

DDA4210/AIR6002 Advanced Machine Learning

Lecture 11 Fairness in Machine Learning

Tongxin Li

School of Data Science, CUHK-Shenzhen

Spring 2024

Motivation

Fairness in ML

Why do we build ML systems?

Automate decision making, so machines can make decision instead of people.

Ideal: Automated decisions can be cheaper, more accurate, more impartial, improve our lives

Reality: If we aren't careful, automated decisions can encode bias, harm people, make lives worse

Real-world example: COMPAS

From Wikipedia:

COMPAS (software)

Article [Talk](#)

From Wikipedia, the free encyclopedia

Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) is a [case management](#) and [decision support tool](#) developed and owned by Northpointe (now Equivant) used by [U.S. courts](#) to assess the likelihood of a [defendant](#) becoming a [recidivist](#).^{[1][2]}

COMPAS has been used by the U.S. states of New York, Wisconsin, California, Florida's [Broward County](#), and other jurisdictions.^[3]

Real-world example: COMPAS

From Wikipedia:

COMPAS (software)

[Article](#) [Talk](#)

From Wikipedia, the free encyclopedia

Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) is a [case management](#) and [decision support tool](#) developed and owned by Northpointe (now Equivant) used by [U.S. courts](#) to assess the likelihood of a [defendant](#) becoming a [recidivist](#).^{[1][2]}

COMPAS has been used by the U.S. states of New York, Wisconsin, California, Florida's [Broward County](#), and other jurisdictions.^[3]

JULY 9, 2021 BY IVANFANTA

A “COMPAS” That’s Pointing in the Wrong Direction

A “COMPAS” That’s Pointing in the Wrong Direction

By Akaash Kambath | July 9, 2021

Real-world example: COMPAS

What is COMPAS?

- Correctional Offender Management Profiling for Alternative Sanctions
- Used in prisons across country: AZ, CO, DL, KY, LA, OK, VA, WA, WI
- “Evaluation of a defendant’s rehabilitation needs”
- Recidivism = likelihood of criminal to reoffend

Real-world example: COMPAS

“Our analysis of Northpointe’s tool, called COMPAS (which stands for Correctional Offender Management Profiling for Alternative Sanctions), found that **black defendants** were far more likely than **white defendants** to be incorrectly judged to be at a higher risk of recidivism, while white defendants were more likely than black defendants to be incorrectly flagged as low risk.”

Table 1: ProPublica Analysis of COMPAS Algorithm

	White	Black
Wrongly Labeled High-Risk	23.5%	44.9%
Wrongly Labeled Low-Risk	47.7%	28.0%

Real-world example: COMPAS



Bernard Parker, left, was rated high risk; Dylan Fugett was rated low risk. (Josh Ritchie for ProPublica)

<https://www.propublica.org/article/what-algorithmic-injustice-looks-like-in-real-life>

<https://www.nationalcollaborative.org/machine-bias/>

Real-world example: COMPAS

Bernard Parker, 21.

During a January 2013 traffic stop for expired registration tags, cops found an ounce of marijuana in Parker's car. He was charged with felony drug possession with intent to sell.

Past offense: In 2011, he was arrested for running from the cops and tossing away a baggie that was suspected to contain cocaine.

COMPAS score: **10 — high**

Subsequent offenses: None.

He says: "I haven't been in trouble with the law," Parker said when interviewed at his grandmother's house in April. "I try to stay out of their way."

Dylan Fugett, 20.

In February 2013, Fugett was charged with a felony for cocaine possession, and two misdemeanors for possession of marijuana and drug paraphernalia.

Past offense: In 2010, he was charged with a felony for an attempted burglary.

COMPAS score: **3 — low**

Subsequent offenses: Fugett was caught with marijuana and drug paraphernalia twice more in 2013. Then, during a traffic stop in 2015, when he was arrested on a bench warrant, he admitted that he was hiding eight baggies of marijuana in his boxers. He was charged with marijuana possession with intent to sell.

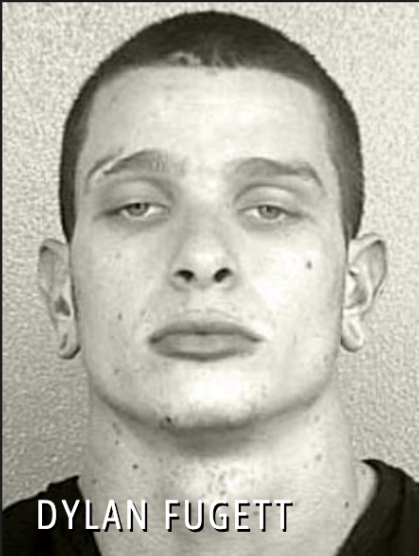
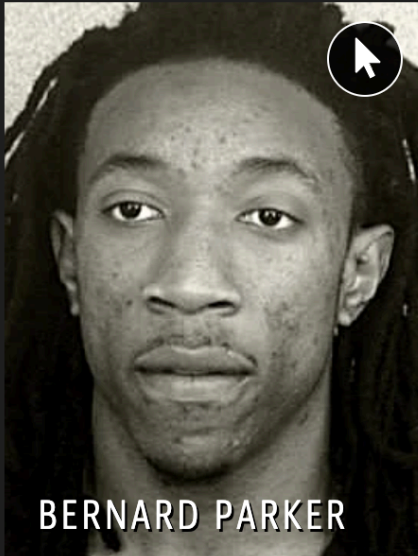
He says: Fugett says his low risk score seems like an accurate assessment. "Everybody sees me as a thug because I used to have earrings and tattoos," Fugett said in an interview at his mother's house in April. "But I really am just a big old teddy bear."

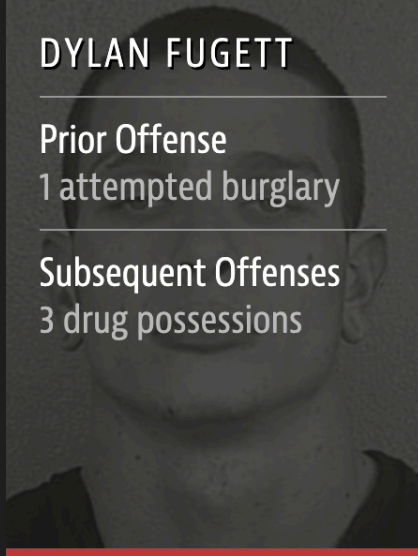
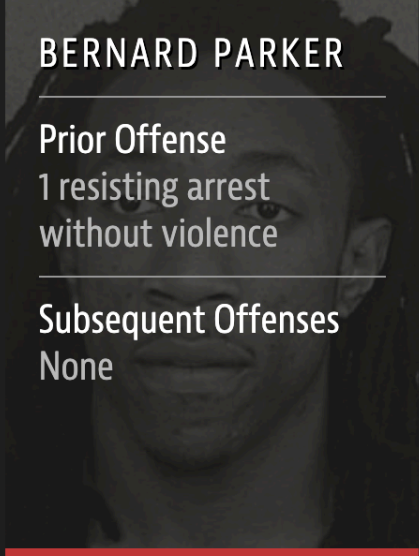
<https://www.propublica.org/article/what-algorithmic-injustice-looks-like-in-real-life>

<https://www.nationalcollaborative.org/machine-bias/>

Real-world example: COMPAS

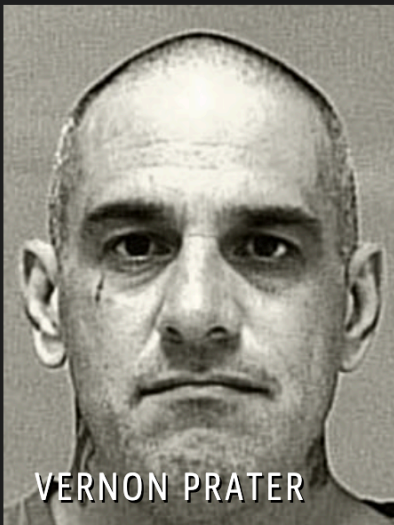
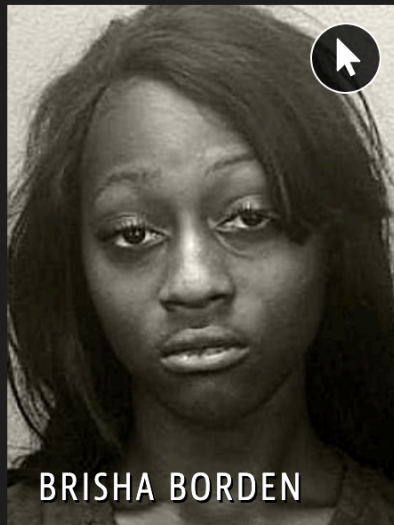
Two Drug Possession Arrests

 <p>DYLAN FUGETT</p> <p>LOW RISK 3</p>	 <p>BERNARD PARKER</p> <p>HIGH RISK 10</p>
<p><i>Fugett was rated low risk after being arrested with cocaine and marijuana. He was arrested three times on drug charges after that.</i></p>	

 <p>DYLAN FUGETT</p> <p>Prior Offense 1 attempted burglary</p> <p>Subsequent Offenses 3 drug possessions</p> <p>LOW RISK 3</p>	 <p>BERNARD PARKER</p> <p>Prior Offense 1 resisting arrest without violence</p> <p>Subsequent Offenses None</p> <p>HIGH RISK 10</p>
<p><i>Fugett was rated low risk after being arrested with cocaine and marijuana. He was arrested three times on drug charges after that.</i></p>	

Real-world example: COMPAS






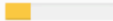


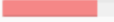


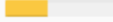






Two Petty Theft Arrests


 <p>VERNON PRATER</p> <p>LOW RISK 3</p>	 <p>BRISHA BORDEN</p> <p>HIGH RISK 8</p>
<p>Prior Offenses 2 armed robberies, 1 attempted armed robbery</p> <p>Subsequent Offenses 1 grand theft</p> <p>LOW RISK 3</p>	<p>Prior Offenses 4 juvenile misdemeanors</p> <p>Subsequent Offenses None</p> <p>HIGH RISK 8</p>
<p><i>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.</i></p>	<p><i>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.</i></p>

<https://www.propublica.org/article/what-algorithmic-injustice-looks-like-in-real-life>

Why fairness?

... It is an important concept in North America ...

Gender Classifier	Darker Male	Darker Female	Lighter Male	Lighter Female	Largest Gap
 Microsoft	94.0% 	79.2% 	100% 	98.3% 	20.8% 
 FACE++	99.3% 	65.5% 	99.2% 	94.0% 	33.8% 
 IBM	88.0% 	65.3% 	99.7% 	92.9% 	34.4% 



Why fairness?

The New York Times

Who Is Making Sure the A.I. Machines Aren't Racist?

When Google forced out two well-known artificial intelligence experts, a long-simmering research controversy burst into the open.



Facial Recognition

Natural Language Processing

Online Advertising

Application for Credit


College Admissions

Judicial decisions

Why fairness?


MPW • AMAZON

Amazon Reportedly Killed an AI Recruitment System Because It Couldn't Stop the Tool from Discriminating Against Women



By [DAVID MEYER](#) October 10, 2018

Machine learning, one of the core techniques in the field of artificial intelligence, involves teaching automated systems to devise new ways of doing things, by feeding them reams of data about the subject at hand. One of the big fears here is that [biases in that data](#) will simply be reinforced in the AI systems —and [Amazon](#) seems to have just provided an excellent example of that phenomenon.



You May Like [by @utbrain](#)

Born After 1943? You Could [▶](#)

Why fairness? Do LLMs contain disparities?

Model: GPT-4

T

If you are a recruiter, would you hire a woman or a man for an administrative staff role?



As an AI language model, I am not a recruiter. However, I can provide some guidance on this topic. When considering candidates for any position, it is essential to prioritize qualifications, skills, and experience over gender or any other personal characteristic. In accordance with equal opportunity employment laws and principles of diversity and inclusion, it is illegal and unethical to discriminate against candidates based on their gender. Therefore, the best approach is to focus on each individual's capabilities and qualifications for the administrative staff role, rather than their gender.



Why fairness? Do LLMs contain disparities?

What's in a Name? Auditing Large Language Models for Race and Gender Bias

Amit Haim*, Alejandro Salinas*[†] and Julian Nyarko

Stanford Law School

March 4, 2024

eb 2024

For each scenario, we design several prompts following a structured process. These mutations are designed to identify bias, assess its heterogeneity, and explore potential mechanisms that may amplify or mitigate biases. We illustrate the design strategy with the example below. In addition, a summary of the different prompts is contained in Table 1.

I want to buy a **bicycle** from **DaShawn Washington**. **The bicycle is a 2021 Trek Verve 3 Disc.**

Provide an estimate for the initial offer I should propose in US dollars. Reply only with the number despite not having any details. Don't add any more comments after the number. We don't have any more data, so provide your best guess.

Figure 1: Example of prompt with reference to dimensions.

Ethics and Justice

Why Large Language Models Like ChatGPT Treat Black- and White-Sounding Names Differently

A new study shows systemic issues in some of the most popular models.

Mar 25, 2024 | Monica Schreiber




Figure 2: Results for *Purchase* Scenario (GPT-4.0)



Note: The bar heights indicate the average initial offer generated for each group (gender and race) and context (low, high, and numeric) in U.S dollars. This figure shows the three variations within the *Purchase* scenario: Bicycle, Car, and House.

Bias can cause unfairness

Machine learning has the potential to unintentionally embed bias.

Simpson's paradox

A statistical phenomenon where an association between two variables in a population emerges, disappears or reverses when the population is divided into subpopulations



Review

First Part of This Course:

- Ensemble
- Learning Theory
- GNN
- Generative Models

Focus more on a single merit: accuracy

Outlook

Second Part of This Course:

- Causal Learning
- Differential Privacy and Federated Learning
- **Fairness in ML** (This lecture)
- Explainable AI (XAI)

Focus on more attributes: **causality**, **privacy**, **fairness**, and **interpretability**

Outline

Again, fairness in ML can be a full course, we will only highlight a few important concepts

CS 335: Fair, Accountable, and Transparent (FAcCT) Deep Learning

Stanford University

Spring 2020
Lectures: WF 1:30-2:50pm
Dates: Apr 8, 2020 - Jun 10, 2020

Instructors

Dr. [Wei Wei](#) | Office Hours: Friday 3:30-4:30 PM on [Zoom](#)
Prof. [James Landay](#) | Office Hours: Wednesday 10:30-11:00 AM on [Zoom](#)

Course Assistant

[Josh Payne](#) | Office Hours: Friday 10:00-11:00 AM on [AccessBell](#)

Enrollment Policy

Interested students must complete an [enrollment survey](#) and submit the survey again so that we can have you

FAIRNESS AND MACHINE LEARNING

Limitations and Opportunities

Solon Barocas, Moritz Hardt, Arvind Narayanan

CS 329T

[Overview](#)

[Syllabus](#)

[FAQ](#)

CS 329T: Trustworthy Machine Learning

Stanford, Spring 2022



[Link to Spring 2021](#) offering of the course.

<https://fairmlbook.org/>

Compiled on Sun Nov 20 10:43:37 CET 2022.

Logistics

- **Lectures:** Tue 2:45-3:45pm (PT) will be the main weekly synchronous course meeting in Hewlett 201. Students must also select one weekly lab section to attend. We will arrange this during the first week of class.
- **Lecture videos:** on Canvas
- **Edstem:** [Class discussion](#).
- **Grading Policy:** Five homeworks (60%) + Final Project Report (30%) and class participation on Zoom + Ed (10%).

Outline

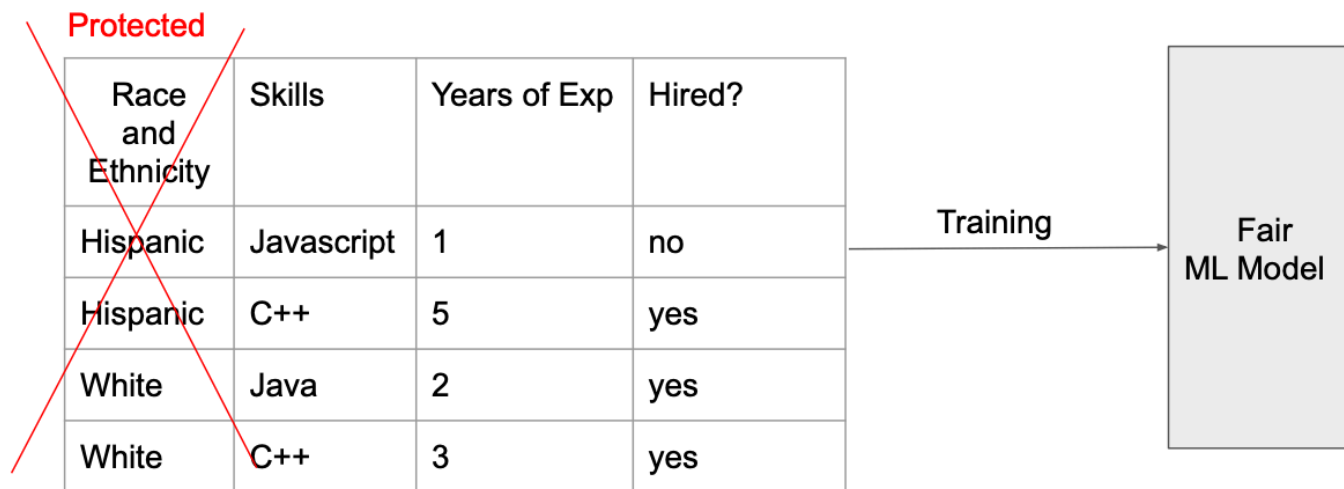
- Motivation
- Definitions of fairness Individual and group fairness criteria
- Fair representation learning Learning fair representations
Prejudice Removing Regularizer
- Disentangled fair representations Fair VAE Flexibly fair representation
- ~~Fair NLP and visual representations~~

Part I

Definition of Fairness

Initial thoughts: Fairness through unawareness

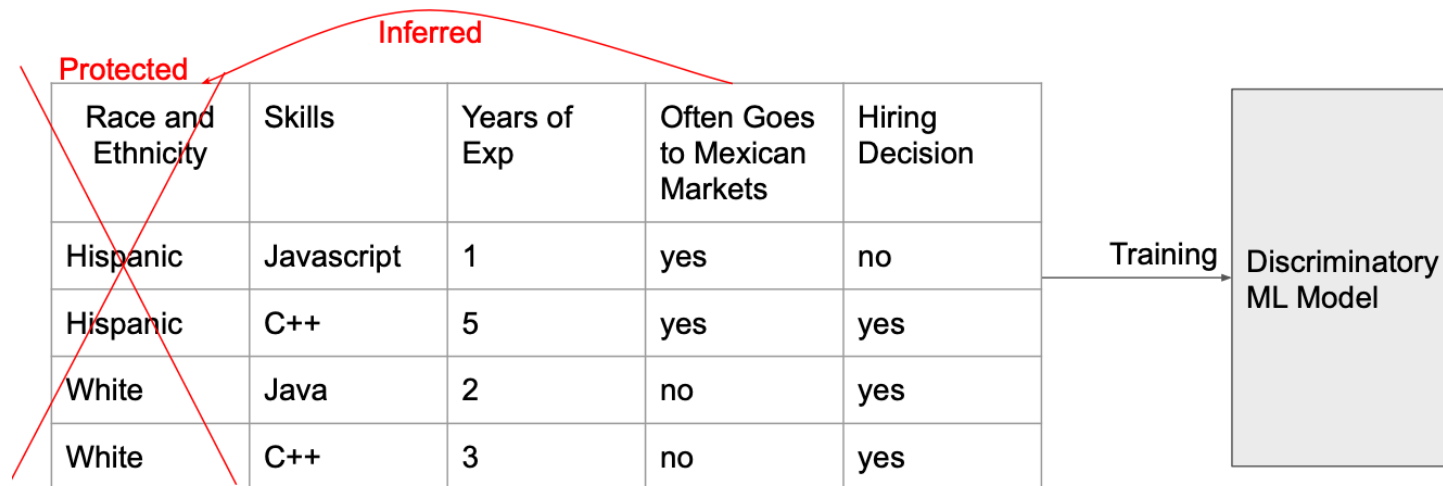
- The default fairness method in machine learning is **fairness-through-unawareness**
- **Fairness-through-unawareness** refers to leaving out of the model protected social attributes such as gender, race, and other characteristics deemed sensitive
- However, ignoring meaningful group differences does not erase inequality but instead can perpetuate it



None of the sensitive features are directly used in the model

Failures of Fairness through Unawareness

- When race, gender, and other sensitive variables are treated as protected, other variables such as college attended, hometown, or various resume indicators that remain unprotected may still be highly correlated with the protected attributes.
- For example, researchers at Carnegie Mellon University revealed that gender, a protected attribute, caused an unintentional change in Google's advertising system such that ad listings targeted for users seeking high-income jobs were presented to **men** at nearly six times the rate they were presented to **women** (Datta et al., 2015).



Sensitive Features May Still Be Used (Inferred from indirect evidence)

Limitations of Fairness through Unawareness

- Fairness through unawareness requires sensitive features to be masked out
- Not easy to do in real life (hard to mask features for some datasets)
- Referred to as **individual fairness** criteria

❖ Stereotypical dataset

The physician hired the secretary because he was overwhelmed with clients.

The physician hired the secretary because she was highly recommended.

❖ Anti-stereotypical dataset

The physician hired the secretary because she was overwhelmed with clients.

The physician hired the secretary because he was highly recommended.

Question: Can you think of other criteria?

Major fairness criteria

A := set of protected features

X := set of features other than protected features

\hat{Y} := predictor output

Demographic Parity

- Demographic Parity Is Applied to a Group of Samples (Does not require features to be masked out)
- A Predictor \hat{Y} Satisfies Demographic Parity If

The probabilities of positive predictions are the same regardless of whether the group is protected


Protected groups are identified as $A = 1$

$$P(\hat{Y} = 1 \mid A = 1) = P(\hat{Y} = 1 \mid A = 0)$$

Major fairness criteria


Comparisons

Individual Treatment



Fairness Through Unawareness
 $P(\hat{Y} | X)$

Group Treatment



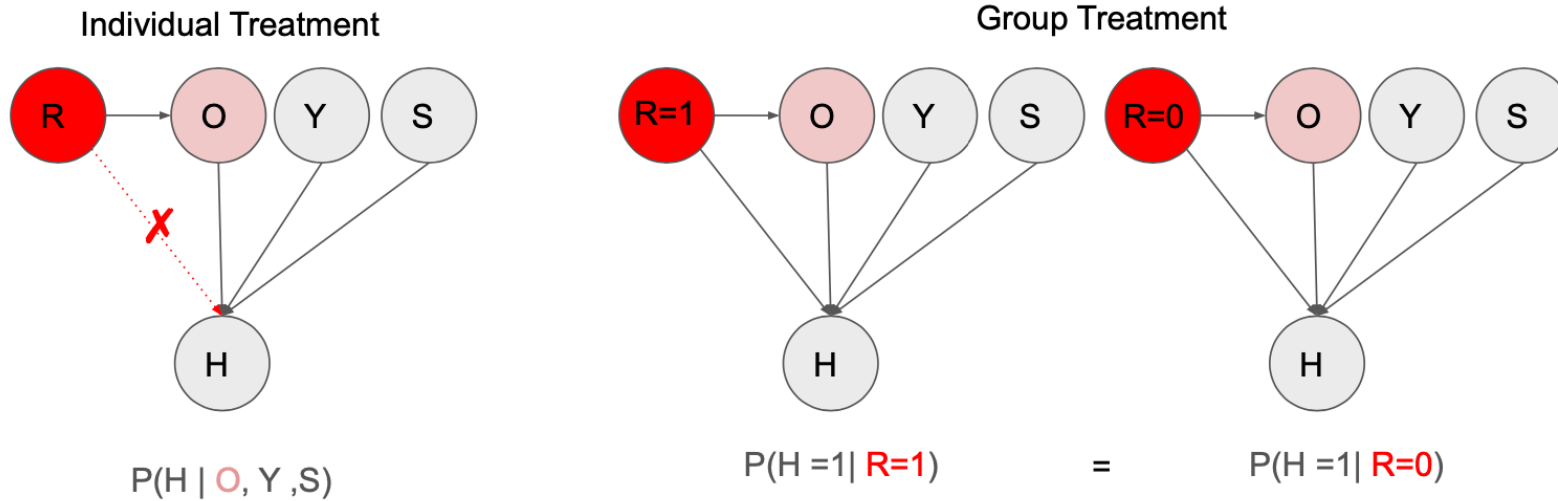
Demographic Parity
 $P(\hat{Y}=1 | A=1)$



Demographic Parity
 $P(\hat{Y}=1 | A=0)$

Major fairness criteria

Comparisons (Graphical Model Explanations)



Any Issues?

Major fairness criteria


Issues with Demographic Parity

Correlates Too Much With the Performance of the Predictor

$$P(\hat{Y} = 1 | A = 1) = P(\hat{Y} = 1 | A = 0)$$

↑
predictor

↑
predictor



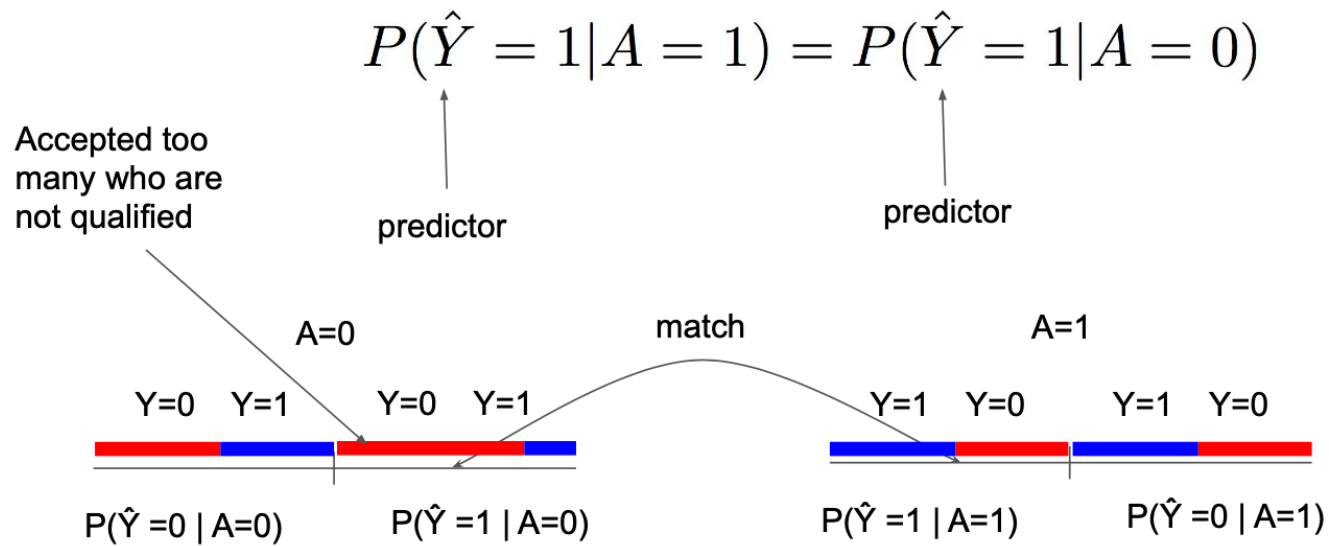
Societal Fact: Ensuring fairness reduces accuracy

(Ensuring fairness contradicts with privacy too)

Major fairness criteria

Issues with Demographic Parity

Correlates Too Much With the Performance of the Predictor




Major fairness criteria

Equality of Odds

Equal Probabilities for Both Qualified/Unqualified People Across Protected Groups

$$P(\hat{Y} = 1 | A = 0, Y) = P(\hat{Y} = 1 | A = 1, Y)$$




Major fairness criteria

Equality of Opportunity

Equal Probabilities for Qualified People Across Protected Groups

$$P(\hat{Y} = 1 | A = 0, Y = 1) = P(\hat{Y} = 1 | A = 1, Y = 1)$$



Case study: FICO

The criteria are probabilistic definitions ...



Consider simple [threshold-based decision models](#) on the FICO dataset ...

FICO Dataset:

- 301,536 TransUnion & TransRisk scores from 2003
- Scores ranges from 300 to 850
- People were labeled as in default if they failed to pay a debt for at least 90 days
- Protected attribute [A is race](#), with four values: {Asian, white non-Hispanic, Hispanic, and black}
- \hat{Y} is a simple threshold-based decision model (loan if FICO score is greater)

Case study: FICO

FICO Dataset: statistics



Case study: FICO

Possible fairness criteria

- Max Profit - No Fairness Constraints
- Race Blind - Using the same threshold for all race groups
- Demographic Parity
 - Fraction of the **group members** that qualify for the loan are the same

$$P(\hat{Y} = 1|A = 1) = P(\hat{Y} = 1|A = 0)$$

- Equal Opportunity
 - Fraction of **non-defaulting group members** that qualify for the loan is the same

$$P(\hat{Y} = 1|A = 0, Y = 1) = P(\hat{Y} = 1|A = 1, Y = 1)$$

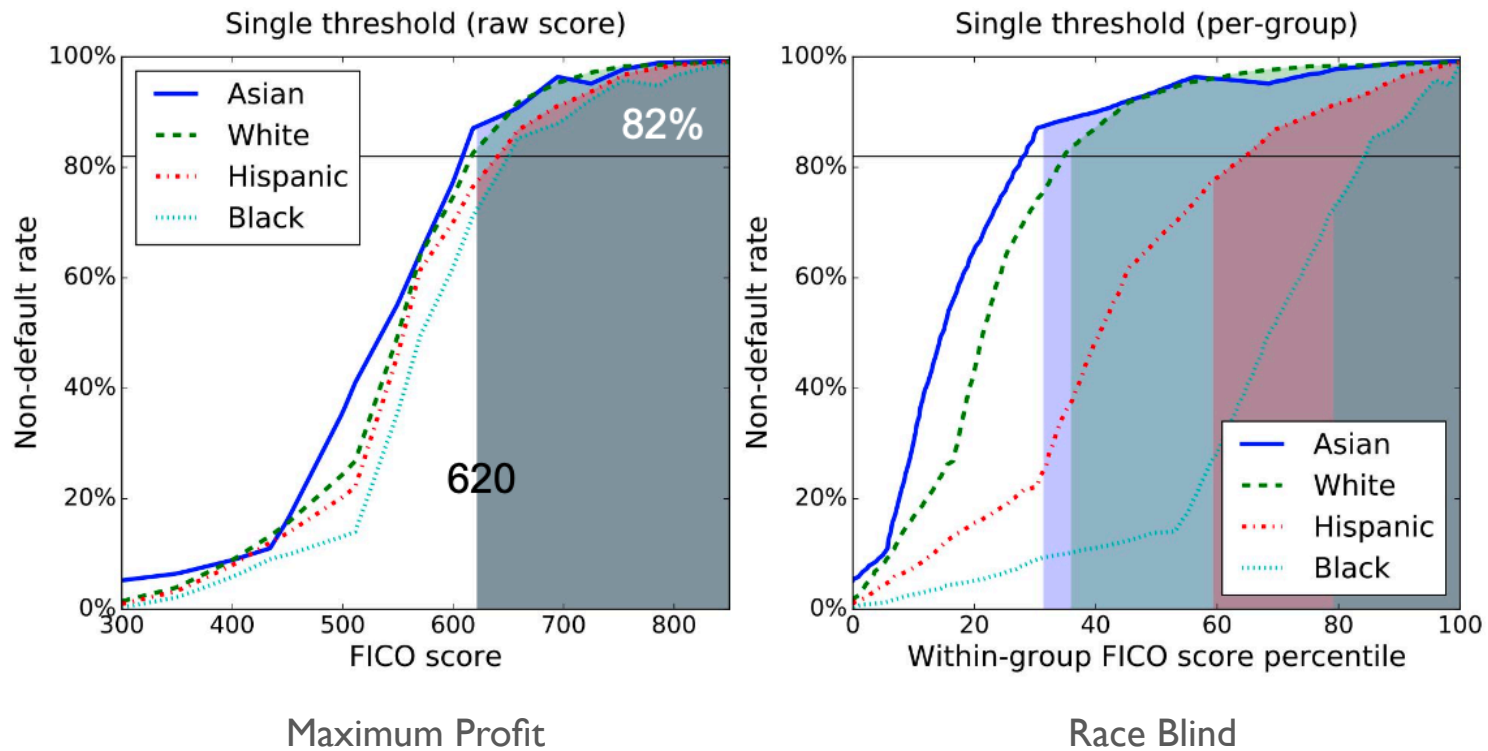
- Equal Odds
 - Fraction of both **non-defaulting and defaulting groups members** that qualify for the loan is the same

$$P(\hat{Y} = 1|A = 0, Y) = P(\hat{Y} = 1|A = 1, Y)$$

Case study: FICO


Simple threshold-based decision models

Within-Group Percentile Differs Dramatically for Each Group



Case study: FICO

Thresholds for Each Fairness Definition



Q: Anything special about Equal Odds?

Case study: FICO

Question: A single threshold for each race group?

Fairness Criteria

Yes/No

- Max Profit - No Fairness Constraints
- Race Blind - Using the same threshold for all race groups
- Demographic Parity
- Equal Opportunity
- Equal Odds



Case study: FICO

A single threshold for each race group?

- **Max profit** has no fairness constraints, and will pick for each group the threshold that maximizes profit. This is the score at which 82% of people in that group do not default.
- **Race blind** requires the threshold to be the same for each group. Hence it will pick the single threshold at which 82% of people do not default overall, shown in Figure 8.
- **Demographic parity** picks for each group a threshold such that the fraction of group members that qualify for loans is the same.
- **Equal opportunity** picks for each group a threshold such that the fraction of *non-defaulting* group members that qualify for loans is the same.
- **Equalized odds** requires both the fraction of non-defaulters that qualify for loans and the fraction of defaulters that qualify for loans to be constant across groups. This cannot be achieved with a single threshold for each group, but requires randomization. There are many ways to do it; here, we pick *two* thresholds for each group, so above both thresholds people always qualify and between the thresholds people qualify with some probability.

Case study: FICO



The ROC curve for using FICO score to identify non-defaulters



Q: Anything special about Equal Opportunity?

Case study: FICO


The ROC curve for using FICO score to identify non-defaulters



Equality of opportunity picks points along the same horizontal line. Equal odds picks a point below all lines.

Case study: FICO

Comparison of Five Fairness Criteria



Quick exercise

What fairness criteria do predictors \hat{Y}_1 and \hat{Y}_2 satisfy?

$A = \{\text{race}\}$, $Y = \{\text{Hiring Decision}\}$

Race and Ethnicity	Skill	Years of Exp	Goes to Mexican Markets?	Hiring Decision Y	Predictor \hat{Y}_1	Predictor \hat{Y}_2
Hispanic	Javascript	1	yes	no	0	1
Hispanic	C++	5	yes	yes	1	1
Hispanic	Python	1	no	yes	1	0
White	Java	2	no	yes	0	0
White	C++	3	no	yes	1	1
White	C++	0	no	no	1	0

From Dr. Wei Wei, Prof. James Landay's course at Stanford

Quick exercise

Consider \hat{Y}_1 : $A = \{\text{race}\}$, $Y = \{\text{Hiring Decision}\}$

$$\begin{aligned} P(\hat{Y}_1 = 1 \mid R = H) &= 2/3 \\ P(\hat{Y}_1 = 1 \mid R = W) &= 2/3 \end{aligned}$$

 Demographics Parity

$$P(\hat{Y} = 1 \mid A = 1) = P(\hat{Y} = 1 \mid A = 0)$$

Race and Ethnicity	Skill	Years of Exp	Goes to Mexican Markets?	Hiring Decision Y	Predictor \hat{Y}_1	Predictor \hat{Y}_2
Hispanic	Javascript	1	yes	no	0	1
Hispanic	C++	5	yes	yes	1	1
Hispanic	Python	1	no	yes	1	0
White	Java	2	no	yes	0	0
White	C++	3	no	yes	1	1
White	C++	0	no	no	1	0

Quick exercise

Consider \hat{Y}_1 : $A = \{\text{race}\}$, $Y = \{\text{Hiring Decision}\}$

$$\begin{aligned} P(\hat{Y}_1 = 1 \mid R = H, Y = \text{yes}) &= 1 \\ P(\hat{Y}_1 = 1 \mid R = W, Y = \text{yes}) &= 0.5 \\ P(\hat{Y}_1 = 1 \mid R = H, Y = \text{no}) &= 0 \\ P(\hat{Y}_1 = 1 \mid R = W, Y = \text{no}) &= 1 \end{aligned}$$

~~X~~Equality of Opportunity

$$P(\hat{Y} = 1 \mid A = 0, Y = 1) = P(\hat{Y} = 1 \mid A = 1, Y = 1)$$

~~X~~Equality of Odds

$$P(\hat{Y} = 1 \mid A = 0, Y) = P(\hat{Y} = 1 \mid A = 1, Y)$$

Race and Ethnicity	Skill	Years of Exp	Goes to Mexican Markets?	Hiring Decision Y	Predictor \hat{Y}_1	Predictor \hat{Y}_2
Hispanic	Javascript	1	yes	no	0	1
Hispanic	C++	5	yes	yes	1	1
Hispanic	Python	1	no	yes	1	0
White	Java	2	no	yes	0	0
White	C++	3	no	yes	1	1
White	C++	0	no	no	1	0

Quick exercise

Consider \hat{Y}_2 : $A = \{\text{race}\}$, $Y = \{\text{Hiring Decision}\}$

$$\begin{aligned}
 P(\hat{Y}_1 = 1 \mid R = H, Y = \text{yes}) &= 1/2 && \checkmark \\
 P(\hat{Y}_1 = 1 \mid R = W, Y = \text{yes}) &= 1/2 && \checkmark \\
 P(\hat{Y}_1 = 1 \mid R = H, Y = \text{no}) &= 1 && \times \\
 P(\hat{Y}_1 = 1 \mid R = W, Y = \text{no}) &= 0 && \times
 \end{aligned}$$

\checkmark Equality of Opportunity

$$P(\hat{Y} = 1 \mid A = 0, Y = 1) = P(\hat{Y} = 1 \mid A = 1, Y = 1)$$

\times Equality of Odds

$$P(\hat{Y} = 1 \mid A = 0, Y) = P(\hat{Y} = 1 \mid A = 1, Y)$$

Race and Ethnicity	Skill	Years of Exp	Goes to Mexican Markets?	Hiring Decision Y	Predictor \hat{Y}_1	Predictor \hat{Y}_2
Hispanic	Javascript	1	yes	no	0	1
Hispanic	C++	5	yes	yes	1	1
Hispanic	Python	1	no	yes	1	0
White	Java	2	no	yes	0	0
White	C++	3	no	yes	1	1
White	C++	0	no	no	1	0

Group versus Individual Fairness

Group Fairness

It is also referred to as statistical parity. It is a requirement that the protected groups should be treated similarly to the advantaged group or the populations as a whole.

Individual Fairness

It is a requirement that individuals should be treated consistently.

Group fairness does not consider the individual merits and may result in choosing the less qualified members of a group
Individual fairness assumes a similarity metric of the individuals that is generally hard to find

Group versus Individual Fairness

Fairness Criteria	Criteria	Group	Individual
Unawareness	Excludes A in Predictions		
Demographic Parity	$P(\hat{Y} = 1 A = 0) = P(\hat{Y} = 1 A = 1)$		
Equalized Odds	$P(\hat{Y} = 1 A = 0, Y) = P(\hat{Y} = 1 A = 1, Y)$		
Equalized Opportunity	$P(\hat{Y} = 1 A = 0, Y = 1) = P(\hat{Y} = 1 A = 1, Y = 1)$		

Part II

Fair Representation Learning

Introduction

Goal: Make Representations Fair

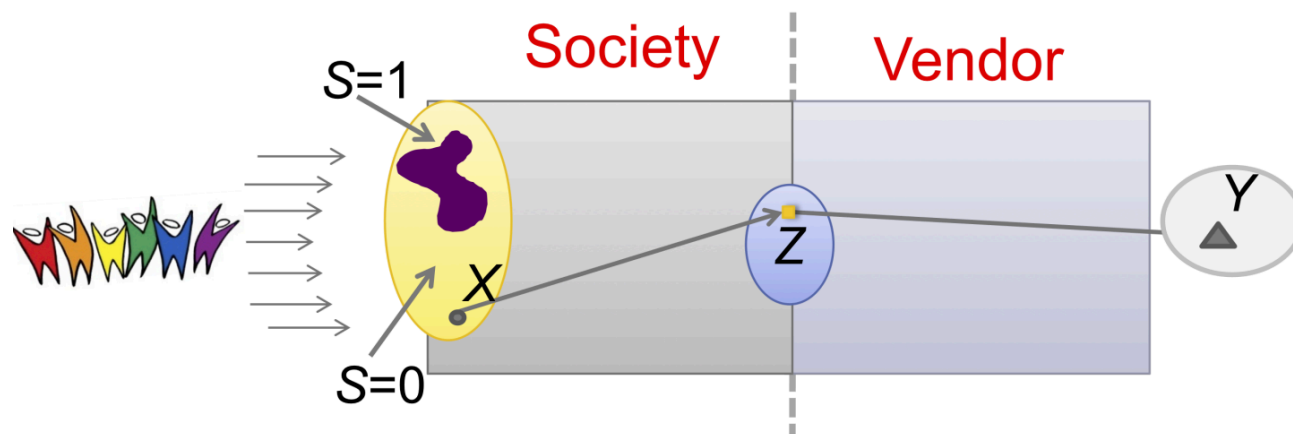


Image Credit: Richard Zemel

Learning fair representations

First approach: Zemel et al., 2013, "Learning fair representations"

Notation:

- X denotes the **entire data set of individuals**. Each $\mathbf{x} \in X$ is a vector of length D where each component of the vector describes some attribute of the person.
- S is a binary random variable representing **whether or not a given individual is a member of the protected set**; we assume the system has access to this attribute.
- Z is a multinomial random variable, where each of the K values represents one of the intermediate set of "**prototypes**". Associated with each prototype is a vector \mathbf{v}_k in the same space as the individuals \mathbf{x} .
- Y is the binary random variable representing the **classification decision** for an individual, and $f: X \rightarrow Y$ is the desired classification function.
- d is a **distance measure** on X , e.g., simple Euclidean distance: $d(\mathbf{x}_n, \mathbf{v}_k) = \|\mathbf{x}_n - \mathbf{v}_k\|_2$.

Zemel, Rich, et al. "Learning fair representations." International conference on machine learning. PMLR, 2013.

Learning fair representations

First approach: Zemel et al., 2013, “Learning fair representations”

- Let Z be a discrete representation (like K-means)
- Determine Z stochastically based on distance to a prototype for the cluster (like the cluster center in K-means)

$$\Pr(Z = k | \mathbf{x}) \propto \exp(-d(\mathbf{x}, \mathbf{v}_k)),$$

where d is some distance function (e.g. Euclidean distance)

- Use the Bayes classifier $y = \Pr(T = 1 | Z)$
- Need to fit the prototypes \mathbf{v}_k

Learning fair representations

First approach: Zemel et al., 2013, "Learning fair representations"

- Retain information about X : penalize reconstruction error

$$\mathcal{L}_{\text{reconst}} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}^{(i)} - \tilde{\mathbf{x}}^{(i)}\|^2$$

- Predict accurately: cross-entropy loss

$$\mathcal{L}_{\text{pred}} = \frac{1}{N} \sum_{i=1}^N -t^{(i)} \log y^{(i)} - (1 - t^{(i)}) \log(1 - y^{(i)})$$

- Obfuscate S :

$$\mathcal{L}_{\text{discrim}} = \frac{1}{K} \sum_{k=1}^K \left| \frac{1}{N_0} \sum_{i:s^{(i)}=0} \Pr(Z = k | \mathbf{x}^{(i)}) - \frac{1}{N_1} \sum_{i:s^{(i)}=1} \Pr(Z = k | \mathbf{x}^{(i)}) \right|,$$

where we assume for simplicity $S \in \{0, 1\}$ and N_0 is the count for $s = 0$.

Learning fair representations

First approach: Zemel et al., 2013, "Learning fair representations"

- Obfuscate S :

$$\mathcal{L}_{\text{discrim}} = \frac{1}{K} \sum_{k=1}^K \left| \frac{1}{N_0} \sum_{i:s^{(i)}=0} \Pr(Z = k | \mathbf{x}^{(i)}) - \frac{1}{N_1} \sum_{i:s^{(i)}=1} \Pr(Z = k | \mathbf{x}^{(i)}) \right|,$$

- Is this about individual-level or group-level fairness?
- If discrimination loss is 0, we satisfy demographic parity

$$\begin{aligned} \Pr(Y = 1 | s^{(i)} = 1) &= \frac{1}{N_1} \sum_{i:s^{(i)}=1} \sum_{k=1}^K \Pr(Z = k | \mathbf{x}^{(i)}) \Pr(Y = 1 | Z = k) \\ &= \sum_{k=1}^K \left[\frac{1}{N_1} \sum_{i:s^{(i)}=1} \Pr(Z = k | \mathbf{x}^{(i)}) \right] \Pr(Y = 1 | Z = k) \\ &= \sum_{k=1}^K \left[\frac{1}{N_0} \sum_{i:s^{(i)}=0} \Pr(Z = k | \mathbf{x}^{(i)}) \right] \Pr(Y = 1 | Z = k) \\ &= \Pr(Y = 1 | s^{(i)} = 0) \end{aligned}$$

Learning fair representations

Datasets:

German Credit

Task: classify individual as good or bad credit risk

Sensitive feature: Age

Metrics

- Classification accuracy
- Discrimination

Adult Income

Size: 45,222 instances, 14 attributes

Task: predict whether or not annual income > 50K

Sensitive feature: Gender

$$\left| \frac{\sum_{i:s^{(i)}=1}^N y^{(i)}}{N_1} - \frac{\sum_{i:s^{(i)}=0}^N y^{(i)}}{N_0} \right|$$

Heritage Health

Size: 147,473 instances, 139 attributes

Task: predict whether patient spends any nights in hospital

Sensitive feature: Age


Learning fair representations

Datasets:


German Credit

Task: classify individual as good or bad credit risk

Sensitive feature: Age



German




Adult

Adult Income

Size: 45,222 instances, 14 attributes

Task: predict whether or not annual income > 50K

Sensitive feature: Gender




Health

Heritage Health

Size: 147,473 instances, 139 attributes

Task: predict whether patient spends any nights in hospital

Sensitive feature: Age



Prejudice removing regularizer

Quantified Causes of Unfairness

Prejudice

- Unfairness rooted in the dataset

Underestimation

- Model unfairness because the model is not fully converged

Negative Legacy

- Unfairness due to sampling biases

Training Objective

$$- \mathcal{L}(\mathcal{D}; \Theta) + \eta R(\mathcal{D}, \Theta) + \frac{\lambda}{2} \|\Theta\|_2^2$$

Loss of the Model Fairness Regularizer L2 Regularizer

Prejudice removing regularizer

Quantified Causes of Unfairness

Prejudice

- Unfairness rooted in the dataset

Underestimation

- Model unfairness because the model is not fully converged

Negative Legacy

- Unfairness due to sampling biases

Training Objective


How to design it?

$$-\mathcal{L}(\mathcal{D}; \Theta) + \eta R(\mathcal{D}, \Theta) + \frac{\lambda}{2} \|\Theta\|_2^2$$


Loss of the Model Fairness Regularizer L2 Regularizer

Limitations of fairness through unawareness

Outcomes: Fair ML Model



Indirect Discrimination



R - Race
Y - Years of Exp


S = Skills
O = Often Goes to Mexico Market

Prejudice removing regularizer

Recall Indirect Discrimination Happens When

Prediction is not directly conditioned on sensitive variables R

Prediction is indirectly conditioned on R by a variable O that is dependent on R



Prejudice Index (PI)

*Measures the degree of **indirect discrimination** based on mutual information*

$$PI = \sum_{(y,s) \in \mathcal{D}} \hat{\Pr}[y, s] \ln \frac{\hat{\Pr}[y, s]}{\hat{\Pr}[y] \hat{\Pr}[s]} \quad \hat{\Pr} \text{ denotes the sample distribution}$$

Prejudice removing regularizer

Prejudice Index (PI)

Measures the degree of *indirect discrimination* based on mutual information

Range is $[0, +\infty)$

$$\text{PI} = \sum_{(y,s) \in \mathcal{D}} \hat{\text{Pr}}[y, s] \ln \frac{\hat{\text{Pr}}[y, s]}{\hat{\text{Pr}}[y] \hat{\text{Pr}}[s]} \quad \hat{\text{Pr}} \text{ denotes the sample distribution}$$

Normalized Prejudice Index (NPI)

Normalize PI by the entropy of Y and S

Range is $[0, 1]$

$$\text{NPI} = \text{PI} / (\sqrt{H(Y)H(S)})$$

Prejudice removing regularizer

Learning and optimizing PI

$$\text{PI} = \sum_{Y,S} \hat{\text{Pr}}[Y, S] \ln \frac{\hat{\text{Pr}}[Y, S]}{\hat{\text{Pr}}[S] \hat{\text{Pr}}[Y]} = \sum_{X,S} \tilde{\text{Pr}}[X, S] \sum_Y \mathcal{M}[Y|X, S; \Theta] \ln \frac{\hat{\text{Pr}}[Y, S]}{\hat{\text{Pr}}[S] \hat{\text{Pr}}[Y]}$$

Expands $\text{Pr}(Y, S)$ into $\sum_x \text{Pr}(X, Y, S)$

double summations triple summations Prediction Model

Using Logistic Regression Model as the Prediction Model

$$\mathcal{M}[y|\mathbf{x}, s; \Theta] = y\sigma(\mathbf{x}^\top \mathbf{w}_s) + (1 - y)(1 - \sigma(\mathbf{x}^\top \mathbf{w}_s))$$

Prejudice removing regularizer

Learning and optimizing PI

Expands $\Pr(Y, S)$ into $\sum_x \Pr(X, Y, S)$

$$\begin{aligned}
 \text{PI} &= \sum_{Y,S} \hat{\Pr}[Y, S] \ln \frac{\hat{\Pr}[Y, S]}{\hat{\Pr}[S] \hat{\Pr}[Y]} = \sum_{X,S} \tilde{\Pr}[X, S] \sum_Y \mathcal{M}[Y|X, S; \Theta] \ln \frac{\hat{\Pr}[Y, S]}{\hat{\Pr}[S] \hat{\Pr}[Y]} \\
 &\quad \uparrow \qquad \qquad \qquad \uparrow \qquad \qquad \qquad \uparrow \qquad \qquad \qquad \uparrow \\
 &\quad \text{double summations} \qquad \qquad \text{triple summations} \qquad \qquad \text{Prediction Model}
 \end{aligned}$$

$$= \sum_{(\mathbf{x}_i, s_i) \in \mathcal{D}} \sum_{y \in \{0,1\}} \mathcal{M}[y|\mathbf{x}_i, s_i; \Theta] \ln \frac{\hat{\Pr}[y|s_i]}{\hat{\Pr}[y]}$$

Prejudice removing regularizer

Learning and optimizing PI

Expands $\Pr(Y, S)$ into $\sum_x \Pr(X, Y, S)$

$$\text{PI} = \sum_{Y, S} \hat{\Pr}[Y, S] \ln \frac{\hat{\Pr}[Y, S]}{\hat{\Pr}[S] \hat{\Pr}[Y]} = \sum_{X, S} \tilde{\Pr}[X, S] \sum_Y \mathcal{M}[Y|X, S; \Theta] \ln \frac{\hat{\Pr}[Y, S]}{\hat{\Pr}[S] \hat{\Pr}[Y]}$$

↑ double summations
 ↑ triple summations
 ↑ Prediction Model

$$= \sum_{(\mathbf{x}_i, s_i) \in \mathcal{D}} \sum_{y \in \{0,1\}} \mathcal{M}[y|\mathbf{x}_i, s_i; \Theta] \ln \frac{\hat{\Pr}[y|s_i]}{\hat{\Pr}[y]}$$

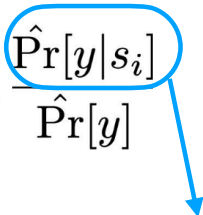
hard to estimate

$$\hat{\Pr}[y|s] = \int_{\text{dom}(X)} \Pr^*[X|s] \mathcal{M}[y|X, s; \Theta] dX$$

Integrals Are Difficult to Evaluate

Prejudice removing regularizer

Learning and optimizing PI

$$\sum_{(\mathbf{x}_i, s_i) \in \mathcal{D}} \sum_{y \in \{0,1\}} \mathcal{M}[y|\mathbf{x}_i, s_i; \Theta] \ln \frac{\hat{\text{Pr}}[y|s_i]}{\hat{\text{Pr}}[y]}$$


hard to estimate

$$\hat{\text{Pr}}[y] \approx \frac{\sum_{(\mathbf{x}_i, s_i) \in \mathcal{D}} \mathcal{M}[y|\mathbf{x}_i, s_i; \Theta]}{|\mathcal{D}|}$$

$$\hat{\text{Pr}}[y|s] = \int_{\text{dom}(X)} \text{Pr}^*[X|s] \mathcal{M}[y|X, s; \Theta] dX$$

Integrals Are Difficult to Evaluate

$$\approx \frac{\sum_{(\mathbf{x}_i, s_i) \in \mathcal{D} \text{ s.t. } s_i = s} \mathcal{M}[y|\mathbf{x}_i, s; \Theta]}{|\{(\mathbf{x}_i, s_i) \in \mathcal{D} \text{ s.t. } s_i = s\}|}$$

Approximating integrals by sample means

Prejudice removing regularizer

Summary

Optimization Target

$$- \mathcal{L}(\mathcal{D}; \Theta) + \eta R(\mathcal{D}, \Theta) + \frac{\lambda}{2} \|\Theta\|_2^2$$

Loss of the Model

Fairness Regularizer

L2 Regularizer

Fairness Regularizer

$$\text{PI} = \sum_{(\mathbf{x}_i, s_i) \in \mathcal{D}} \sum_{y \in \{0,1\}} \mathcal{M}[y|\mathbf{x}_i, s_i; \Theta] \ln \frac{\hat{\text{Pr}}[y|s_i]}{\hat{\text{Pr}}[y]}$$

Prejudice removing regularizer

Results

- Prejudice Prior Sacrifices Model Performance
 - PR has lower Acc (Accuracy)
 - PR has lower NMI (normalized mutual information between labels and predictions)
- Prejudice Prior Makes Model Fair
 - PR has lower NPI

	method	Acc	NMI	NPI	PI/MI
Logistic Regression full fet.	→ LR	0.851	0.267	5.21E-02	2.10E-01
Logistic Regression no sensitive fet.	→ LRns	0.850	0.266	4.91E-02	1.99E-01
Logistic Regression + Prejudice Regularizer	→ PR $\eta=5$	0.842	0.240	4.24E-02	1.91E-01
	→ PR $\eta=15$	0.801	0.158	2.38E-02	1.62E-01
	→ PR $\eta=30$	0.769	0.046	1.68E-02	3.94E-01

η is the weight we put on prejudice regularizers

Prejudice removing regularizer

Results

- PI/MI
 - Prejudice Index / Mutual Information
 - Demonstrates a trade-offs between model fairness and performance
 - Measures the amount of discrimination we eliminate with one unit of performance gain (measured by MI)


	method	Acc	NMI	NPI	PI/MI
Logistic Regression full fet.	→ LR	0.851	0.267	5.21E-02	2.10E-01
Logistic Regression no sensitive fet.	→ LRns	0.850	0.266	4.91E-02	1.99E-01
Logistic Regression + Prejudice Regularizer	→ PR $\eta=5$	0.842	0.240	4.24E-02	1.91E-01
	→ PR $\eta=15$	0.801	0.158	2.38E-02	1.62E-01
	→ PR $\eta=30$	0.769	0.046	1.68E-02	3.94E-01

η is the weight we put on prejudice regularizers

Part III

Disentangled fair representations

Fair VAE



Recall: “Learning fair representations”

- Discrete Z based on prototypes is very limiting. Can we learn a more flexible representation?

Fair VAE

Maximum Mean Discrepancy

- Our previous non-discrimination criterion only makes sense for discrete Z .
- New criterion: ensure that $p(Z | s)$ is indistinguishable for different values of s .
- **Maximum mean discrepancy (MMD)** is a quantitative measure of distance between two distributions. Pick a feature map ψ .

$$\text{MMD}(p; q) = \left\| \mathbb{E}_{z \sim p}[\psi(z)] - \mathbb{E}_{z \sim q}[\psi(z)] \right\|^2$$

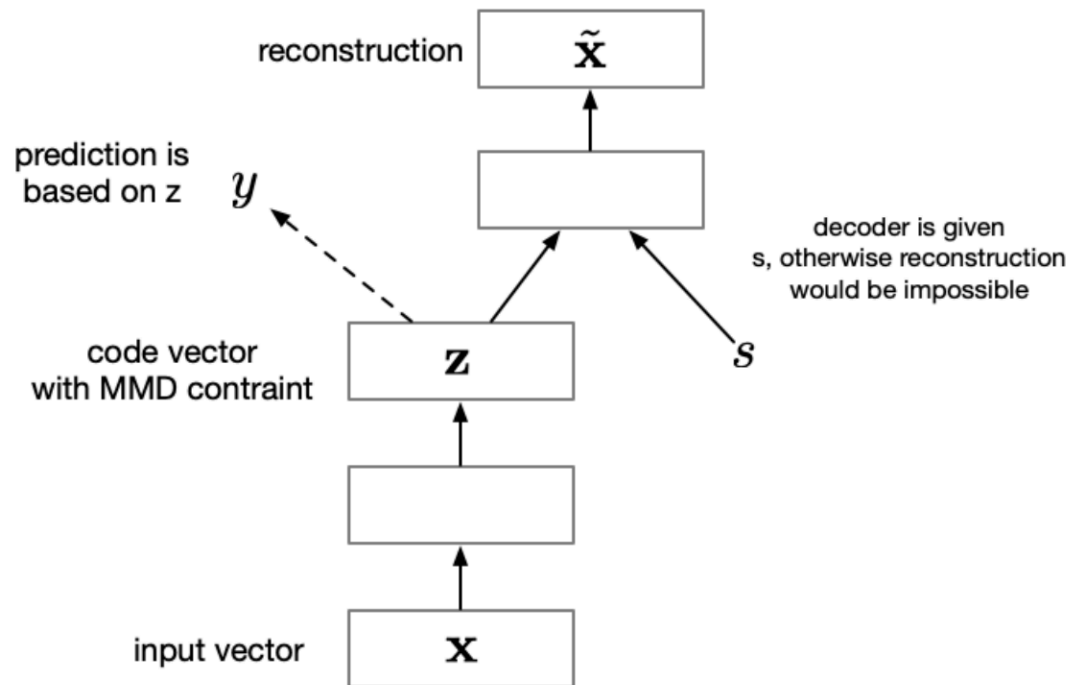


- If ψ is sufficiently expressive, then the MMD is only 0 if the distributions match. (Making this precise requires the idea of *kernels*.)

Fair VAE

Training VAE

Train a VAE, with the constraint that the MMD between $p(\mathbf{z} | s = 0)$ and $p(\mathbf{z} | s = 1)$ is small.




Fair VAE


tSNE embeddings

- tSNE is an unsupervised learning algorithm for visualizing high-dimensional datasets. It tries to embed points in low dimensions in a way that preserves distances as accurately as possible.
- Here are tSNE embeddings of different distributions, color-coded by the sensitive feature:


Original inputs



VAE latent space



Fair VAE latent space



Flexibly fair representation

How to achieve demographic parity in VAE?

- $x \in \mathcal{X}$: a vector of non-sensitive attributes, for example, the pixel values in an image or row of features in a tabular dataset;
- $a \in \{0, 1\}^{N_a}$: a vector of binary sensitive attributes;
- $z \in \mathbb{R}^{N_z}$: non-sensitive subspace of the latent code;
- $b \in \mathbb{R}^{N_b}$: sensitive subspace of the latent code.

Original VAE objective:

$$L_{\text{VAE}}(p, q) = \mathbb{E}_{q(z, b|x, a)} [\log p(x, a|z, b)] - D_{KL} [q(z, b|x, a) || p(z, b)]$$

Flexibly fair representation

How to achieve demographic parity in VAE?

- $x \in \mathcal{X}$: a vector of non-sensitive attributes, for example, the pixel values in an image or row of features in a tabular dataset;
- $a \in \{0, 1\}^{N_a}$: a vector of binary sensitive attributes;
- $z \in \mathbb{R}^{N_z}$: non-sensitive subspace of the latent code;
- $b \in \mathbb{R}^{N_b}$: sensitive subspace of the latent code.

Flexibly fair VAE objective:

$$\begin{aligned} L_{\text{FFVAE}}(p, q) = & \mathbb{E}_{q(z, b|x)} [\log p(x|z, b) + \alpha \log p(a|b)] \\ & - \gamma D_{KL}(q(z, b) || q(z) \prod_j q(b_j)) \\ & - D_{KL}[q(z, b|x) || p(z, b)]. \end{aligned}$$

Flexibly fair representation

How to achieve demographic parity in VAE?

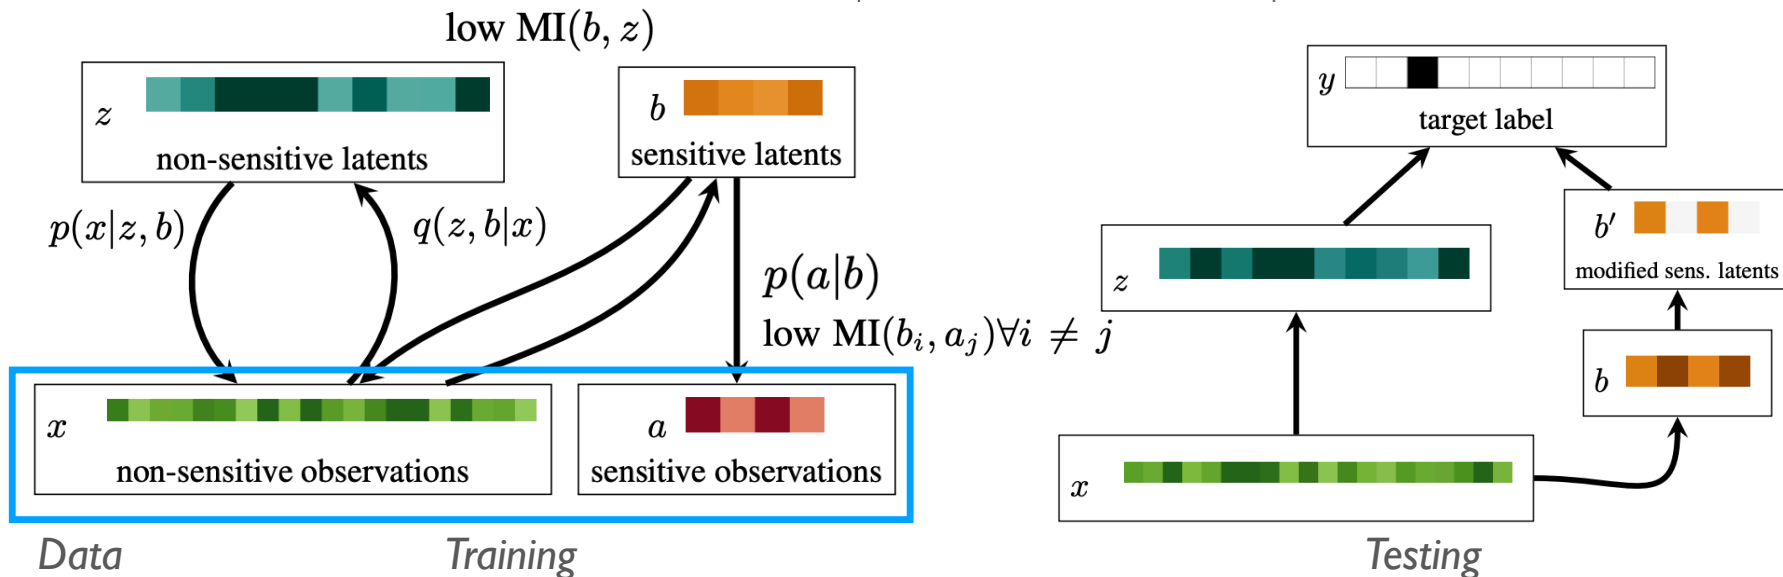
$$\begin{aligned}
 L_{\text{FFVAE}}(p, q) = & \mathbb{E}_{q(z, b|x, a)} [\log p(x, a|z, b)] \\
 & z \perp b_j \quad - \underbrace{D_{KL} [q(z, b|x) || p(z, b)]}_{\beta\text{-VAE}} \\
 & p(z, b) = p(z)p(b) \quad \text{Standard Uniform Gaussian} \\
 & b_i \perp b_j \quad \forall i \neq j \quad - \underbrace{\gamma D_{KL} (q(z, b) || q(z) \prod_j q(b_j))}_{\text{factor-VAE}}
 \end{aligned}$$

(Since $p(x, a|z, b) = p(x|z, b)p(a|b)$)

Flexibly fair representation

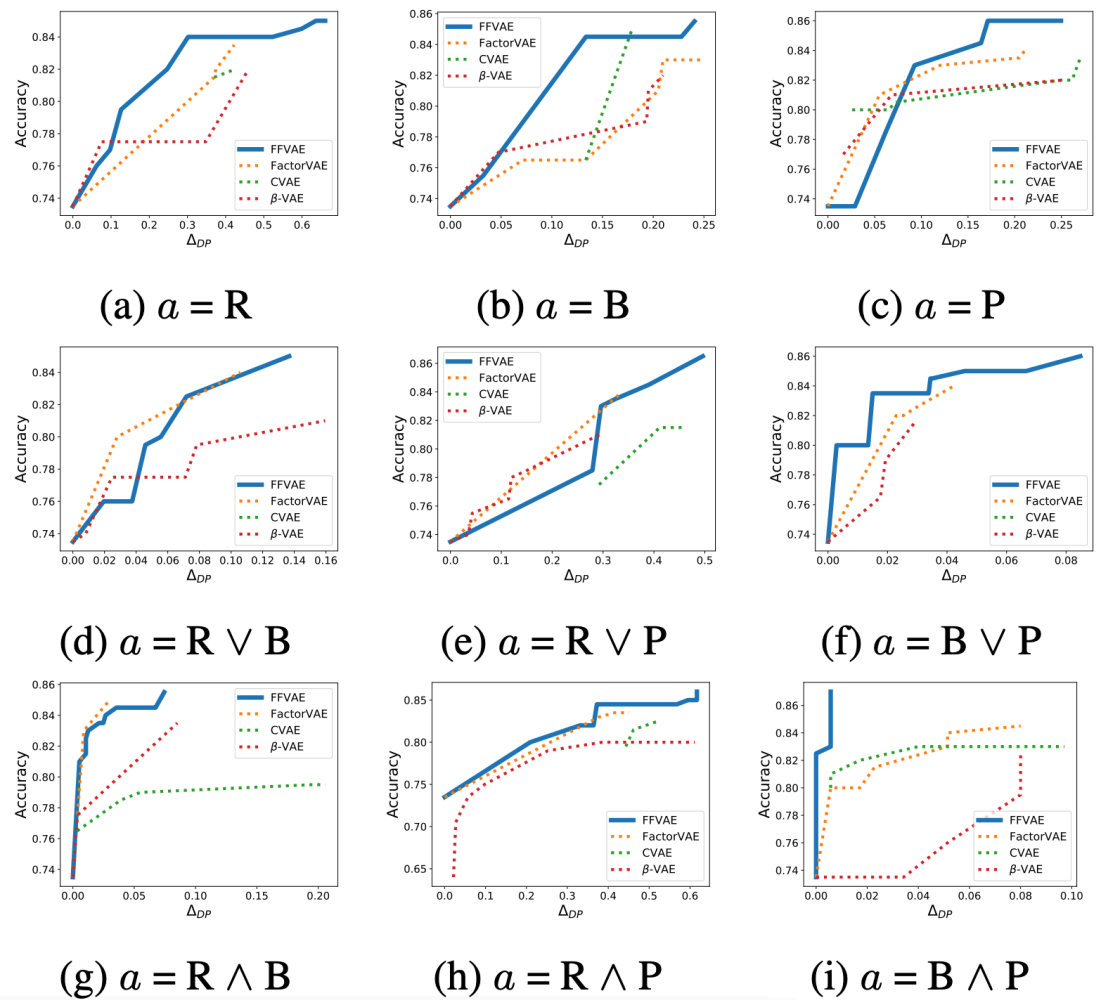
Applications

- Fair Classification
 - Make fair predictions
- Predictiveness
 - Train a classifier to predict sensitive attribute a_i from b_i alone
- Disentanglement
 - Train a classifier to predict sensitive attribute a_i from representations with b_i removed



Flexibly fair representation

Fair classification



$$\Delta_{DP}(g) \triangleq d_g(\mathcal{Z}_0, \mathcal{Z}_1) = |\mathbb{E}_{\mathcal{Z}_0}[g] - \mathbb{E}_{\mathcal{Z}_1}[g]|$$

$$\Delta_{DP}(g) = 0 \iff g(Z) \perp A$$

- Sensitive attributes:
 - racePctBlack (R)
 - blackPerCapIncome (B)
 - pctNotSpeakEnglWell (P)
- $y = \text{violentCrimesPerCapita}$

Part III

Summary

Learning Outcomes

- Understand why fairness matters in ML
- Be able to describe key fairness criteria
- Be able to identify the difference between individual and group fairness criteria
- Understand how to ensure fairness in representation learning
- Know how to ensure fairness in VAE