Introduction to Data Science

Data Ethics - Algorithmic Bias

Joanna Bieri DATA101

Important Information

- Email: joanna_bieri@redlands.edu
- Office Hours: Duke 209

  Click Here for Joanna's Schedule

If you start having trouble with git!!!

Some people have reported that GIT is disappearing or giving errors on when they try to use it in Jupyter Lab.

Here is another option for interacting with git:

Git Desktop

Data Science Ethics - Algorithmic Bias

This lecture follows closely to Data Science in a Box Unit 3 - Deck 3. It has been updated to fit the prerequisites and interests of our class and translated to Python.

Data Science Ethics - Algorithmic Bias

In small groups:

- What do we mean about Algorithmic Bias
- Where are these algorithms being used?
- What are the human and societal effects?

How do you train yourself to make the right decisions (or reduce the likelihood of accidentally making the wrong decisions) at those points?

How do you respond when you see bias in someones work? How could you take action to educate others?

Where are your ethical lines?

The Hathaway Effect

We will start with a lighthearted example to just explain what we mean by algorithmic bias.

The company Berkshire Hathaway is owned by Warren Buffett. This analysis looked at certain time points in Anne Hathaway's career and compared it to Berkshire Hathaway's stock

The Hathaway Effect

- Oct. 3, 2008: Rachel Getting Married opens, BRK.A up 0.44%
- Jan. 5, 2009: Bride Wars opens, BRK.A up 2.61%
- Feb. 8, 2010: Valentine's Day opens, BRK.A up 1.01%
- March 5, 2010: Alice in Wonderland opens, BRK.A up 0.74%
- Nov. 24, 2010: Love and Other Drugs opens, BRK.A up 1.62%
- Nov. 29, 2010: Anne announced as co-host of the Oscars, BRK.A up 0.25%

The Hathaway Effect

Dan Mirvish. The Hathaway Effect: How Anne Gives Warren Buffett a Rise. The Huffington Post. 2 Mar 2011.

The Hathaway Effect

Dan Mirvish. The Hathaway Effect: How Anne Gives Warren Buffett a Rise. The Huffington Post. 2 Mar 2011.

Observational study - maybe people are searching for Anne Hathaway and then getting results for Berkshire Hathaway, could there be downstream effects. Read the article and you be the judge. Do you think there is a real effect here?

The Hathaway Effect

How does this illustrate Algorithmic Bias?

- The algorithms, a search engine and trading bots, ware trained to just return the best results for their purpose.
- The algorithms do not know anything about the difference between Anne and Berkshire.
- This could have unintended consequences. Anne's success leads to success for Berkshire (maybe?).
- BUT imagine how this could go really wrong.

Algorithmic bias and gender - Google Translate

A basic translator that can take sentences and translate them from one language to another. On the left are sentence fragments in Turkish and on the right the English translation.

It is having to choose a gender when translating. How did it do? Do you notice some things that are biased?

Algorithmic bias and gender - Google Translate

Turkish | Chinese | Spanish | Detect language — English | Chinese (Simplified) | Spanish | Translate

Turkish	English
o bir asker	he is a soldier
o bir öğretmen	She's a teacher
O bir doktor	He is a doctor
o bir hemşire	she is a nurse
o bir yazar	he is a writer
o bir kopek	he is a dog
o bir dadı	she is a nanny
o bir kedi	it is a cat
o bir rektör	he is a rector
o bir başkanı	he is a president
o bir girişimci	he is an entrepreneur
o bir Şarkıcı	she is a singer
o bir Öğrenci	he is a student
o bir Tercüman	he is a translator
o çalışkan	he is hard working
o tembel	she is lazy
o bir ressam	he is a painter
o bir kuaför	he is a hairdresser
o bir garson	he is a waiter
O bir mühendis	He is an engineer
o bir mimar	he is an architect

Algorithmic bias and gender - Amazon's experimental hiring algorithm

- Used AI to give job candidates scores ranging from one to five stars – much like shoppers rate products on Amazon
- Amazon's system was not rating candidates for software developer jobs and other technical posts in a gender-neutral way. - The system taught itself that male candidates were preferable.
- WHY? Because it was trained on past hiring decisions that were biased.

Algorithmic bias and gender - Amazon's experimental hiring algorithm

Gender bias was not the only issue. Problems with the data that underpinned the models' judgments meant that unqualified candidates were often recommended for all manner of jobs, the people said.

Jeffrey Dastin. Amazon scraps secret AI recruiting tool that showed bias against women.
Reuters. 10 Oct 2018.

Algorithmic bias and gender - Amazon's experimental hiring algorithm

- Algorithms can only pick up on features in the data.
- They do not consider the human behind the numbers.
- They are trained to minimize their "cost function".

Algorithmic bias and race - Facial recognition

Interview

'A white mask worked better': why algorithms are not colour blind

Ian Tucker

When Joy Buolamwini found that a robot recognised her face better when she wore a white mask, she knew a problem needed fixing

Sun 28 May 2017 13.27 BST

Joy Buolamwini is a graduate researcher at the MIT Media Lab and founder of the Algorithmic Justice League - an organisation that aims to challenge the biases in decision-making software. She grew up in Mississippi, gained a Rhodes scholarship, and she is also a Fulbright fellow, an Astronaut scholar and a Google Anita Borg scholar. Earlier this year she won a $50,000 scholarship funded by the makers of the film *Hidden Figures* for her work fighting coded

Algorithmic bias and race - Facial recognition

Ian Tucker. 'A white mask worked better': why algorithms are not colour blind. The Guardian. 28 May 2017.

- Joy Buolamwini graduate researcher at the MIT media lab and a leader in the Algorithmic Justice League.
- She noted that because the facial recognition algorithms were trained on mostly white faces they did not do a good job of recognizing patterns in non-white faces.
- When she put on a white mask with no human features the algorithm was better at tracking her face than when she wasn't wearing a mask.

Algorithmic bias and race - Criminal Sentencing

Software is being used across the country to predict future criminals. And it's biased against blacks.

Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine Bias. 23 May 2016. ProPublica.

Algorithmic bias and race - A tale of two convicts

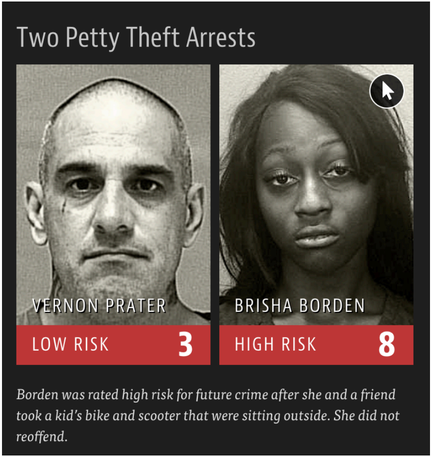

- Risk is ranking whether they will commit a crime again

Algorithmic bias and race - A tale of two convicts

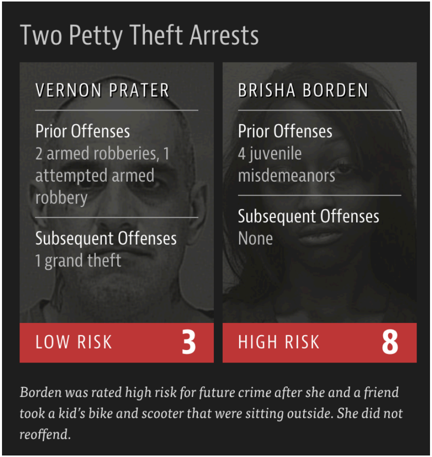

Borden was rated high risk for future crime after she and a friend took a kid's bike and scooter that were sitting outside. She did not reoffend.

Algorithmic bias and race - A tale of two convicts

Although these measures were crafted with the best of intentions, I am concerned that they inadvertently undermine our efforts to ensure individualized and equal justice," he said, adding, "they may exacerbate unwarranted and unjust disparities that are already far too common in our criminal justice system and in our society." - Then U.S. Attorney General Eric Holder (2014)

ProPublica analysis - Data:

Risk scores assigned to more than 7,000 people arrested in Broward County, Florida, in 2013 and 2014 + whether they were charged with new crimes over the next two years

ProPublica analysis - Results:

- 20% of those predicted to commit violent crimes actually did
- Algorithm had higher accuracy (61%) when full range of crimes taken into account (e.g. misdemeanors)

	WHITE	AFRICAN AMERICAN
Labeled Higher Risk, But Didn't Re-Offend	23.5%	44.9%
Labeled Lower Risk, Yet Did Re-Offend	47.7%	28.0%

ProPublica analysis - Results:

- Algorithm was more likely to falsely flag black defendants as future criminals, at almost twice the rate as white defendants
- White defendants were mislabeled as low risk more often than black defendants
- This effects peoples lives!!!
- Actual decisions are being made about the liberty and freedom of these people.

Data Ethics and Algorithms

What are our societal believes about these algorithms?

- Maybe people believe that if you take humans out of the decision making, the decisions will be MORE fair.
- Computers are not biased.
- It is better/faster to let the algorithm decide.

Data Ethics and Algorithms

What really happens behind the algorithm?

- Algorithms must be trained to make decisions, so we use data from our own society.
- Our society has a history of bias in multiple ways:
  - Who is represented in the data is highly dependent on existing bias and access in society.
  - Who collected or owns the data could make a big difference.
- Algorithms encode this bias and so biased decisions can come out.

Data Ethics in your Work

- At some point during your data science learning journey you will learn tools that can be used unethically
- You might also be tempted to use your knowledge in a way that is ethically questionable either because of business goals or for the pursuit of further knowledge (or because your boss told you to do so)