Our questions focus on whether privacy is violated with the current ad tech system and who is responsible for any privacy violations, harms, or diminishment of value/rights.

Case:

- Finding Consumers, No Matter Where They Hide: Ad Targeting and Location Data
- Tracking Overview: Angwin, J. 2020. “What They Know ... Now” The Markup.
<https://themarkup.org/blacklight/2020/09/22/what-they-know-now>

Concepts:

- **Pg. 32-36 only.** Nissenbaum, Helen. 2011. “A Contextual Approach to Privacy Online.” *Daedalus* 140(4): 32-48.
<http://www.cs.cornell.edu/~shmat/courses/cs5436/contextualapproach.pdf>
1. *MEMO: When does targeted advertising violate consumers’ privacy expectations? Why?*
 2. *Who is responsible for any privacy violations with hyper targeted advertising?*

Class 3

Biased Data Sets.

Facial Recognition & NLP

For analytics that relies on 'learning' from data, how does the choice of data impact the eventual model? Data scientists, who choose the data, have a role in assessing the appropriateness of the data as representative of the 'real world' and as impacting the eventual model used for categorizing and predicting. For us, we will focus on the ways data could be biased, how that bias is different (if at all) from biases in the experienced world, and who is responsible for the moral implications of that bias.

Case:

- Hill, K. 2020. Wrongfully accused By an Algorithm. NY Times. <https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html?action=click&module=RelatedLinks&pgtype=Article>
- SKIM SITE. Buolamwini, Joy, and Timnit Gebru How well do IBM, Microsoft, and Face++ AI services guess the gender of a face? <http://gendershades.org>
- Actual Data. Quach, K. 2020. MIT apologizes, permanently pulls offline huge dataset that taught AI systems to use racist, misogynistic slurs. https://www.theregister.com/2020/07/01/mit_dataset_removed/

Concepts:

- SKIM. Gebru, Timnit, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. "Datasheets for datasets." arXiv preprint arXiv:1803.09010 (2018). <https://arxiv.org/pdf/1803.09010.pdf>

1. *MEMO: Was the arrest morally wrong? Why or why not? Who is at fault?*
2. *Do corporations have an obligation to vet data sets used in creating AI? Why or why not?*

Discuss In Class.

- NLP for class discussion. Ghaffary, S. 2019. Vox. The algorithms that detect hate speech online are biased against black people. <https://www.vox.com/recode/2019/8/15/20806384/social-media-hate-speech-bias-black-african-american-facebook-twitter>

Facial recognition technology can be used not only for identification purposes (or mis-identification), but also to predict or guess something about the person. The idea is that the structure of the face ‘reveals’ hidden attributes about the person. Here we will focus on the use of facial recognition programs to identify race or emotions -- and whether the purpose of the program fits with what we understand to be the purpose of the corporation.

Cases Facial Recognition

- Olson, Parmy. 2020. The quiet growth of race detection software. The Wall Street Journal
- Harwell, D. 2019. A face-scanning algorithm increasingly decides whether you deserve the job. The Washington Post

Extra

- 2020. “Row over AI that ‘identifies gay faces’” BBC <https://www.bbc.com/news/technology-41188560>
- Huawei tested AI software that could recognize Uighur minorities and alert police, report says”. Washington Post. <https://www.washingtonpost.com/technology/2020/12/08/huawei-tested-ai-software-that-could-recognize-uighur-minorities-alert-police-report-says/>
- Williams-Lee, J. 2020. Amazon’s A.I. Emotion-Recognition Software Confuses Expressions for Feelings. Medium. <https://onezero.medium.com/amazons-a-i-emotion-recognition-software-confuses-expressions-for-feelings-53e96007ca63>
- Thompson, C. 2020. Rediscovering “Good” and “Bad” Heads in the Phrenological Present. <https://nursingclio.org/2020/12/08/rediscovering-good-and-bad-heads-in-the-phrenological-present/>

Concepts. Purpose of the firm.

- Freidman – Corporate Responsibility of Corporations is to Increase Profits.
- Freeman et al. Chapter on Stakeholders.

Questions:

1. *MEMO: Should a company develop AI software that purports to identify ethnicity? Why or why not?*
2. *Should a company use AI software that purports to identify ‘emotions’? Why or why not?*
3. *What is a good use of technology that would identify ethnicity? emotions?*

Class 5 Fairness. (Rawls, Nozick, Walzer). Compas

Fairness, Justice, and COMPAS Sentencing Program

The question of fairness is important for us since analytics and AI are frequently used to allocate 'things' to people (employment, vaccines, ads, medical care, punishment, sentencing, money, etc). Who gets what and whether the allocation is fair will impact the legitimacy of the analytics program. Here we study three philosophers of justice - Rawls, Nozick, and Walzer - to analyze whether the COMPAS sentencing algorithm is fair.

Rawls is always concerned with inequalities and finds inequalities suspect and needing a justification. Rawls wants to know if the position to which the inequality is attached was open to everyone and whether the entire system works out for everyone -- but particularly the least fortunate in society. Rawls also gives us the tool of the veil of ignorance to design the system of allocation, where we should pretend to not know who we are and what goods we have when we design the system of allocation. Rawls gives us an argument when we find a system reinforces existing inequalities or advantages one group over others.

Nozick sees inequalities as a natural extension of people wanting to do different things and only questions whether the person handing out the inequality earned it justly (without coercion, fraud, manipulation, etc) and whether they transferred it to another justly (without coercion, fraud, manipulation, etc). Nozick's default is that inequalities are expected and one needs to justify interfering with inequalities. Nozick provides the argument for the person who says "It's theirs and they can give it to whomever they want."

Finally, Walzer introduces the idea of spheres of justice and provides us with a way of seeing 'winners' in the allocation in one sphere of life (e.g., education) should not automatically become winners in another sphere of life (e.g., health care). Walzer provides language and an argument when we see someone 'unjustly' winning or losing in a sphere such as sentencing a prisoner based on attributes from another sphere (income of father).

Case:

- Compas Algorithm. Julia Angwin, Jeff Larson, Surya Mattu, Lauren Kirchner, "Machine Bias," ProPublica (May 23, 2016) available at <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Angwin, J. Larson, J. 2016. Bias in Criminal Risk Scores Is Mathematically Inevitable, Researchers Say. ProPublica. <https://www.propublica.org/article/bias-in-criminal-risk-scores-is-mathematically-inevitable-researchers-say>

1. *MEMO: Is using Compas to make sentencing or parole decisions fair? Why or why not?*

Concepts:

- Nozick: http://econ2.econ.iastate.edu/classes/econ362/hallam/Readings/Nozick_Justice.pdf
- Rawls: <https://eportfolios.macaulay.cuny.edu/thorne15/files/2015/03/Rawls-JUSTICE-AS-FAIRNESS.pdf>
- Walzer. Video: <https://bigthink.com/videos/michael-walzer-what-is-justice> Spheres of Justice. (full). <http://fs2.american.edu/dfagel/www/Philosophers/Walzer/SpheresofJustice-Chapters%201-3.pdf>

Class 6 Surveillance/Power and Shared Responsibility. ** Last day for Memo.

Data on where we go, who we see, and what we do is collected not only by social network sites and trackers on our websites. Location data aggregators collect detailed GPS location data on individuals and are able to track individual's behavior offline. Detailed data is not only used in advertising, alternative credit or trust scores, based on where we go and who we see, can be used to assess the trustworthiness of individuals who are not known to the company for employment, lending, education, or just matching with other users. Here we will analyze whether detailed location data should be used for trust scores and under what conditions companies should use trust scores to assess users, applicants, or employees. We will use David Lyon thoughts on surveillance to examine who is impacted by the use of these scores. People that study surveillance are concerned about the power dynamic between the watcher and the surveilled as well as how the surveilled may change their thoughts and behavior merely by thinking they could be surveilled.

Case – Location and Trust Scores

- Thompson, S. Warzel, C. Twelve Million Phones, One Dataset, Zero Privacy. 2019 NY Times: <https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html>

Concepts

- Pg. 57-67 Only. Lyon, D. From Big Brother to Electronic Panopticon. *The Electronic Eye: The Rise of Surveillance Society* (Minneapolis: University of Minnesota Press, 1994): 57-80. <https://home.fnal.gov/~annis/digirati/otherVoices/Lyon.html> or <https://www.jstor.org/stable/pdf/10.5749/j.cttsqw8.7.pdf?refreqid=excelsior%3A72520fa63646e674483b1b1918203b34>

1. *MEMO: What is worrisome about the collection or use of information in these articles? Why?*
2. *Should a company use a trust score for potential employees that tracks their location? for the promotion of current employees? Customers? Why or why not?*
3. *Who is impacted by the decision to use trust scores with location data?*

Extra Cases:

- Christopher Mims, The Secret Trust Scores Companies Use to Judge Us All, Wall Street Journal, April 6, 2019. <https://www.wsj.com/articles/the-secret-trust-scores-companies-use-to-judge-us-all-11554523206>
- Hill, K. 2020. The Secretive Company That Might End Privacy as We Know It. *The New York Times*

Within both employment and housing law, companies are responsible for both the disparate treatment (intentionally treating people differently based on special classes) as well as disparate impact (having an adverse impact on people based on their protected class). Protected classes can include race, color, religion, ethnicity, sex, national origin, disability status, as well as others. While disparate treatment may seem obvious (i.e. you can't treat someone differently just because they are female or because they are from a specific country you don't like), disparate impact is the use of a 'facially neutral' practice that has an unjustified adverse impact on members of a protected class. The easiest example is if someone had a height requirement for admittance into a school of 6' tall. While 'facially neutral' (it doesn't say anything about a protected class), the rule would have an adverse impact on women.

Here we need to better understand how algorithms, analytics, and AI could run into the rules around disparate treatment and impact since so many AI programs are used in areas where these laws apply.

Case: Hiring

- Amazon scraps secret AI recruiting tool that showed bias against women | Article. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G> .
- Heaven, Will Douglas. 2021. "Bias isn't the only problem with credit scores—and no, AI can't help." *MIT Technology Review*.
- FOR CLASS EXERCISE: 2019. Schools are using software to help pick who gets in. What could go wrong? Fast Company. <https://www.fastcompany.com/90342596/schools-are-quietly-turning-to-ai-to-help-pick-who-gets-in-what-could-go-wrong>

Concepts:

- Selection only. Barocas Solon, and Selbst Andrew D. 2016. "Big Data's Disparate Impact." *California Law Review*. P. 672-712.
1. *MEMO: Is hiring via algorithm fair? What could go wrong given Barocas and Selbst's paper?*
 2. *What would Walzer, Rawls, or Nozick think about Amazon's program?*
 3. *IN CLASS: If we were to design a project for admissions to college, how would you design it? We will work on this in class.*

Class 8 Economics of Trust Content Moderation

How much to modify in an automated system is value-laden; we decide when we override a programmed decision or when/where we insert human judgement into an AI enabled decision. The governance of automated systems, to trim back problematic outcomes, is a constant problem for companies whether curating videos, comments, user content, or music. Here we examine how to design a content moderating program that aligns with the goals of the firm. We will ask *which stakeholders' concerns matter in designing this program?*

Robert Frank is an economist at Cornell who argues (in the assigned reading) that acting in a trustworthy manner, taking into consideration stakeholders, being a responsible corporation, etc is economically advantageous. Such acts, if visible, allows stakeholders - such as employees, customers, suppliers, investors - to more easily find companies with matching values and allows others to make commitments faster and with fewer safeguards (because they trust you).

For us, we are concerned with how to assess data analytics where the incentives may not align with the long term value of the firm or with the interests of the company's stakeholders.

Case – Content Moderation

- CASE: Recommending an Insurrection
 1. *MEMO: Should Facebook do more to control their content? Why? Who should they take into consideration when making this decision?*
 2. *What should be the end goal of a recommender system? Of a content moderation system?*
 3. *Who should matter when making a recommender system? Content moderation system? Give an example of content and who should be considered.*

Concepts. Economics of Trust

- Frank, Robert H. "Can socially responsible firms survive in a competitive environment." Codes of Conduct: Behavioral Research into Business Ethics (Russell Sage Foundation, New York) (1996): 86-103.

Class 9 Transparency v accountability People Analytics

When decisions are deemed fair or legitimate, people care not only as to the outcomes of 'who-gets-what', but also the process by which the decision was made and how they were treated along the way. This is distributional (who gets what), procedural (how it was decided), and interactional fairness (how I was treated along the way). For decisions leveraging data analytics, explaining procedural and interactional fairness can be a challenge since explaining how the decision was made may be obscured by the designer of the program. However, much work has been done in CS to explore how AI-enabled decisions can be more transparent and how designers and users of data analytics tools should be held accountable. We will focus on two examples of people analytics (i.e., the analysis of employees, students, users based on their behavior as captured in data) and how the programs were designed for transparency, accountability, and (as we will read) 'contestability'.

Case. Uber/workers Workers AI

- Houston teachers to pursue lawsuit over secret evaluation system.
<https://www.houstonchronicle.com/news/houston-texas/houston/article/Houston-teachers-to-pursue-lawsuit-over-secret-11139692.php>

Concepts

- Hao, Karen. 2019. When algorithms mess up, the nearest human gets the blame. Technology Review. <https://www.technologyreview.com/2019/05/28/65748/ai-algorithms-liability-human-blame/>
- Mulligan, Deirdre K. and Kluttz, Daniel and Kohli, Nitin. Shaping Our Tools: Contestability as a Means to Promote Responsible Algorithmic Decision Making in the Professions (July 7, 2019). <https://www.cambridge.org/core/books/after-the-digital-tornado/shaping-our-tools-contestability-as-a-means-to-promote-responsible-algorithmic-decision-making-in-the-professions/311281626ECA50F156A1DDAE7A02CECB/core-reader>

1. *MEMO (choose one case): Do the teachers need to know about the use of AI or how it works?*
2. *Why are the teachers of AI upset? Could this have been avoided?*
3. *Where could you see the need for AI transparency? Accountability? Contestability?*

Extra:

- Labor Unions Work to Find Ways to Bargain With AI's Black Box - Bloomberg.
<https://www.bloomberg.com/news/articles/2020-11-30/labor-unions-work-to-find-ways-to-bargain-with-ai-s-black-box>

Class 10

Creating & Measuring Accuracy. Sheriff/school/parole

Measuring whether data analytics 'works' is not clear cut. The outcome being measured ('fitness score') has a relationship with the overall goal of the project ('good employee'); and that outcome is what is measured as being 'accurate' or not. We examine two related issues with 'accuracy' and AI programs: (1) whether enough information is gathered to measure 'accuracy' including false positives and false negatives; (2) whether is labeling the outcome variable (e.g., 'possible cheater'), we then put in place steps to make that outcome more likely ('caught cheater').

We first examine accuracy in regards to algorithms that may create poverty traps. We then examine how to measure accuracy with labeling students as potential criminals. Constructing the outcome of interest and how to measure accuracy is an important step within data analytics with moral implications.

Case:

- Pasco's sheriff uses grades and abuse histories to label schoolchildren potential criminals. <https://projects.tampabay.com/projects/2020/investigations/police-pasco-sheriff-targeted/school-data/>

Concepts:

- 2020 The problem with metrics is a big problem for AI. <https://www.fast.ai/2019/09/24/metrics/>

Questions:

1. *MEMO: Should schools provide student indicators for potential criminals to police? Why or why not?*

Class 11 Summary – The Everything Goes Wrong Case.

The cheating detection software used to proctor online exams exemplifies many of the issues we have identified in class. We will use this as a summary class – to show the many ways design of data analytics programs can go wrong.

- Students: Harwell. 2020. Washington Post. Cheating-detection companies made millions during the pandemic. Now students are fighting back.
<https://www.washingtonpost.com/technology/2020/11/12/test-monitoring-student-revolt/>

Questions:

1. *MEMO: Why are the students upset with the cheating detection program? Use any theory from the class to make your case.*

Class 12 Office Hours Only

Every company designates internal groups or entire divisions as responsible for researching the quality, safety, and societal impacts of their product or service. Tech companies have not settled on an approach as to who should be conducting research into the moral implications of data analytics - including how AI-enabled decisions may diminish rights, treat people unfairly, have a disparate impact on populations, or violates norms or even laws. We will use the case of Google Research and their work on the ethics of AI to examine who should be responsible for investigating the ethics of AI and data analytics more broadly. We will also skim a few alternative approaches, including in pharmaceuticals and manufacturing, to research in to the societal impacts of products and services more generally.

Case: Google Research

- CASE: Google Research: Who is Responsible for Ethics of AI?

Questions

1. *Should Google Research be filtering research conducted using their data and systems? Should the outcome of that research matter?*
2. *What obligation do corporations have to identify implications of their analytics initiatives, products, services? Why?*
3. *What is the role of ethics of AI and DA in an organization? Product development? Research?*

****Assigned day to do help with peer and online participation. These students will have a list of students in class and online to mark participation and will help give voice to those online in the classroom.**