



Fairness in ML

A general introduction about Fairness in Algorithmic ML

Adrián Arnaiz-Rodríguez

1y PhD Student

ELLIS Alicante

Universidad de Alicante

17 January 2022

- 
- › Introduction to Algorithmic Fairness
 - › Bias
 - › Fairness definitions
 - › Limitations in definitions
 - › Imposing Fairness
 - › Current prominent approaches
 - › Datasets
 - › Libraries
 - › History and conceptual concerns
 - › General conclusions

Algorithmic bias problem and fairness at a glance

ML is used for critical decision making

How bias appears in society:

- Sources of bias
- Examples of bias



Challenges of AI

- Uncover bias/unfairness
- Measure bias (definitions Fairness)
- Mitigate bias
- Real world applications

How do we formulate the bias-fairness problem in every problem set up?

How do we detect the bias and how to solve it?

How could we define and measure bias or fairness?

Which are the ethical principles that follows each definition of bias and fairness?

Which are the implications in the real-world problems and, specifically in our own value system?




What is fairness for you?



Justice, equality and equity

Inequality

Unequal access to opportunities



With apologies to Shel Silverstein from @lunchbreath

2019 Design In Tech Report | Addressing Imbalance

Equity

Custom tools that identify and address inequality




With apologies to Shel Silverstein from @lunchbreath

2019 Design In Tech Report | Addressing Imbalance

Equality?

Evenly distributed tools and assistance




With apologies to Shel Silverstein from @lunchbreath

2019 Design In Tech Report | Addressing Imbalance

Justice

Fixing the system to offer equal access to both tools and opportunities



With apologies to Shel Silverstein from @lunchbreath

2019 Design In Tech Report | Addressing Imbalance






Introduction to Algorithmic Fairness

ML for critical decision making

- ML models are becoming the main tools for addressing complex societal problems in many consequential areas of our lives

- Education
- Justice: pretrial and detention
- Security
- Health
- Child Maltreatment screening
- Social Services
- Hiring
- Finance
- Advertising



- Each one with its own objectives

- Reduce cost
- Maximize social benefit
- ...

✓ Privacy	✓ Reliability
✓ Transparency	✓ Autonomy
✓ Accountability	✓ Fairness

Ethical implications

Many of these concepts do not have universally accepted definitions



Harms from Algorithmic Decision-Making





INDIVIDUAL HARMS			COLLECTIVE SOCIAL HARMS
ILLEGAL DISCRIMINATION	UNFAIR PRACTICES		
HIRING			LOSS OF OPPORTUNITY
EMPLOYMENT			
INSURANCE & SOCIAL BENEFITS			
HOUSING			
EDUCATION			
CREDIT			ECONOMIC LOSS
DIFFERENTIAL PRICES OF GOODS			
LOSS OF LIBERTY			SOCIAL STIGMATIZATION
INCREASED SURVEILLANCE			
STEREOTYPE REINFORCEMENT			
DIGNATORY HARMS			

Chart Contents Courtesy of Megan Smith, Former CTO of the United States

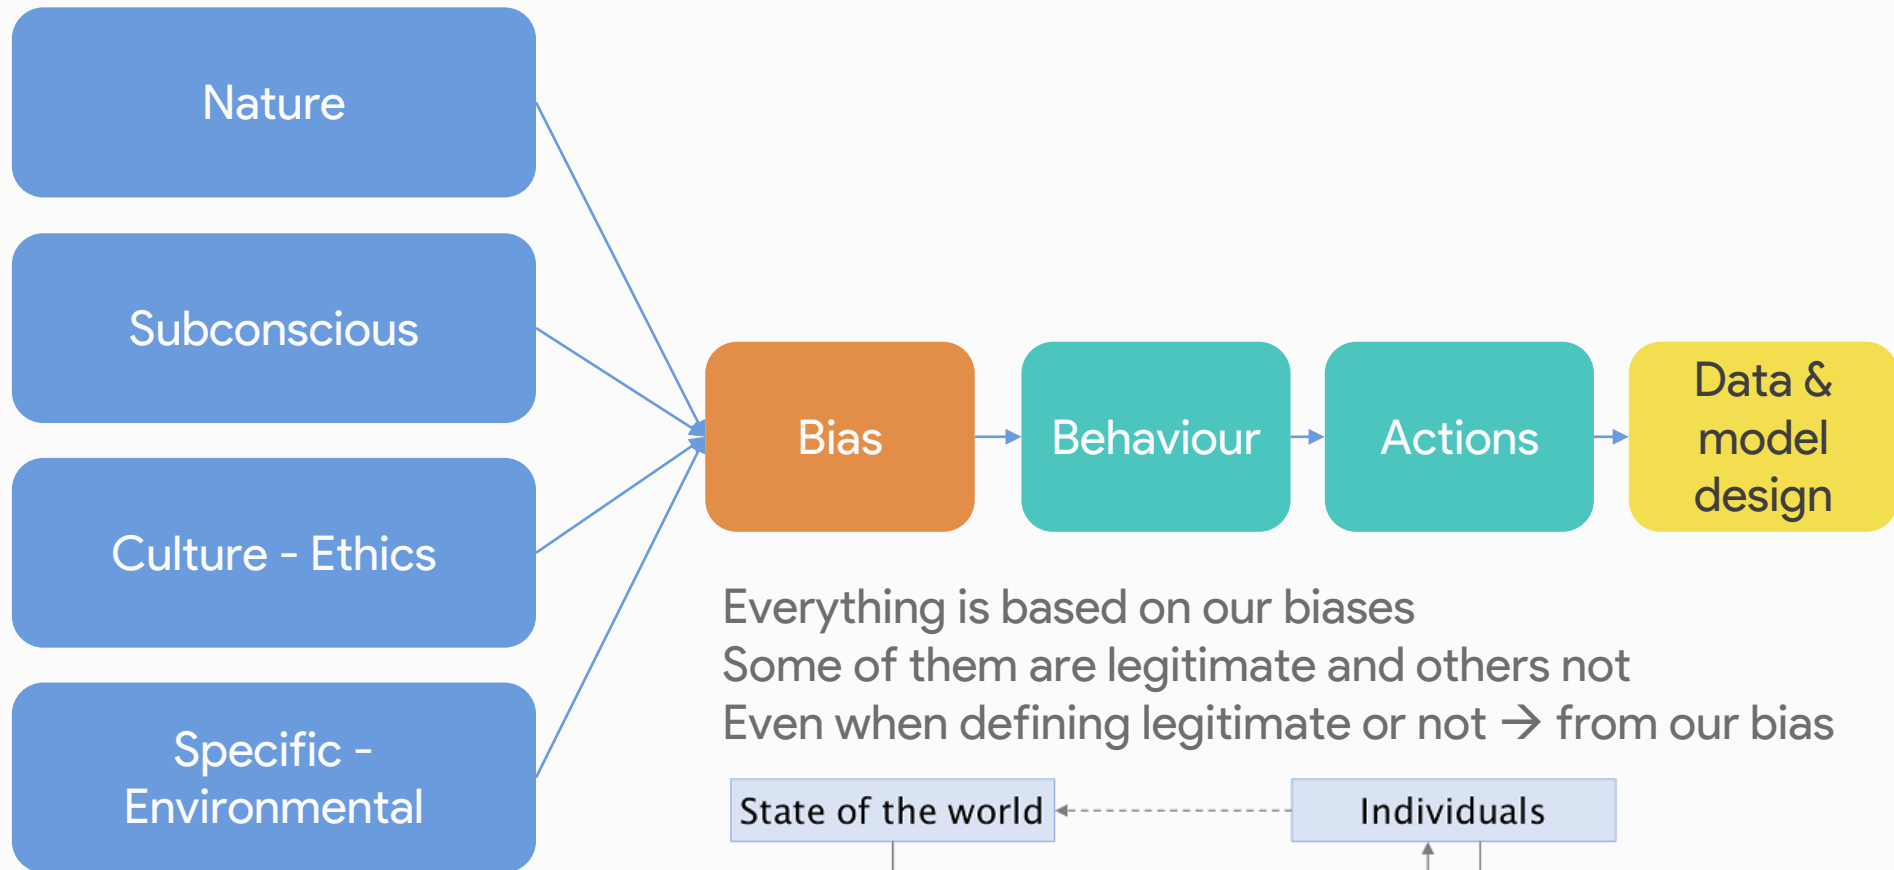


ML for critical decision making - examples


- Finance
 - *A. Byanjankar, M. Heikkilä, and J. Mezei. Predicting credit risk in peer-to-peer lending: A neural network approach. In IEEE Symposium Series on Computational Intelligence, 2015*
- Hiring
 - *M. Bogen and A. Rieke. Help wanted: An examination of hiring algorithms, equity, and bias. Technical report, Upturn, 2018*
- Pretrial and detention
 - *J. Angwin, J. Larson, S. Mattu, and L. Kirchner. Machine Bias: There's software used across the country to predict future criminals. And it's biased against blacks., 2016.*
- Child maltreatment screening
 - *A. Chouldechova, E. Putnam-Hornstein, D. Benavides-Prado, O. Fialko, and R. Vaithianathan. A case study of algorithmassisted decision making in child maltreatment hotline screening decisions. In Proceedings of the 1st Conference on Fairness, Accountability and Transparency, pages 134–148, 2018.*
- Education
 - *L. Oneto, A. Siri, G. Luria, and D. Anguita. Dropout prediction at university of genoa: a privacy preserving data driven approach. In European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, 2017.*
- Social Services
 - *V. Eubanks. Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor. St. Martin's Press, 2018*



Bias is implicit in every decision we make



Everything is based on our biases
Some of them are legitimate and others not
Even when defining legitimate or not → from our bias



Human centric ML approaches

AI systems learning moral notions

*AI-based systems can **learn moral notions** or ethical behaviors and then **autonomously behave ethically***

- Comparative Moral Turing Test
- Ethical Turing Test
- Evaluate the morality of the choices of automated systems
- **Branch quite unexplored:** difficult connection between philosophy, ethic and technical problems
- AGI related

How humans should design AI systems to minimize harms

*Designing for **minimizing harms** derived from **poor design, bad applications and misuse** of the systems*

- **Algorithmic Fairness**
- Privacy Preserving Data Mining – Federated Learning
- Explainable AI [2] & Interpretable AI
- Adversarial Learning
- Many more examples due to many different ML methods and problems addressed

HCML Perspective: building responsible AI including human relevant requirements, but also considering broad societal issues [1]

- Safety, **Fairness**, privacy, accountability & interpretability - Ethics and legislation



Human centric ML approaches

AI systems learning moral notions

*AI-based systems can **learn moral notions** or ethical behaviors and then **autonomously behave ethically***

- Comparative Moral Turing Test
- Ethical Turing Test
- Evaluate the morality of the choices of automated systems
- **Branch quite unexplored:** difficult connection between philosophy, ethic and technical problems

How humans should design AI systems to minimize harms

*Designing for **minimizing harms** derived from **poor design, bad applications and misuse** of the systems*

- **Algorithmic Fairness**
- Privacy Preserving Data Mining – Federated Learning
- Explainable AI [2] & Interpretable AI
- Adversarial Learning
- Many more examples due to many different ML methods and problems addressed

HCML Perspective: building responsible AI including human relevant requirements, but also considering broad societal issues [1]

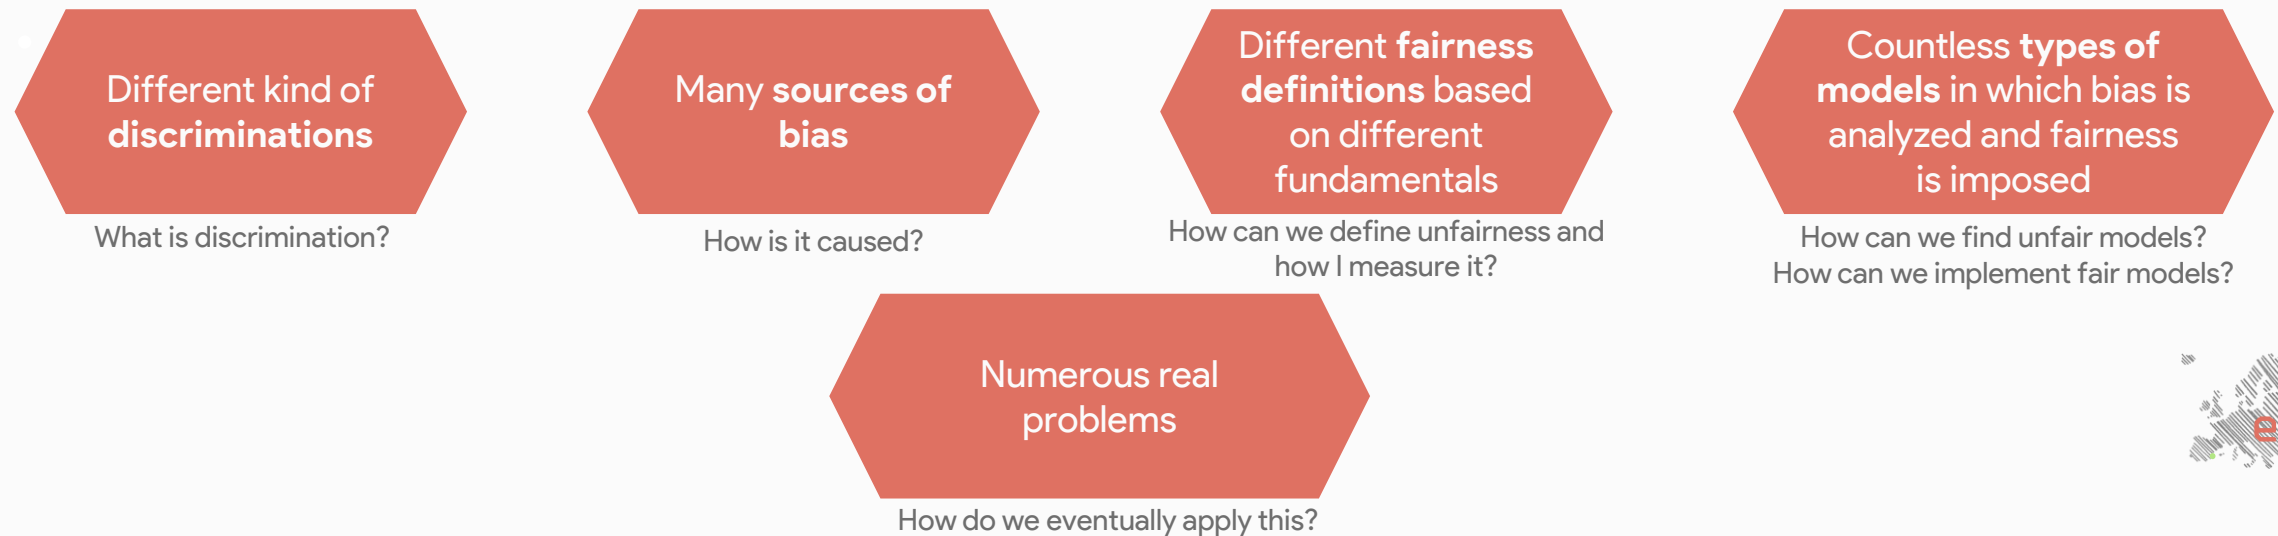
- Safety, Fairness, privacy, accountability & interpretability - Ethics and legislation



What should we consider to formally defining fairness?

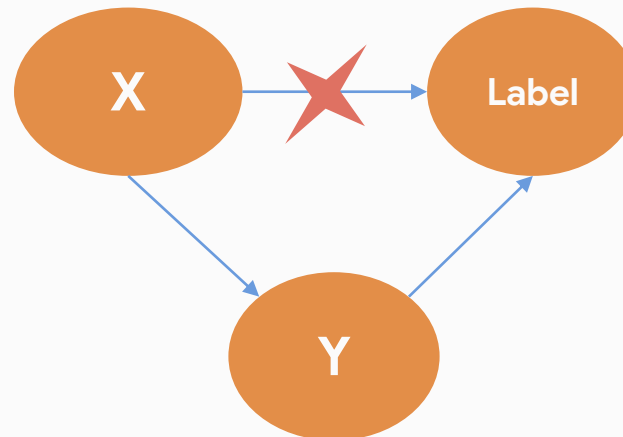
- How we define different discriminations?
- What are the main sources of bias?
- How we define fairness and measure it?
- How do we find bias in our models?
- How we mitigate bias / impose fairness in our models?
 - What kind of different approaches are there?
- What are some examples of real applications?

Hints on the complexity of formally defining fairness



Algorithmic Fairness

- Algorithmic Fairness deals with the problem of developing AI-based systems able to treat:
 - Subgroups in the population equally → **Group fairness**
 - Similar individuals in a similar way → **Individual Fairness**
- Subgroups → determined by means of sensitive attributes, considered for decisions
 - Gender, incomes, ethnicity, and sexual or political orientation and so on*




How do we define equally?
How do we define similar?



Algorithmic Fairness

- How to enhance ML models with fairness requirements, not unethically biasing decisions



- Ensure that the outputs of a model DO NOT depend on sensitive attributes
 - *In a way that is considered unfair - differences due to such traits cannot be reasonably justified*

$$F(X) = R, A \in X \rightarrow R \perp A$$


- Many approaches: properties of the model outputs with respect to the sensitive attributes
- **Relationships among all relevant variables in the data → unfairness underlying**
 - If not → COMPAS: biased recidivism application even not using sensitive data




Two Petty Theft Arrests

<p>VERNON PRATER</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>BRISHA BORDEN</p> <p>Prior Offenses 4 juvenile misdemeanors</p> <p>Subsequent Offenses None</p> <p>HIGH RISK 8</p>
---	--

Black Defendants' Risk Scores



White Defendants' Risk Scores



Two Drug Possession Arrests

<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>
--	---

Prediction Fails Differently for Black Defendants

	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%

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

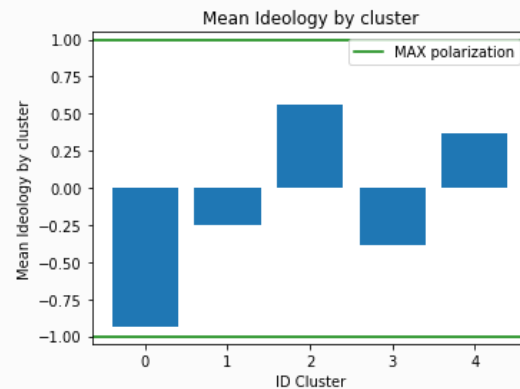

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

May 23, 2016

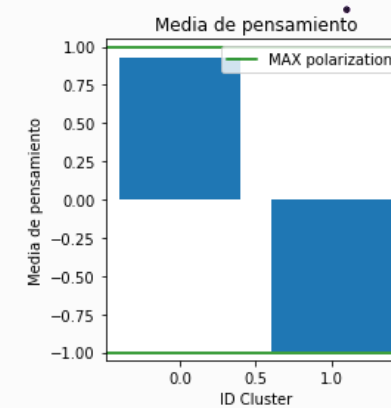

Correctional Offender Management Profiling for Alternative Sanctions - COMPAS

Not only fair decisions: echo chambers

- US House of Representatives 1973 VS 2016
- Two politicians are linked if they have supported 3+ initiatives together



1973



2016



Before kicking off: spoiler

- There are quite a lot different approaches to mitigating unfairness.
- No single approach is universally best → No free lunch 😞
- Choosing the most appropriate one will require:

Expert judgement

Knowledge of relevant legal
and compliance requirements

Context in which we are
working

Takeaway: Choosing Fairness metric and method highly depends on the context

No universal fairness definition or bias mitigation / imposing fairness approach





Bias

Different types


Bias & Sources

1. How law define bias?

- Disparate treatment
- Disparate impact


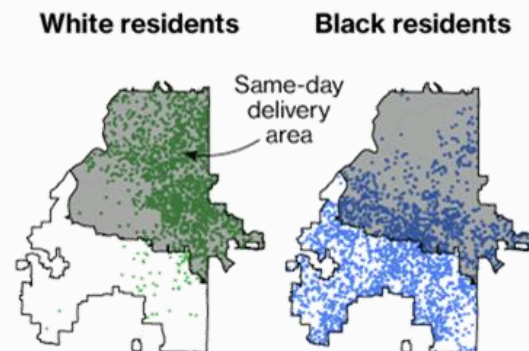
2. Bias in in ML

- By source
- By interaction



Disparate Treatment and Impact

- Anti-discrimination laws in various countries prohibit unfair treatment of individuals
- Legal or ethical support and formalize it quantitatively
 - **Disparate treatment:**
 - Decisions are (partly) based on the subject's sensitive attribute
 - Explicit or intentional
 - **Disparate impact:**
 - Outcomes or implemented policy disproportionately hurt people with certain sensitive attribute
 - Implicit or unintentional



Sources of Bias – Data

Bias in historical data

- Skewed towards groups or imbalanced limited information
 - Amazon, COMPAS or 2018-CEO-image-search
- Easy to ignore biases and surrogate variables for protected attributes
- Label imperfectly observed: Label bias
- Record of crimes comes from crimes observed by police

Bias in data collection mechanisms

- Inherent biases in the data collection mechanisms
- Lack of representativeness
- Crowdsourcing from a technology that only uses a type of people → Autonomous car related with wealthier

Bias in alternate sources of data

- “New” sources of data: worldwide web, social media, blogs
- Digital footprint variables: computer brand or type of device
 - Proxies of protected attributes
 - Socio-economic variables → surrogates for protected groups

Selective labels - Unobservable Outcomes

- Observed outcomes are consequence of the existing choices of the human decision-makers
 - → Label distribution based on previous policy
- Was former policy accurate or biased?
- Would they have defaulted if had they been approved for a mortgage? → Counterfactual
- Tainted samples → Decision-maker bias
- *We observe loan defaults only for those who received a mortgage → we do not have any information for those who were denied*
- *We observe whether a defendant fails to return for their court appearance only if the human judge decides to release the defendant on bail*

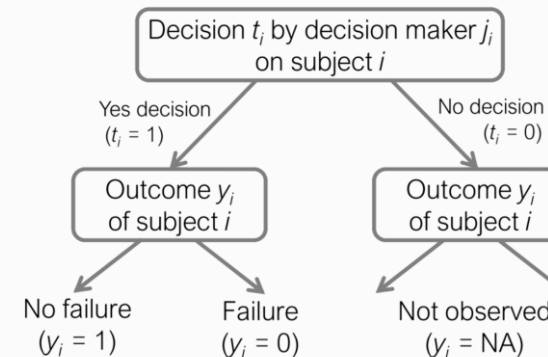
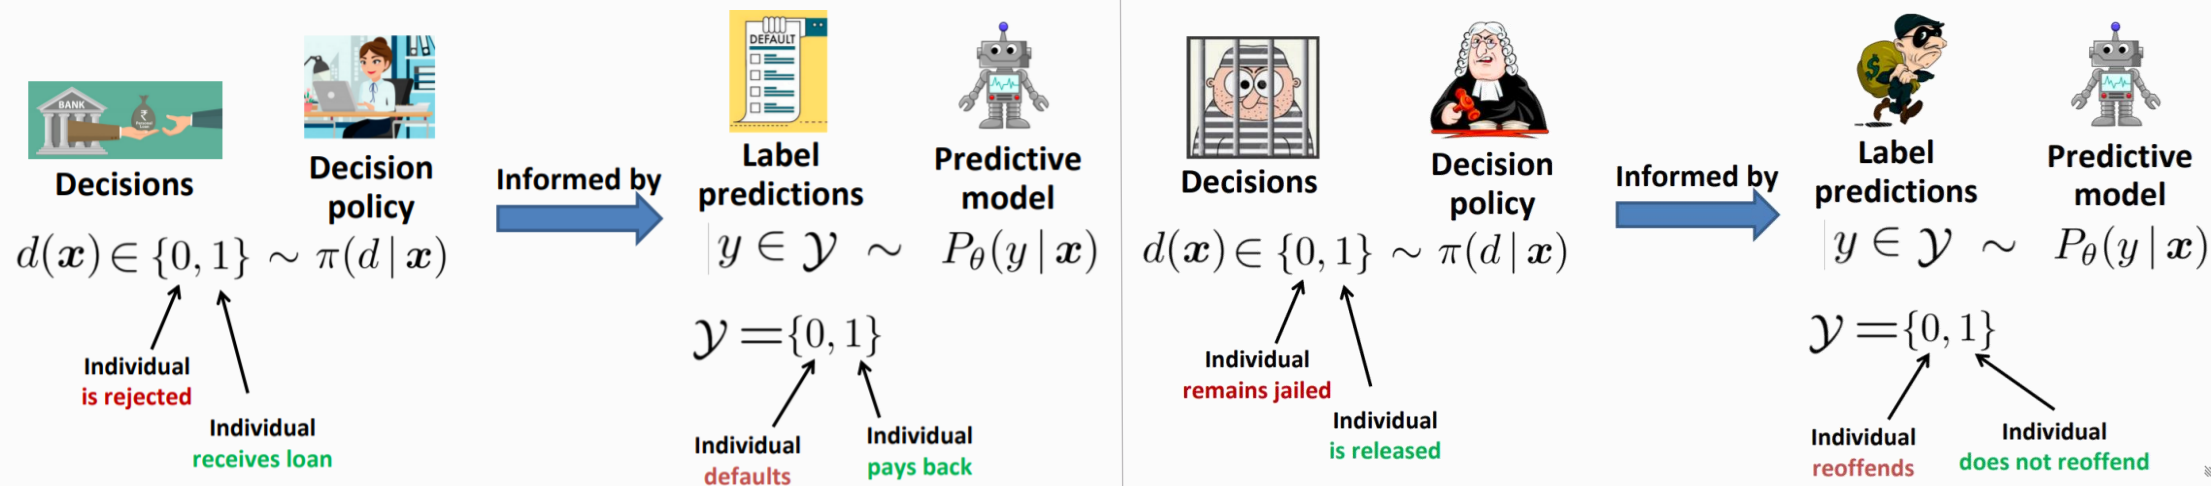
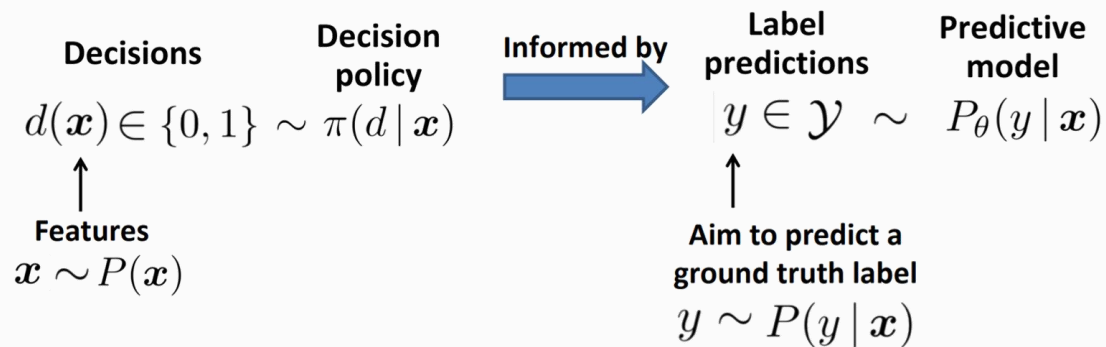


Figure 1: Selective labels problem.

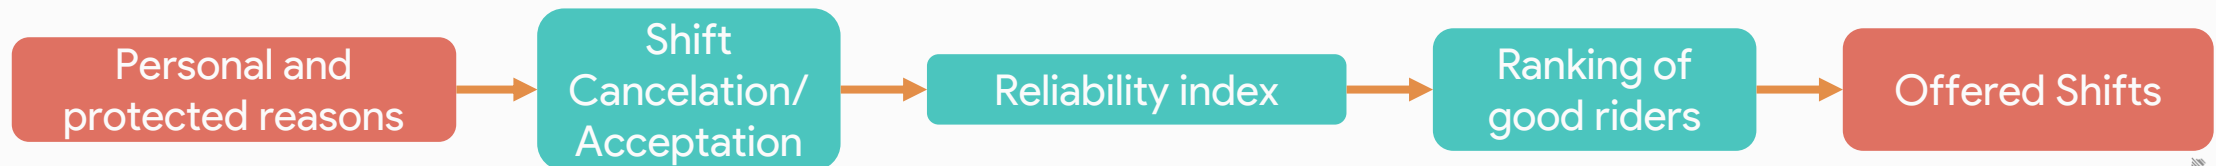


Examples of selective label



Sources of Bias – Algorithm

- The **automated** nature of modern ML
 - Millions of automated data-transformations to get a tiny improvement in predictive performance
 - Don't carefully review the selected variables → surrogate variables and proxy discrimination
- **Overfitting** and **hyperparameter tuning** can amplify biases
- **Opacity** and **lack of interpretability** of complex ML algorithms
 - If one can identify the input-output relationships → easier to isolate potential algorithmic bias
- **Inherent biases in programmers** conveyed to the algorithm
- **Unexpected decisions** in traditional programming
 - Deliveroo riders affected by the ranking algorithm → Reliability index



Sources of Bias – By interaction

- **Data to Algorithm**

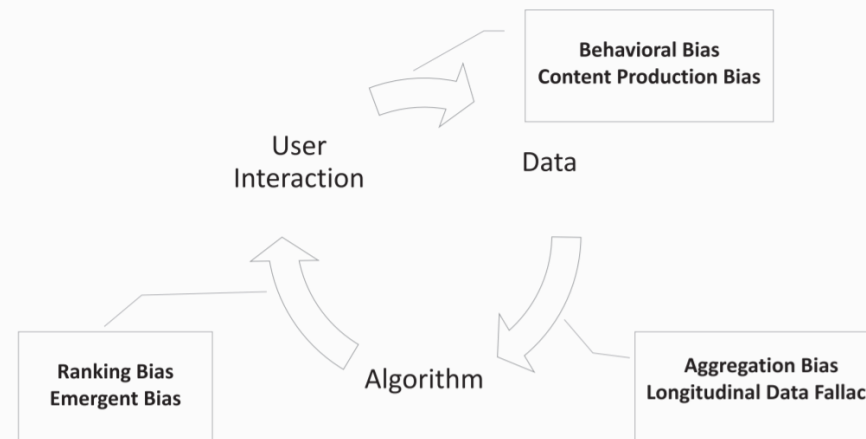
- Measurement Bias
- Omitted Variable Bias
- Representation Bias
- Aggregation Bias
 - E.g., Simpson paradox
- Sampling Bias
- Longitudinal Data Fallacy
- Linking Bias
- Proxie

- **Algorithm to User**

- Algorithmic Bias
- User Interaction Bias - Ranking
- Popularity Bias
- Emergent Bias
- Evaluation Bias

- **User to Data**

- Historical Bias
- Population Bias
- Self-selection Bias
- Social Bias
- Behavioral Bias
- Survivorship bias
- Temporal Bias
- Content production bias





$$P(S=s | A=a) = P(S=s | A=b)$$


Fairness definitions and metrics

Several notions of fairness
already exist in the literature

Recap: Algorithmic Fairness

- Algorithmic Fairness deals with the problem of developing AI-based systems able to treat:

- Subgroups in the population equally → **Group fairness**
- Similar individuals in a similar way → **Individual Fairness**
- Other newer approaches



- Subgroups → determined by means of sensitive attributes, considered for decisions
 - *Gender, incomes, ethnicity, and sexual or political orientation and so on*
- Ensure that the outputs of a model DO NOT depend on sensitive attributes
 - *In a way that is considered unfair - differences due to such traits cannot be reasonably justified*

$$F(X) = R, A \in X \rightarrow R \perp A$$

- Many approaches: properties of the model outputs with respect to the sensitive attributes

How do we define equally?

How we define similar?



Decision Rules: Classification

- Each individual has a set of features:
 - $x_i \in \mathbb{R}^p$
- x can be partitioned into protected and unprotected features:
 - $x = (x_u, x_p)$
 - Set of protected features: $A \in X \rightarrow$ different A values leads to different protected groups
- Target of prediction
 - $y \in \{0, 1\}$
- Training samples
 - $D = \{(x_i, y_i)\}_i^N$
- Random variables X and Y that take on values $X = x$ and $Y = y$ for an individual drawn randomly from the population of interest
- Binary classification
 - $f: \mathbb{R}^p \rightarrow \{0, 1\}$, where $\hat{y} = f(x)$ or, in population level $\hat{Y} = f(X)$
- Risk score
 - True risk score: $r(x) = \Pr(Y = 1|X = x)$
 - Model approximation of risk score $s(x)$ instead of binary decision and $d(x) = 1$ iff $s(x) > t$
 - $R = E[Y|X]$

In binary classification \rightarrow probability of decision S



Confusion matrix

Event	Condition	Notion $P(event condition)$
$\hat{Y} = 0$	$Y = 0$	True Negative rate
$\hat{Y} = 1$	$Y = 0$	False Positive rate
$\hat{Y} = 0$	$Y = 1$	False Negative rate
$\hat{Y} = 1$	$Y = 1$	True Positive rate

Classical clf criteria

Event	Condition	Notion $P(event condition)$
$Y = 0$	$\hat{Y} = 0$	Positive predicted value
$Y = 1$	$\hat{Y} = 1$	Negative predicted value

Additional clf criteria

		Predicted Label		
		$\hat{y} = 1$	$\hat{y} = -1$	
True Label	$y = 1$	True positive	False negative	$P(\hat{y} \neq y y = 1)$ False Negative Rate
	$y = -1$	False positive	True negative	$P(\hat{y} \neq y y = -1)$ False Positive Rate
		$P(\hat{y} \neq y \hat{y} = 1)$ False Discovery Rate	$P(\hat{y} \neq y \hat{y} = -1)$ False Omission Rate	$P(\hat{y} \neq y)$ Overall Misclass. Rate

Confusion matrix allow us to go further accuracy in error explanations related with joint distributions of (X, \hat{Y}, Y)

		Predicted Label	
		Positive	Negative
True Label	Positive	True Positives $PPV = \frac{TP}{TP + FP}$ $TPR = \frac{TP}{TP + FN}$	False Negative $FOR = \frac{FN}{FN + TN}$ $FNR = \frac{FN}{FN + TP}$
	Negative	False Positive $FDR = \frac{FP}{FP + TP}$ $FPR = \frac{FP}{FP + TN}$	True Negatives $NPV = \frac{TN}{TN + FN}$ $TNR = \frac{TN}{TN + FP}$



More confusion matrix measures

$$Pr(\hat{Y} = y | Y = y)$$

$$Pr(Y = y | \hat{Y} = y)$$

		Predicted condition			
		Positive (PP)	Negative (PN)		
Total population = P + N				Informedness, bookmaker informedness (BM) = TPR + TNR - 1	Prevalence threshold (PT) = $\frac{\sqrt{TPR \times FPR} - FPR}{TPR - FPR}$
Actual condition	Positive (P)	True positive (TP), hit	False negative (FN), type II error, miss, underestimation	True positive rate (TPR), recall, sensitivity (SEN), probability of detection, hit rate, power = $\frac{TP}{P} = 1 - FNR$	False negative rate (FNR), miss rate = $\frac{FN}{P} = 1 - TPR$
	Negative (N)	False positive (FP), type I error, false alarm, overestimation	True negative (TN), correct rejection	False positive rate (FPR), probability of false alarm, fall-out = $\frac{FP}{N} = 1 - TNR$	True negative rate (TNR), specificity (SPC), selectivity = $\frac{TN}{N} = 1 - FPR$
Prevalence = $\frac{P}{P + N}$		Positive predictive value (PPV), precision = $\frac{TP}{PP} = 1 - FDR$	False omission rate (FOR) = $\frac{FN}{PN} = 1 - NPV$	Positive likelihood ratio (LR+) = $\frac{TPR}{FPR}$	Negative likelihood ratio (LR-) = $\frac{FNR}{TNR}$
Accuracy (ACC) = $\frac{TP + TN}{P + N}$		False discovery rate (FDR) = $\frac{FP}{PP} = 1 - PPV$	Negative predictive value (NPV) = $\frac{TN}{PN} = 1 - FOR$	Markedness (MK), deltaP (Δp) = PPV + NPV - 1	Diagnostic odds ratio (DOR) = $\frac{LR+}{LR-}$
Balanced accuracy (BA) = $\frac{TPR + TNR}{2}$		F₁ score = $\frac{2PPV \times TPR}{PPV + TPR} = \frac{2TP}{2TP + FP + FN}$	Fowlkes–Mallows index (FM) = $\sqrt{PPV \times TPR}$	Matthews correlation coefficient (MCC) = $\frac{\sqrt{TPR \times TNR \times PPV \times NPV} - \sqrt{FNR \times FPR \times FOR \times FDR}}$	Threat score (TS), critical success index (CSI), Jaccard index = $\frac{TP}{TP + FN + FP}$

- Confusion matrix allow us to go further accuracy in error explanations related with joint distributions of (X, \hat{Y}, Y)
- However, it may seem quite unmanageable to try all possible combinations
- **How do we leverage all this measures for fairness?** → Add sensitive attribute to conditional probabilities



Considering a binary protected attribute $A \in \{a, b\}$ and $\hat{Y} = d$

Group fairness: main definitions

Predicted Outcome (\hat{Y}) $\rightarrow A \perp S$

- **Demographic parity** [1] $\rightarrow A \perp S$ (independence)

$$P(d=1|A=a) = P(d=1|A=b)$$

Predicted (\hat{Y}) and Actual Outcomes (d)

- **Predictive parity** [2] - Same PPV $\rightarrow A \perp Y | S$ (sufficiency)

$$P(Y=1 | d=1, A=a) = P(Y=1 | d=1, A=b)$$

- Predictive equality - Same FPR [TNR]

$$P(d=1 | Y=0, A=a) = P(d=1 | Y=0, A=b)$$

- **Equal opportunity** - Same FNR [TPR]

$$P(d=0 | Y=1, A=a) = P(d=0 | Y=1, A=b)$$

- **Equalized odds** [3] - same TPR and FPR $\rightarrow A \perp S | Y$ (separation)

$$P(d=1 | Y=i, A=a) = P(d=1 | Y=i, A=b), \forall i \in \{0, 1\}$$

- Conditional use accuracy equality - same accuracy for G

$$P(Y=1 | d=1, A=a) = P(Y=1 | d=1, A=b) \wedge$$

$$P(Y=0 | d=0, A=a) = P(Y=0 | d=0, A=b)$$

- Overall accuracy equality - general accuracy

$$P(d=Y, A=a) = P(d=Y, A=b).$$

- Treatment equality - same ratio of errors.

$$(FN/FP)_f = (FN/FP)_m.$$

Predicted Probabilities (S) and Actual Outcome (d) $\rightarrow A \perp Y | S$

- **Calibration** - predictive parity but with probabilities $\rightarrow A \perp Y | S$

$$P(Y=1 | S=s, A=a) = P(Y=1 | S=s, A=b), \forall s \in [0, 1]$$

- Well calibration

$$P(Y=1 | S=s, A=a) = P(Y=1 | S=s, A=b) = s, \forall s \in [0, 1]$$

- Balance for positive class - equal average predicted S for actual positives

$$E(S | Y=1, A=a) = E(S | Y=1, A=b)$$

- Balance for negative class - same average predicted S for actual negatives

$$E(S | Y=0, A=a) = E(S | Y=0, A=b)$$

ML model should behave equally, or at least similarly, no matter whether it is applied to one subgroup in the population or to another one

Example of incompatibility
If different base rate $P(Y=1|A=a) \neq P(Y=1|A=b)$
and satisfies predictive parity
 \rightarrow Cannot satisfy Equalized odds

Barocas, S., Hardt, M., & Narayanan, A. (2017). **Fairness in machine learning**. Nips tutorial, 1, 2017

Verma, S., & Rubin, J. (2018). **Fairness definitions explained**. In 2018 IEEE/ACM Fairware. IEEE.

[1] Cynthia Dwork, et al. 2012. **Fairness Through Awareness**. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference

[2] Alexandra Chouldechova. 2016. **Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments**. Big Data.

[3] Moritz Hardt, Eric Price, and Nati Srebro. 2016. **Equality of Opportunity in Supervised Learning**. In Advances in Neural Information Processing Systems



Definition clarification: Formal criteria

$$P(d=[0,1] | Y=[0,1]) \text{ AND } P(Y=[0,1] | d=[0,1])$$

$$P(D=d | Y=y, A=a) = P(D=d | Y=y, A=b)$$

D \ Y	0	1
0	Predictive equality	Equal opportunity
1	Predictive equality Equal odds	Equal opportunity Equal odds

} Group fairness and conditional statistical parity

$$P(Y=y | D=d, A=a) = P(Y=y | D=d, A=b)$$

Y \ D	0	1
0	Conditional use acc	Predictive parity
1		Predictive parity conditional use acc

} Overall accuracy



Definition clarification: Formal criteria

“Many fairness criteria have been proposed over the years, each aiming to formalize different desiderata. We’ll start by jumping directly into the formal definitions of three representative fairness criteria that relate to many of the proposals that have been made.” (Hardt et al., Fairness in Machine Learning book, 2019)

$P(S A)$	$P(S Y, A)$	$P(Y S, A)$
<i>Independence</i>	<i>Separation</i>	<i>Sufficiency</i>
$S \perp A$	$S \perp A Y$	$A \perp Y S$

Demographic parity

$$P(d=1|A=a) = P(d=1|A=b)$$

Positive Predicted Ratio
Equal acceptance rate




Equalized odds

$$P(d=1 | Y=i, A=a) = P(d=1 | Y=i, A=b), i \in 0, 1$$

Equal opportunity

$$P(d=0 | Y=1, A=a) = P(d=0 | Y=1, A=b)$$

TPR - FPR
Equal error rates



Predictive Parity

$$P(Y=1 | d=1, A=a) = P(Y=1 | d=1, A=b)$$

Calibration

$$P(Y=1 | S=s>t, A=a) = P(Y=1 | S=s>t, A=b) \forall t$$

PPV - NPV
Calibration by group



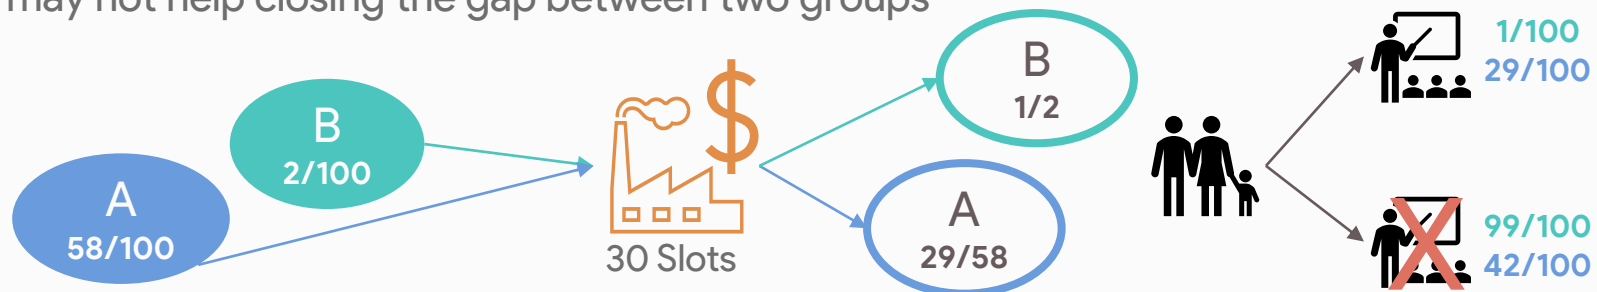
Definition clarification: Formal criteria

List of demographic fairness criteria			
Name	Closest relative	Note	Reference
Statistical parity	Independence	Equivalent	Dwork et al. (2011)
Group fairness	Independence	Equivalent	
<u>Demographic parity</u>	Independence	Equivalent	
Conditional statistical parity	Independence	Relaxation	Corbett-Davies et al. (2017)
Darlington criterion (4)	Independence	Equivalent	Darlington (1971)
Equal opportunity	Separation	Relaxation	Hardt, Price, Srebro (2016)
<u>Equalized odds</u>	Separation	Equivalent	Hardt, Price, Srebro (2016)
Conditional procedure accuracy	Separation	Equivalent	Berk et al. (2017)
Avoiding disparate mistreatment	Separation	Equivalent	Zafar et al. (2017)
Balance for the negative class	Separation	Relaxation	Kleinberg, Mullainathan, Raghavan (2016)
Balance for the positive class	Separation	Relaxation	Kleinberg, Mullainathan, Raghavan (2016)
Predictive equality	Separation	Relaxation	Chouldechova (2016)
Equalized correlations	Separation	Relaxation	Woodworth (2017)
Darlington criterion (3)	Separation	Relaxation	Darlington (1971)
Cleary model	Sufficiency	Equivalent	Cleary (1966)
Conditional use accuracy	Sufficiency	Equivalent	Berk et al. (2017)
<u>Predictive parity</u>	Sufficiency	Relaxation	Chouldechova (2016)
Calibration within groups	Sufficiency	Equivalent	Chouldechova (2016)
Darlington criterion (1), (2)	Sufficiency	Relaxation	Darlington (1971)



Group fairness gaps

- Proved that statistical definitions are insufficient [1, 2, 3, 4]
- Moreover, most valuable statistical metrics assume availability of actual, verified outcomes.
 - Problems with Selective label bias
- Similar individuals may not be treated equally for achieving measures of group fairness
- Demographic Parity [*Independence*]
 - Ignores any possible correlation between Y and A
 - E.g., Perfect predictor ($S=Y$) is not considered fair when base rates differ (i.e., $P[Y=1 | A=a] \neq P[Y=1 | A=b]$)
 - laziness: if we hire the qualified from one group and random people from the other group, we can still achieve demographic parity.
- Equalized Odds – Predictive Parity [*separation and sufficiency*]
 - It may not help closing the gap between two groups



[1] Richard Berka, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2017. *Fairness in Criminal Justice Risk Assessments: The State of the Art*.
[2] Alexandra Chouldechova. 2016. *Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments*. Big Data (2016)
[3] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. *Fairness Through Awareness*. 3rd Innovations in Theoretical CS Conference.
[4] Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2017. *Inherent Trade-Offs in the Fair Determination of Risk Scores*. In ITCS



Individual Fairness

- Group Fairness → *Similar individuals could not be treated equally due to calibrations across groups to achieve group fairness measures*
- Individual Fairness → **treating similar individuals similarly**
 - Difference between individuals similar to difference in predictions
 - More fine-grained than any group-notion fairness: it imposes restriction on for each pair of i .


Our Dataset: $D = \{(x_i, y_i)\}_i^N$

Distance between x_i pairs: $k: V \times V \rightarrow R$.

Mapping from x_i to probability distribution over outcomes $M: V \rightarrow \alpha A$

Distance between distributions of outputs D

Individual fairness $D(M(x), M(y)) \leq k(x, y)$



- ? How to define appropriate distance metrics for the specific problem and application?

Metric Learning

Graph Theory



Representation Learning



Individual Fairness flaws

Graph Theory
Representation Learning
Semi/Self-Supervised Learning

- Big expertise to establish a distance metric between individuals.
 - Metrics can still be implicit biased ☹️
- Testing definitions relies on availability of “similar” individuals
 - Search space very large → e.g., the global population.
 - More work to **narrow down the search space without impeding the accuracy**
- **Distance between data does not only depends on pairwise distances**
 - Relationships among every all the data and topology (*cliques or communities on graphs*)
- Very difficult to find the proper metric (both d and M)
 - Specifically, M → unseen labels → Selective Labels / unobserved variables / substitutes labels





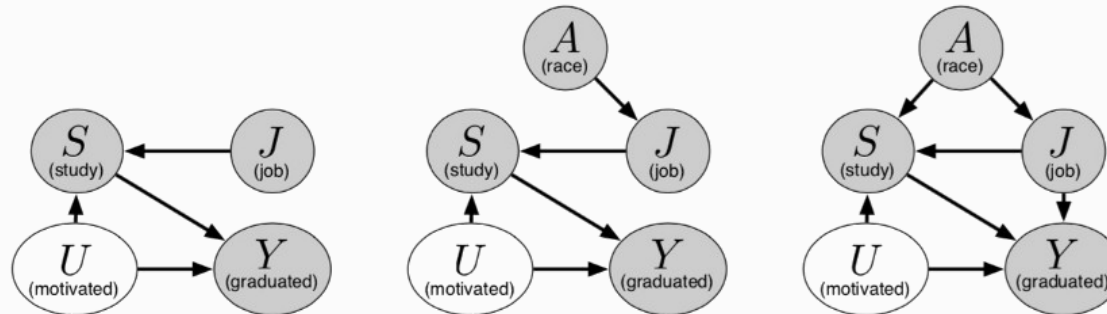
Is individual A closer to B than C? How much?
→ very metric dependent d

Is A closer to B than C regarding their predicted performance?
→ We don't have real ground truth → Selective labels
→ Very metric dependent M



Counterfactual fairness

- *Group*
 - Observational fairness criteria 
 - Cannot find the cause of the unfairness 
- *Individual*
 - Limitation of finding the proper metric.
- **Causality** → Explaining the impact of bias via a causal graph
 - Replacing A, other correlated features with it will also be influenced




Causal graphs: Acyclic graphs
- nodes representing attributes
- edges representing relationships

- Ideal idea? hard to reach a consensus in terms of
 - what the **causal graph** should look like?
 - which **features to use** even if we have such a graph?



Counterfactual fairness

- **Counterfactual** → “Would I have been hired if I were non-black?” “Would I have avoided the traffic jam had I taken a different route this morning?”
 - Decision does not depend on protected attribute
- The counterfactual $Y_{\{X:=1, Z:=Z_{X:=0}\}}$ is the value that Y would obtain had X been set to 1 and had Z been set to the value Z would’ve assumed had X been set to 0
- Fair Causal graph → if Y don’t depend on A, i.e., no A-Y way
 - Make decision only using non-descendants of A in the causal graph
- Difficult task of agreeing on which graph to build and validating it
- Impossible to test an existing classifier against **strict** causal definitions of fairness
- What should we do when not we are not able to built neither validate a causal graph?
 - Counterfactual discrimination criteria → normative fairness criteria



Counterfactual fairness

- Notation of $d(w)$, $d(m)$ be the decision if the individual had been woman or men

- **Individual Counterfactual Fairness**

$d_i(w) = d_i(m)$ for individual i and every other attribute remaining the same, i.e.,

$$P(\hat{Y}_{\{A \leftarrow a\}}(U) = y | X = x, A = a) = P(\hat{Y}_{A \leftarrow b}(U) = y | X = x, A = a)$$

- negative answer to “would the decision have been different if I were not black?”

- **Counterfactual Demographic Parity**

$$E[d(w)] = E[d(m)] \text{ i.e.,}$$

$$E[\hat{Y} | X = x, A = a] = E[\hat{Y} | X = x, A = b] \forall X \text{ and } \forall (a, b)$$

- negative answer to “would the rates of hiring be different if everyone were black?”

Related with *Conditional Demographic Parity*
 $P(d = 1 | L = l, A = a) = P(d = 1 | L = l, A = b)$
which means $\hat{Y} \perp A | X$

- **Conditional Counterfactual Parity**

$$E[d(w) | X] = E[d(m) | X]$$

- “would the rates of hiring be different if everyone were black?” BUT stratified by some factors

- The easiest way to satisfy counterfactual demographic parity is :
prediction only use non-descendants of A in the causal graph



Counterfactual in real world

*“Race plays a significant role in admissions decisions. Consider the example of an **Asian-American applicant who is male, is not disadvantaged**, and has other characteristics that result in a **25% chance of admission**. Simply **changing the race of the applicant to white**— and leaving all his other characteristics the same—would increase his chance of admission to **36%**. Changing his race to **Hispanic** would increase his chance of admission to **77%**. Changing his race to **African-American** would increase his chance of admission to **95%**”.*

(150 Plaintiff’s expert report of Peter S. Arcidiacono, Professor of Economics at Duke University)

- Logistic regression model against Harvard’s past admissions decisions
- Conditional statistical parity is not satisfied

$$P(d=1 | L=l, A=a) = P(d=1 | L=l, A=a)$$



Fairness measurement in benchmarking dataset

- So, is the classifier fair? → Logistic regression on German Credit Dataset

	Definition	Paper	Citation #	Result
3.1.1	Group fairness or statistical parity	[12]	208	×
3.1.2	Conditional statistical parity	[11]	29	✓
3.2.1	Predictive parity	[10]	57	✓
3.2.2	False positive error rate balance	[10]	57	×
3.2.3	False negative error rate balance	[10]	57	✓
3.2.4	Equalised odds	[14]	106	×
3.2.5	Conditional use accuracy equality	[8]	18	×
3.2.6	Overall accuracy equality	[8]	18	✓
3.2.7	Treatment equality	[8]	18	×
3.3.1	Test-fairness or calibration	[10]	57	✓
3.3.2	Well calibration	[16]	81	✓
3.3.3	Balance for positive class	[16]	81	✓
3.3.4	Balance for negative class	[16]	81	×
4.1	Causal discrimination	[13]	1	×
4.2	Fairness through unawareness	[17]	14	✓
4.3	Fairness through awareness	[12]	208	×
5.1	Counterfactual fairness	[17]	14	-

- Depends on the notion of fairness one wants to adopt.
 - More work is needed to clarify which definitions are appropriate to each particular situation

Context matters



Summary of metrics

- Group Fairness

- Independence, separation, sufficiency
- Confusion matrix-related
- Counterfactual parity

- Individual Fairness

- Metrics
- Individual counterfactual

- Counterfactual

- Conceptually
- Applied

- Many more...




Table 1
A synthetic review of most of the notions of fairness.

Notion	Abbreviation	First Appeared
α -Protection	α -P	[159]
Indirect Discriminatory Measure	ELB	[72]
Decision Policy Discrimination	DPD	[131]
Prediction Dependency	PredD	[23]
Dataset Discrimination	DD	[97]
Discrimination Score	DS	[22]
Calders-Verwer Score	CVS	[22, 105]
Statistical Parity	SP	[50]
Mean Difference	MD	[24]
Area Under ROC Curve	AUC	[24]
Disparate Impact	DI	[56]
ϵ -Fairness	ϵ -F	[56]
η -Neutrality	η -N	[60]
Discrimination Correlation Indicator	DCI	[125]
Demographic Parity	DP	[76]
Equal Opportunity	EOP	[76]
Equal Odds	EOD	[76]
Fair Prediction Rule	FairPR	[124]
Indifference	Indiff	[94]
Total Causal Effect	TCE	[206]
Cross-Pair Group Fairness	CPGF	[14]
Hilbert-Schmidt Empirical Cross-Covariance	HSIC	[160]
Expected Statistical Parity	ESP	[37]
Expected Predictive Equality	EPE	[37]
Calibration	Calib	[37]
Balanced Loss	BL	[51]
False Positive Subgroup Fairness	FPSF	[107]
Proxy Discrimination	ProxD	[108]
Proxy Discrimination in Expectation	PDE	[108]
P%-Rule	P-R	[113]
Normalised Disparate Impact	NDI	[137]
α -Discrimination	α -D	[189]
Value Unfairness	ValU	[194]
Absolute Unfairness	AbsU	[194]
Underestimation Unfairness	UeU	[194]
Overestimation Unfairness	OeU	[194]
Preferred Impact	PrefI	[199]
Preferred Treatment	PrefT	[199]
Disparate Mistreatment	DM	[197]
Absolute Value Difference	AVD	[13]
Squared Difference	SD	[13]
Balance	Bal	[28]
Relaxed Equal Odds with Calibration	REOC	[163]
Path Specific Effect	PSE	[143]
Natural Direct Effect	NDE	[143]
Mean Difference Discrimination Score	MDDS	[168]
k-way Discrimination Score	k-DS	[168]
Maximum Discrimination	MaxD	[168]
Discrimination In Prediction	DiscrP	[208]
Loss-Averse Statistical Parity	L-ASP	[5]
Loss-Averse Equal Opportunity	L-AEOP	[5]
Difference of Equal Opportunity	DEO	[33]
Hirschfeld-Gebelein-Rényi	HGR	[134]
Coefficient of Determination	Cod R ²	[114]
Difference of Equal Opportunity	DEOp	[151]
Difference of Equal Odds	DEOd	[151]
Subgroup Risk	SR	[188]
Strong Demographic Parity	SDP	[91]
Strong Pairwise Demographic Disparity	SPDD	[91]


(Continued)

Table 1
(Continued)

Notion	Abbreviation	First Appeared
Group Fairness in Expectation	GFE	[59]
Prejudice Index	PI	[105]
Fair-Factorization	FF	[106]
Resilience to Random Bias	RRB	[58]
Normalised Discounted Difference	rND	[192]
Normalised Discounted Ratio	rRD	[192]
Normalised Discounted KL-divergence	rKL	[192]
Explanatory Conditional Discrimination	ECD	[210]
Expected Conditional Statistical Parity	ECSF	[37]
Individual Proxy Discrimination	IPD	[108]
Disparate Treatment	(DispT)	[197]
Disparity Amplification	DA	[78]
k-Neighbours Difference	k-ND	[126]
Fairness Lipschitz Property	FLP	[50]
Cross-Pair Individual Fairness	CPIF	[14]
Decision Boundary Covariance	DBC	[198]
Random Bias Individual Fairness	RBI	[57]
Inconsistency Score	IS	[168]
(α, γ) -Approximately Metric-Fair	(α, γ) -AMF	[196]
Constant Relative Risk Aversion	CRRA	[81]
Rawlsian Equal Opportunity	R-EOP	[82]
Egalitarian Equal Opportunity	e-EOP	[82]
Generalised Entropy Index	GEI	[179]
Counterfactual Fairness	CF	[117]
ϵ, δ -Approximate Counterfactual Fairness	ϵ, δ -ACF	[171]
Counterfactual Direct Effect	CF-DE	[204]
Counterfactual Indirect Effect	CF-IE	[204]
Counterfactual Spurious Effect	CF-SE	[204]
Chebyshev Demographic Parity	CDP	[207]
Maximum Mean Discrepancy	MMD	[68]
Fairness Ramp-Constraint	FRC	[65]
δ -fairness	δ -F	[96]
Impartiality Score	IS	[94]
Formal Equality of Opportunity	FE	[94]
Full Substantive Equality of Opportunity	F-SEO	[94]
Log-Linear Interaction	LLI	[190]
Markov Decision Fairness	MDF	[88]
Approximate-Choice Markov Decision Fairness	α -CF	[88]
Approximate-Action Markov Decision Fairness	α -AF	[88]
Indirect Influence	II	[11]
ϵ -Loss Fair	ϵ -LF	[49]
α -MultiCalibration	α -MC	[80]
Covariance Constraint	CC	[149]
Metric MultiFairness	MMC	[109]
ϵ -Loss General Fair	ϵ -LGF	[153]
Mutual Information	MI	[186]
Kullback-Leibler Divergence	KL-D	[186]
Wasserstein Distance	WD	[201]
Path Specific Counterfactual Fairness	PSCF	[26]



Metrics clarification



- **Theory:** Formal criteria aforementioned:

- $A \perp S | X - A \perp S - A \perp S | Y - A \perp Y | S$

- **Applied:** Majumder, S. et al (2021)

- 26 classification metrics → 7 clusters
 - 4 dataset metrics → 3 clusters

RQ1: Do current fairness metrics agree with each other?

No → 51% agreement

RQ2: Can we group (cluster) fairness metrics based on similarity?

Yes → minimizing intra-cluster disagreement

RQ4: Can we achieve fairness based on all the metrics at the same time?

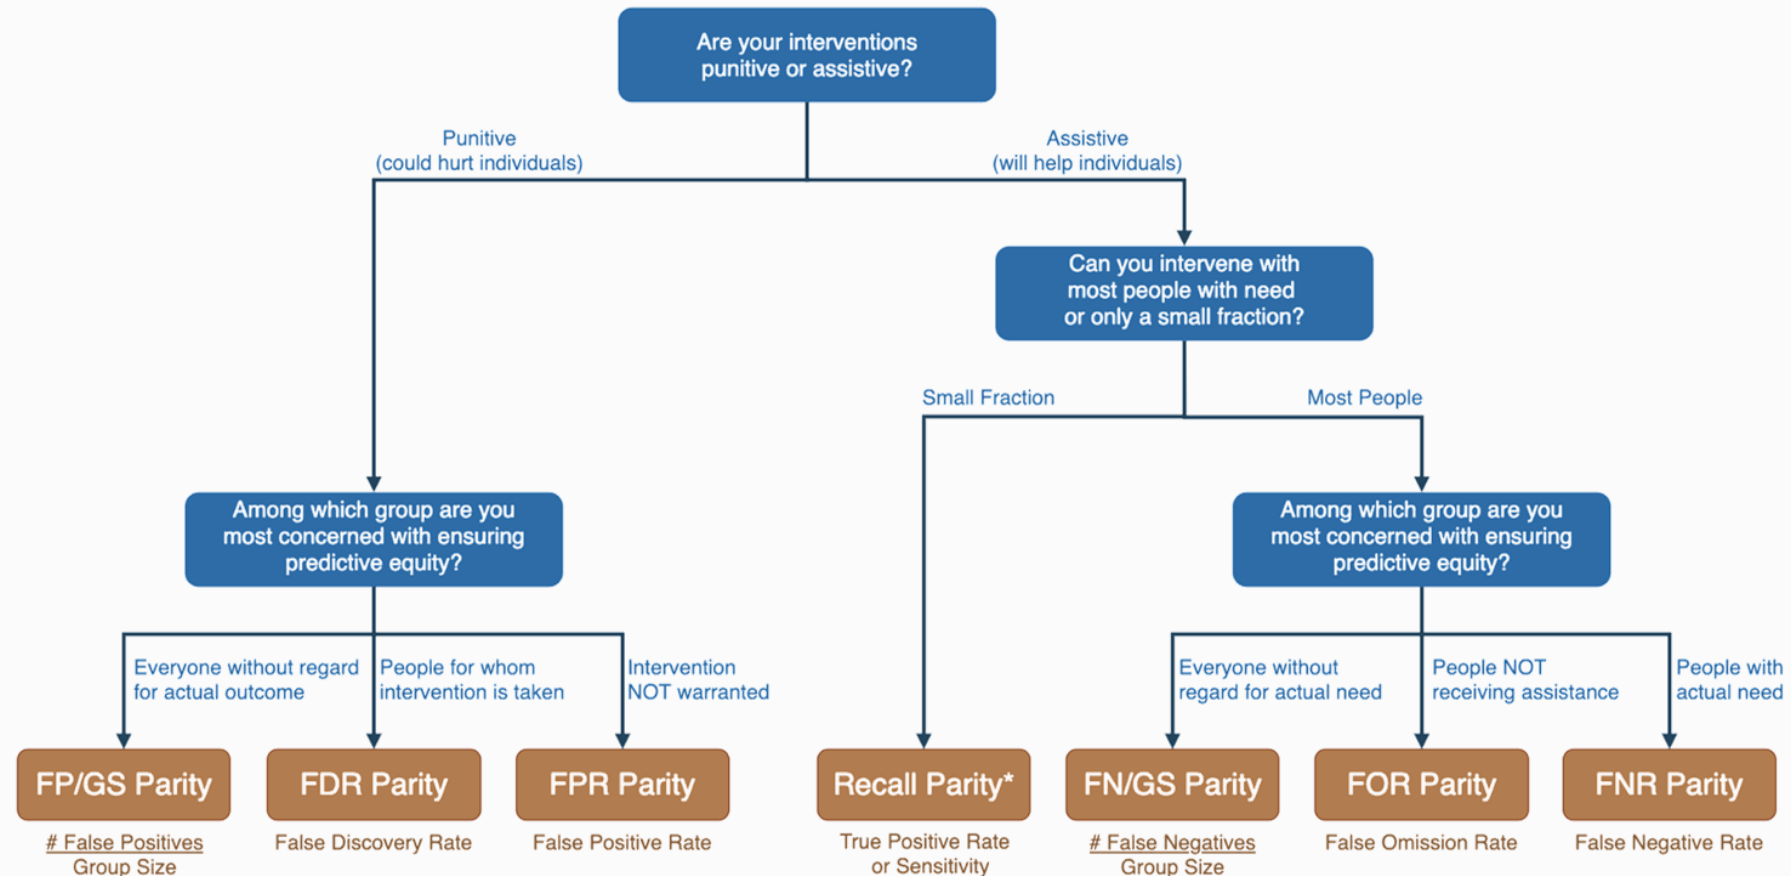
No. Each cluster and metric measure on thing, sometimes opposite

Again, choose depends on the context

Cluster Id	MID	Metrics	Datasets							Metric Type
			Adult	Compas	German	Health	Bank	Student	Titanic	
0	C3	false_omission_rate_difference	Unfair	Fair	Fair	Unfair	Fair	Fair	Unfair	Mis-classification
0	C7	false_omission_rate_ratio	Unfair	Fair	Fair	Unfair	Fair	Unfair	Unfair	
0	C11	error_rate_difference	Unfair	Fair	Fair	Unfair	Fair	Fair	Fair	
0	C12	error_rate_ratio	Unfair	Fair	Fair	Unfair	Fair	Fair	Fair	
		Percentage of agreement	100%	100%	100%	100%	100%	75%	50%	
1	C10	average_abs_odds_difference	Unfair	Unfair	Unfair	Unfair	Unfair	Fair	Unfair	Differential Fairness
1	C25	differential_fairness_bias_amplification	Unfair	Unfair	Unfair	Unfair	Unfair	Fair	Unfair	
		Percentage of agreement	100%	100%	100%	100%	100%	100%	100%	
2	C16	generalized_entropy_index	Fair	Unfair	Fair	Fair	Fair	Fair	Unfair	Individual Fairness
2	C19	theil_index	Unfair	Unfair	Fair	Unfair	Unfair	Fair	Unfair	
2	C20	coefficient_of_variation	Unfair	Unfair	Unfair	Unfair	Unfair	Unfair	Unfair	
		Percentage of agreement	67%	100%	67%	67%	67%	67%	100%	
3	C4	false_discovery_rate_difference	Fair	Fair	Fair	Fair	Fair	Fair	Unfair	Mis-classification
3	C8	false_discovery_rate_ratio	Fair	Fair	Fair	Fair	Fair	Unfair	Unfair	
		Percentage of agreement	100%	100%	100%	65%	100%	50%	100%	
4	C0	true_positive_rate_difference	Unfair	Unfair	Fair	Unfair	Unfair	Fair	Unfair	Confusion Matrix Based Group Fairness
4	C1	false_positive_rate_difference	Fair	Unfair	Unfair	Unfair	Unfair	Fair	Unfair	
4	C2	false_negative_rate_difference	Unfair	Unfair	Unfair	Unfair	Unfair	Fair	Unfair	
4	C5	false_positive_rate_ratio	Fair	Unfair	Unfair	Unfair	Unfair	Fair	Unfair	
4	C6	false_negative_rate_ratio	Unfair	Unfair	Unfair	Unfair	Unfair	Unfair	Unfair	
4	C9	average_odds_difference	Unfair	Unfair	Unfair	Unfair	Unfair	Fair	Unfair	
4	C14	disparate_impact	Unfair	Unfair	Unfair	Unfair	Unfair	Unfair	Unfair	
4	C15	statistical_parity_difference	Unfair	Unfair	Unfair	Unfair	Unfair	Fair	Unfair	
		Percentage of agreement	75%	100%	88%	100%	100%	75%	100%	
5	C17	between_all_groups_generalized_entropy_index	Fair	Fair	Fair	Fair	Fair	Fair	Fair	Between Group Individual Fairness
5	C18	between_group_generalized_entropy_index	Fair	Fair	Fair	Fair	Fair	Fair	Fair	
5	C21	between_group_theil_index	Fair	Fair	Fair	Fair	Fair	Fair	Fair	
5	C22	between_group_coefficient_of_variation	Fair	Fair	Fair	Fair	Fair	Fair	Unfair	
5	C23	between_all_groups_theil_index	Fair	Fair	Fair	Fair	Fair	Fair	Fair	
5	C24	between_all_groups_coefficient_of_variation	Fair	Fair	Fair	Fair	Fair	Fair	Unfair	
		Percentage of agreement	100%	100%	100%	100%	100%	100%	67%	
6	C13	selection_rate	Unfair	Unfair	Unfair	Unfair	Unfair	Unfair	Unfair	Intermediate Metric
		Percentage of agreement	100%	100%	100%	100%	100%	100%	100%	
Percentage of metrics marking dataset as unfair			58%	54%	34%	65%	50%	23%	77%	

Metrics clarification

FAIRNESS TREE (Zoomed in)





Impossibility Theorem

Why different definitions are not compatible?
Inherent Trade-off of fairness

Fairness limitations

- Accuracy VS Fairness
- Group Fairness Impossibility Theorem
- Group VS Individual



Accuracy vs Fairness Tradeoff


Impose constraints on the accuracy with fairness metrics leads to not aligned objectives

Tradeoff depends on how “similar” Y and A are → e.g., if aligned, then linear penalty

The more aligned, the more one will penalize the other

We will have solutions in the pareto front


— Acc=0.87; p%-rule=45%
 - - - Acc=0.82; p%-rule=70%
 Acc=0.74; p%-rule=98%



$\phi = \pi/4$


(a) Maximizing accuracy under fairness constraints

— Acc=0.87; p%-rule=24%
 - - - Acc=0.71; p%-rule=62%
 Acc=0.60; p%-rule=99%



$\phi = \pi/8$


— Acc=0.87; p%-rule=45%
 - - - Acc=0.78; p%-rule=76%
 Acc=0.64; p%-rule=93%



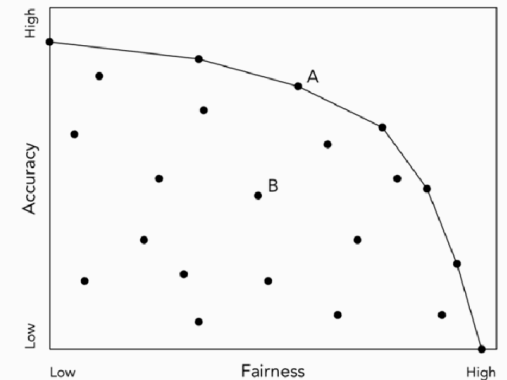
$\phi = \pi/4$

(b) Maximizing fairness under accuracy constraints

— Acc=0.87; p%-rule=24%
 - - - Acc=0.78; p%-rule=54%
 Acc=0.53; p%-rule=99%



$\phi = \pi/8$



$$p\%rule = \min\left(\frac{P\{\hat{Y} = 1 | A = a\}}{P\{\hat{Y} = 1 | A = b\}}, \frac{P\{\hat{Y} = 1 | A = b\}}{P\{\hat{Y} = 1 | A = a\}}\right) \geq \frac{p}{100}$$



Formal criteria's impossibility theorem



- Independence vs sufficiency – DP vs PP
 - If $A \perp\!\!\!\perp Y \rightarrow$ either DP or PP, but NOT BOTH

Independence	Separation	Sufficiency
$A \perp S$	$A \perp S Y$	$A \perp Y S$

$\neg \perp \rightarrow$ dependent || $\perp \rightarrow$ Independent
 Demographic Parity - DP
 Equalized odds - EO
 Predictive Parity - PP

- Independence vs Separation – DP vs EO
 - If $Y \perp\!\!\!\perp A$ & $Y \perp\!\!\!\perp S \rightarrow$ either DP or EO, but NOT BOTH

- Separation vs sufficiency – EO vs PP

- If $P(a, s, y) > 0 \forall AxSxY$ (all events in the joint distribution of have positive probability) AND
- If $A \perp\!\!\!\perp Y \rightarrow$ either EO or PP, but NOT BOTH
- If predictor satisfy EO, PP requires equal PPV, and therefore need equal base rates \rightarrow Not usually happen
- i.e., If different base rates $P(Y=1 | A=a) \neq P(Y=1 | A=b) \rightarrow$ either EO or PP, but NOT BOTH

Group	a	b
Outcome		Unequal base rates
Predictor		

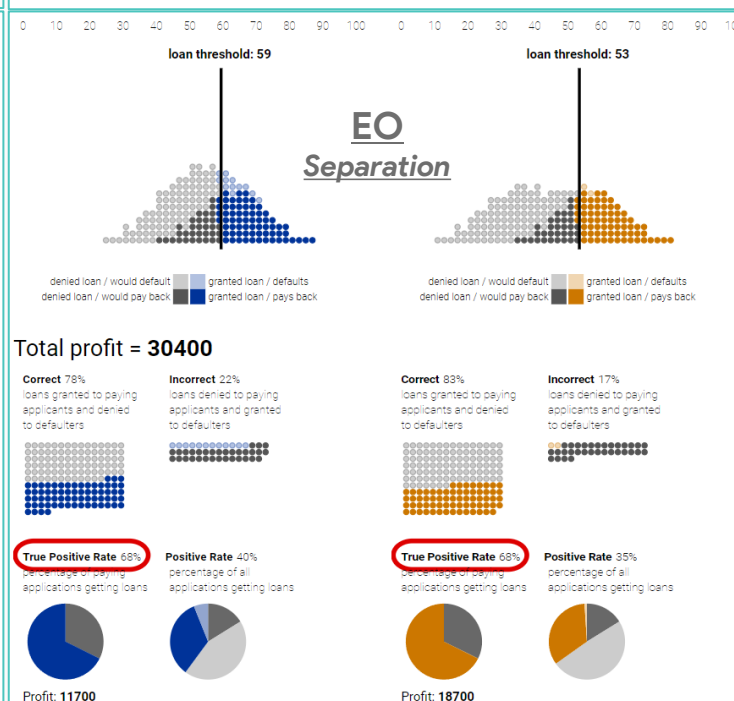
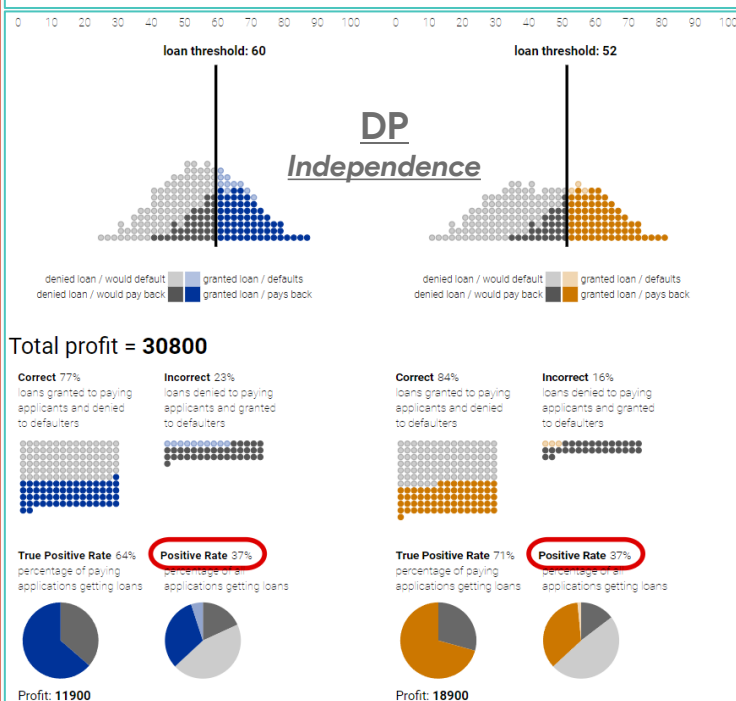
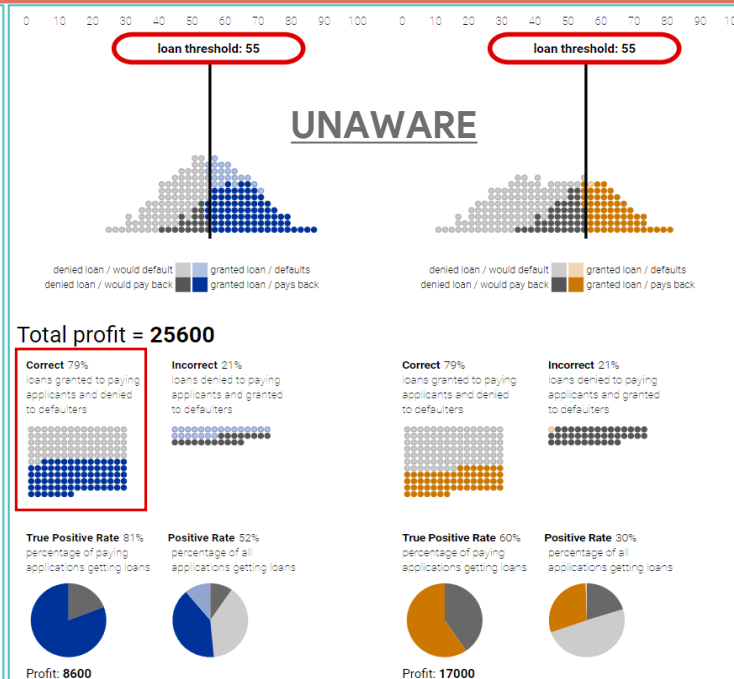
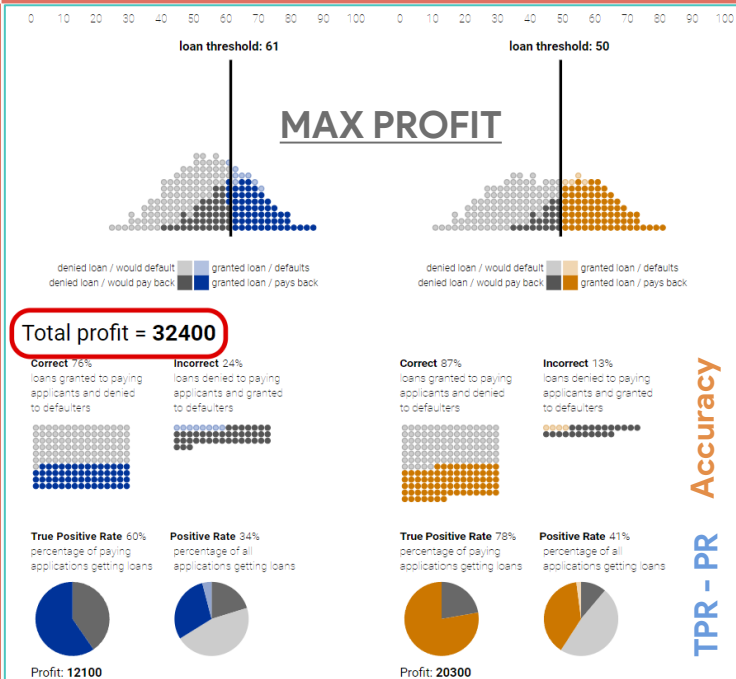
Group	a	b
Outcome		Unequal base rates
Predictor		

Make 2 FP to achieve EO
 Equal TPR and TNR between groups

Group	a	b
Outcome	Unequal base rates	
Predictor		
NPV	2/5	1/3

Negative Predictive Parity violated
 Not possible to preserve NPV without sacrificing EO/PP





Formal criteria's relationship

$$P(y|s, A) \times P(s|A) = P(s|y, A) \times P(y|A)$$

Predictive Parity Demographic Parity Equalized odds Base Rate

Proofs based on Positive Predicted Value, TPR and FPR

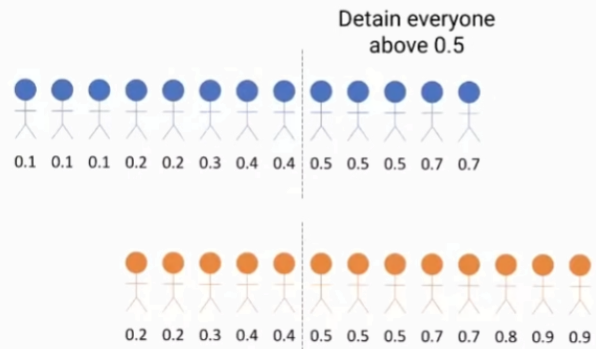
If unequal base rates && not perfect classifier
→ Sufficiency implies that Error parity Fails

Loan granting: 2 groups with different base rates

- Maximize profit → violate TPR and PR
- Unaware → orange gets fewer loans - also violate TPR and PR
- Demographic Parity (PR) → Violates TPR (EO)
- Equalized odds (EO) → Violates PR (DP)



Metrics not sufficient on their own




Detention rate	False pos. rate
38%	25%
61%	42%

— Impedence and error rate parity [EO, FPR] violated



Statistical fairness criteria on their own cannot be a proof of fairness, just a piece of it



Detention rate	False pos. rate
38%	25%
61% 42%	42% 26%

Garg, P., Villasenor, J., & Foggo, V. (2020). *Fairness metrics: A comparative analysis*. In 2020 IEEE Big Data. IEEE.

del Barrio, E., Gordaliza, P., & Loubes, J. M. (2020). Review of mathematical frameworks for fairness in machine learning. arXiv

Castelnovo, A., Crupi, R., Greco, G., & Regoli, D. (2021). The zoo of Fairness metrics in Machine Learning. arXiv preprint arXiv:2106.00467

Chiappa, S., & Isaac, W. S. (2018). A causal bayesian networks viewpoint on fairness. In IFIP International Summer School on Privacy and Identity Management. Springer, Cham.

Oneto, L., & Chiappa, S. (2020). Fairness in Machine Learning. ArXiv, abs/2012.15816.

Martin Wattenberg, Fernanda Viégas, and Moritz Hardt Attacking discrimination with smarter ML. <https://research.google.com/bigpicture/attacking-discrimination-in-ml/>

Moritz Hardt - MLSS 2020, Tübingen. https://youtu.be/lqq_S_7lfOU?t=4056

<http://www-student.cse.buffalo.edu/~atri/algo-and-society/support/notes/fairness/index.html>





Imposing fairness

How to plug chosen fairness definition into
the training on ML algorithms?

How to satisfy Fairness criteria

Pre-processing

- From feature space to a **representation** → **Independence SIA**
- Model learned from this representation will be fair [*Data processing inequality* Information Theory]
- **Model agnostic**
- **Information loss in latent space**

In-processing

- **Fairness constraints** in the optimization process
- **Powerful** → **fairness during the optimization process**
- **Loss of generality** → **each type of model and specific task uses its own regularize**

Post-processing

- Taking a trained classifier → adjust it depending on the sensitive attribute and randomness
- independence is achieved
- **Works for black-box models and no re-training needed**
- **Useful when no access to training data, complex-no access to training pipeline**
- **Not that efficient due to the same reasons**



Lots of them... again

• Method family

- Pre
- In
- Post

• Task

- Binary classification
- Multiclassification
- Regression

• Protected attribute


- Binary
- Categorical
- Numerical

Table 2
A synthetic review of most of the papers available in the literature

Paper	Method Family	Task	Protected Attribute	Notion of Fairness	Theoretical Results	Experimental Results	Comparison Against	Code Available	Paper	Method Family	Task	Protected Attribute	Notion of Fairness	Theoretical Results	Experimental Results	Comparison Against	Code Available
[16]	PreP, InP	BC, MC, R	B, C	DP, EOp	✓				[97]	PreP	BC	B, C	DD	✓			✓
[164]	InP	BC	B, C	MMD	✓		[76, 197]		[158]	PostP	BC	B, C	α -P	✓			✓
[194]	InP	BC, MC, R	B, C	ValU, AbsU, UeU, OeU	✓				[23]	PreP	BC	B, C	PredD			[97]	
[137]	PreP	BC, MC	B	SP, NDI	✓				[22]	PreP, InP, PostP	BC	B	DS (CVS)	✓			✓
[149]	InP	BC	B, C	CC	✓		[197]		[100]	InP, PostP	BC	B	DS	✓		[22, 23, 97]	✓
[163]	PostP	BC	B	REOC	✓		[76, 197]		[98]	PreP	BC	B, C	DS	✓		[23, 97]	✓
[1]	PostP	BC, MC	B, C, N	II	✓		[83]	✓	[126]	PreP	BC, MC	B, C, N	k-ND	✓			✓
[25]	PreP	BC	B, C	α -P, DP	✓		[202]		[210]	PreP	BC	B	ECD	✓		[23, 98]	✓
[37]	InP	BC	B, C	ESP, ECSP, EPE	✓				[72]	PreP	BC	B, C	ELB	✓			✓
[107]	InP	BC	B, C	SP, FPSF	✓		[2]		[105]	InP	BC	B, C	PI	✓			✓
[80]	InP	BC	B, C, N	α -MC	✓				[50]	InP	BC	B, C, N	FLP, SP	✓		[22]	✓
[69]	InP	BC	B	ESP, EPE	✓				[71]	PreP	BC	B, C	ELB	✓			✓
[208]	PreP, PostP	BC	B	DiscrP	✓		[76, 197]		[101]	InP	BC	B, C	DS	✓		[22, 99, 100]	✓
[4]	InP	BC, MC	B, C	Eod	✓				[99]	PreP	BC	B	DS	✓		[22, 100]	✓
[49]	PreP, InP	BC	B, C	ϵ -LF	✓		[76, 197]		[131]	PreP	BC	B, C	DPD	✓			✓
[2]	PostP	BC	B, C	DP, Eod	✓		[76, 99]		[73]	PostP	BC	B, C	α -P	✓			✓
[64]	InP	MAB	B, C	FLP	✓				[24]	InP	BC, MC, R	B, C	MD, AUC	✓			✓
[78]	InP	BC, MC, R	B, C, N	DA	✓				[106]	InP	BC	B	FF	✓		[22]	✓
[109]	PostP	BC	B, C, N	MMC	✓				[202]	PreP, InP	BC	B	SP	✓		[97, 103]	✓
[129]	InP	BC	B, C	Eod	✓				[102]	PreP, PostP	BC	B	ECD	✓		[23, 98]	✓
[128]	PreP	BC	B	DP, EOp, Eod	✓		[52]		[132]	PreP	BC	B, C	α -P	✓			✓
[145]	InP	BC, MC	B, C	DP, Eod	✓				[74]	PreP	BC	B, C	α -P	✓			✓
[203]	PreP, InP	BC, MC, R	B, C	DP, Eod	✓		[2, 128, 145, 198, 203]		[56]	PreP	BC	B	DI, ϵ -F	✓		[97, 103, 202]	✓
[133]	PreP, InP	BC, MC, R	B, C	PSE, NDE	✓				[123]	PreP	BC, MC, R	B, C, N	DP, MMD	✓		[202]	✓
[143]	InP	BC, MC, R	B, C	MDDS, k-DS, MaxD, IS	✓		[97, 103, 202]		[52]	PreP, InP	BC	B	SP	✓			✓
[168]	PostP	BC, MC, R	B, C, N	MDDS, IS	✓		[52, 103, 123, 168, 202]		[60]	InP	BC, MC, R	B, C, N	η -N	✓		[22, 104]	✓
[167]	PreP, InP	BC	B	(α, γ) -AMF	✓				[75]	PreP	BC	B, C	α -P	✓			✓
[196]	InP	BC, MC, R	B, C, N	DP, Eod	✓				[125]	InP	BC, MC	B, C	DCI	✓			✓
[185]	PreP, InP	BC, MC, R	B, C	DI	✓		[56]		[55]	PreP	BC	B, C	SP, DI	✓		[202]	✓
[67]	PreP	BC	B	MI, Eod	✓				[57]	PreP, InP	BC	B	DP, RBIF	✓		[202]	✓
[63]	PreP, InP	BC, MC, R	B, C	L-ASP, L-AEOP	✓				[76]	PostP	BC	B, C	EOp, Eod	✓			✓
[5]	PostP	BC	B	MI, KL-D	✓				[65]	InP	BC	B, C	FRC	✓		[198]	✓
[39]	InP	BC, MC	B, C	MI, KL-D	✓				[96]	InP	BC	B, C	δ -F	✓			✓
[40]	InP	BC, MC	B, C	MI	✓				[95]	InP	BC	B, C	δ -F	✓			✓
[81]	InP	BC, MC, R	B, C, N	R-EOP, e-EOP	✓		[81]		[58]	PostP	BC, MC, R	B, C, N	RRB	✓		[97, 100, 202]	✓
[114]	InP	BC, MC, R	B, C, N	GEI	✓		[197]		[124]	PreP	BC, MC, R	B, C, N	FairPR	✓			✓
[186]	PreP	BC, MC, R	B, C, N	WD	✓				[94]	InP	BC, MC, R	B, C, N	FEO, F-SEO	✓			✓
[178]	PreP	BC, MC, R	B, C, N	WD	✓				[190]	PostP	BC, MC	B, C	LLI	✓			✓
[82]	InP	BC, MC, R	B, C	GFE	✓				[206]	PreP	BC	B, C	TFE	✓		[56, 210]	✓
[179]	InP	BC, MC, R	B, C	PSE	✓				[198]	InP	BC	B, C	DBC	✓		[98, 103]	✓
[201]	PreP	BC, MC, R	B, C	DEO	✓				[51]	InP	BC, MC	B, C	BL	✓			✓
[59]	InP, PostP	BC, MC, R	B, C	ϵ -LF	✓				[14]	InP	BC, MC, R	B	CPIF, CPGF	✓			✓
[144]	InP	BC, MC, R	B, C	α -MC	✓				[88]	InP	BC	B, C	MDF, α -CF, α -AF	✓			✓
[33]	PostP	BC	B	HGR	✓		[13, 49]		[93]	PreP	BC	B, C, R	FairPR	✓			✓
[86]	InP	BC	B	EOD, EOp	✓		[76]		[108]	InP	BC, MC	B, C	ProxD, PDE, IPD	✓			✓
[110]	PostP	BC	B	ϵ -LGF	✓		[197, 198]		[113]	PreP, InP	BC, MC, R	B, C, N	P-R	✓		[24, 56, 202]	✓
[134]	InP	BC, MC, R	B, C, N	DEOp, DEOd	✓				[117]	InP	BC, MC, R	B, C	CF	✓			✓
[147]	PostP	BC	B	DP	✓		[52, 128]		[171]	InP	BC, MC, R	B, C	ϵ, δ -ACF	✓			✓
[153]	PreP, InP	BC, MC, R	B, C, N	SR	✓		[49]		[189]	InP, PostP	BC, MC	B	α -D	✓			✓
[151]	InP	BC	B, C	PSCF, MMD	✓				[199]	InP, PostP	BC, MC	B, C	PrefI, PrefT	✓			✓
[152]	PreP, InP	BC, MC, R	B, C	SDP, SPDD, WD	✓		[76]		[197]	InP	BC	B, C	DM	✓			✓
[188]	InP	BC, MC, R	B, C, N	DBC, DI, DM	✓		[37, 49, 76, 98, 103]		[192]	InP	BC	B	rND, rKL, rRD	✓			✓
[26]	InP	BC, MC, R	B, C	SP, DI, FLP	✓				[207]	PreP	BC	B	CDP	✓			✓
[91]	InP, PostP	BC	B, C, R	Eod	✓				[13]	InP	BC	B	AVD, SD	✓		[56, 210]	✓
[200]	InP	BC, MC	B, C		✓				[28]	PreP	C	B	Bal	✓		[76, 197]	✓
[138]	PreP	BC	B		✓				[160]	PreP, InP	BC, MC, R	B, C, N	HSIC	✓			✓
[89]	PostP, InPro	BC, MC	B		✓												

Pre-processing: Fair Representation Learning

- Approaches
 - Awareness
 - Representation Learning
 - Re-weighting
 - Resampling → Over/Under – SMOTE, etc



- $Z \rightarrow$ Latent representation
 - $\max_{Z=g(X)} I(X; Z)$
 - subject to $I(A; Z) < e$
 - $S \perp A$

$$\alpha \text{Loss}_{\text{similarity}} + \beta \text{Loss}_{\text{fairness}} + \gamma \text{Loss}_{\text{prediction}}$$

- Strict approach → Optimizes only Statistical Parity or Individual Fairness
 - Info of Y not used
- No need to access A at test time nor Y at representation time
- If Y is used → **hybrid** approach with potential better results [$S \perp A | Y$ and $Y \perp A | S$]

$$D = \{(a_i, x_i, y_i)\}_{i=1}^N$$

$$x_i \in R^d$$

$$g: R^d \rightarrow R^r \text{ i.e., } g(x_i) = z_i$$

$$z_i \in R^z$$

$$z_i \perp a_i$$

$$Z \perp A$$

If model involved [hybrid]:
 $f(g(X))$

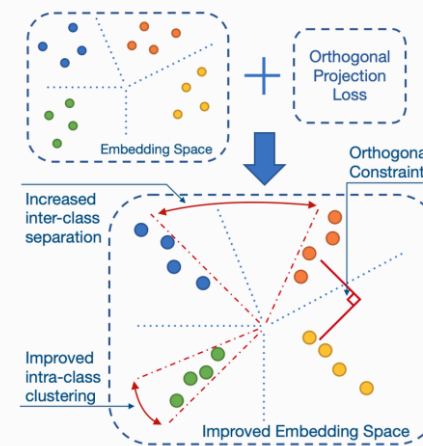



Pre-processing: Fair Representation Learning

Lots of works using NN

$\max I(A, g(X))$ while $\min I(A, g(X))$ and may $\max(g(X), Y)$

$$Loss_C = |x - x'|^2 - \lambda Loss_A(z)$$



$$Loss_C = \alpha |x - x'|^2 + \lambda Loss_A(Z_A) + \beta L_{\perp}$$

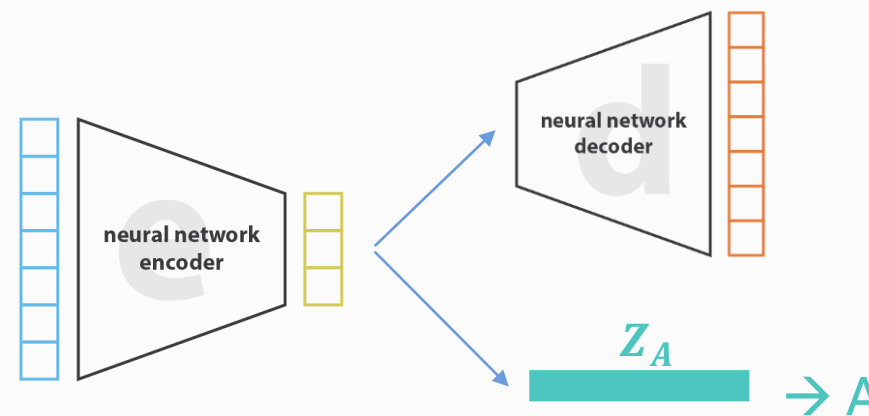
$$\alpha Loss_{similarity} + \beta Loss_{fairness} + \gamma Loss_{prediction}$$

`aif360.algorithms.preprocessing.LFR`

```
class aif360.algorithms.preprocessing.LFR(unprivileged_groups, privileged_groups, k=5, Ax=0.01, Ay=1.0, Az=50.0,
print_interval=250, verbose=0, seed=None) [source]
```

Learning fair representations is a pre-processing technique that finds a latent representation which encodes the data well but obfuscates information about protected attributes [2] ... rubric:: References

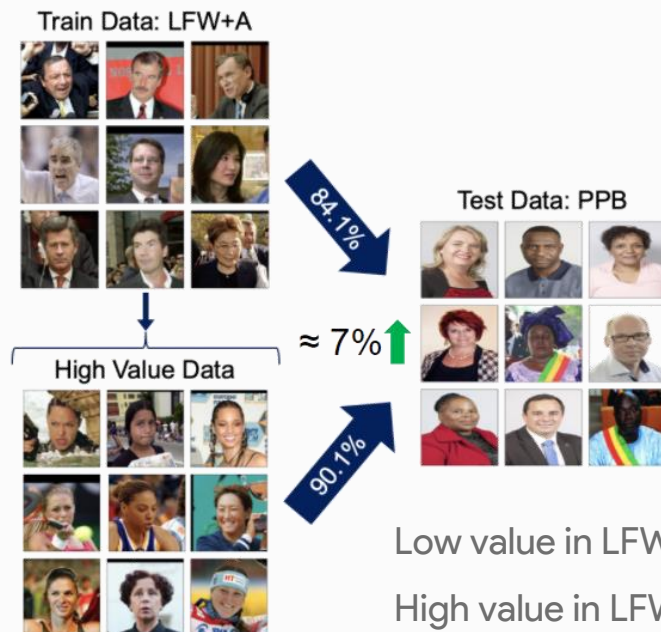
[2] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork, "Learning Fair Representations." International Conference on Machine Learning, 2013.



Pre-processing: Reweighting

- Weight the examples (group, label) to ensure fairness in classification
- Unbalanced learning-related → e.g., Fair-SMOTE
- Advanced example → SHAPLEY values

Domain adaptation: gender detection



`aif360.algorithms.preprocessing.Reweighting`

```
class aif360.algorithms.preprocessing.Reweighting(unprivileged_groups, privileged_groups) [source]
```

Reweighting is a preprocessing technique that Weights the examples in each (group, label) combination differently to ensure fairness before classification [4].

References

- [4] F. Kamiran and T. Calders, "Data Preprocessing Techniques for Classification without Discrimination," Knowledge and Information Systems, 2012.



In-processing

```
class aif360.algorithms.inprocessing.PrejudiceRemover(eta=1.0, sensitive_attr="", class_attr="")
```

Prejudice remover is an in-processing technique that adds a discrimination-aware regularization term to the learning objective [6].

References

[6] T. Kamishima, S. Akaho, H. Asoh, and J. Sakuma, "Fairness-Aware Classifier with Prejudice Remover Regularizer," Joint European Conference on Machine Learning and Knowledge Discovery in Databases, 2012.

- Add penalty to objective function during learning → Regularizer
- Prior work: **Prejudice remover** (Kamishima et al., 2012)
 - **Prejudice remover regularizer**: Based on the **degree of indirect prejudice** (PI)

Mutual Information between Y and S

$$PI = \sum_{(y,a) \in \mathcal{D}} \hat{P}[y, s] \ln \frac{\hat{P}[y, s]}{\hat{P}[y] \hat{P}[s]}$$

Prejudice remover regularizer

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

S: protected/sensitive attribute

$$\sum_{(y_i, \mathbf{x}_i, s_i)} \ln \mathcal{M}[y_i|\mathbf{x}_i, s_i; \Theta] + \eta R_{PR}(\mathcal{D}, \Theta) + \frac{\lambda}{2} \sum_{s \in \mathcal{S}} \|\mathbf{w}_s\|_2^2$$

Logistic Regression

Prejudice remover regularizer

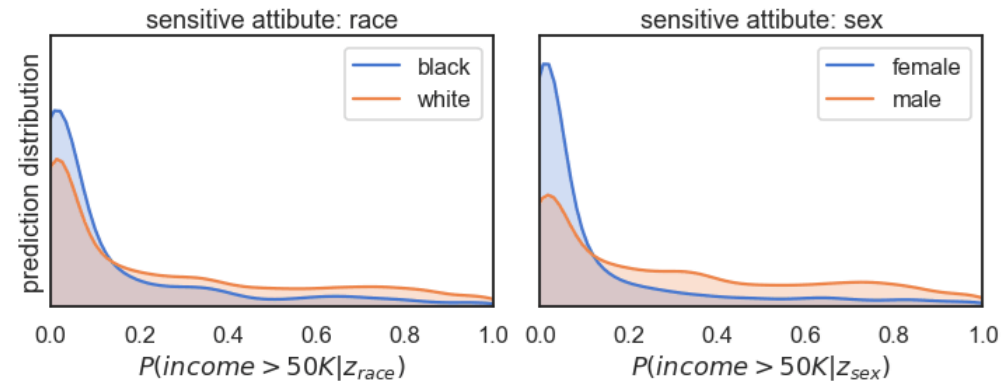
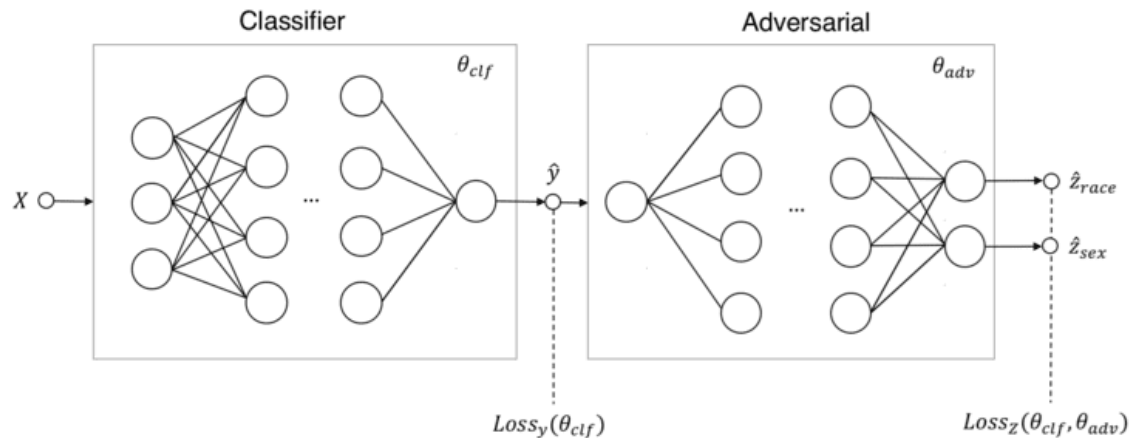
L2 Regularization



In-processing: Adversarial debiasing

- Make the best possible predictions while ensuring that A cannot be derived from them
 - Demographic Parity
 - Adversary gets \hat{Y}
 - Equality Of Odds
 - Adversary gets \hat{Y} and Y
 - Equality Of Opportunity
 - On a given class $y \rightarrow$ restrict adversary's training set to X where $Y = y$

$$\min_{\theta_{clf}} [Loss_y(\theta_{clf}) - \lambda Loss_z(\theta_{clf}, \theta_{adv})]$$



Training iteration #1

Prediction performance:
 - ROC AUC: 0.90
 - Accuracy: 84.9

Satisfied p%-rules:
 - race: 44%-rule
 - sex: 35%-rule

`aif360.algorithms.inprocessing.AdversarialDebiasing`

```
class aif360.algorithms.inprocessing.AdversarialDebiasing(unprivileged_groups, privileged_groups, scope_name,
    sess, seed=None, adversary_loss_weight=0.1, num_epochs=50, batch_size=128, classifier_num_hidden_units=200, debias=True)
[source]
```

$$p\%rule = \min\left(\frac{P\{\hat{Y} = 1 | A = a\}}{P\{\hat{Y} = 1 | A = b\}}, \frac{P\{\hat{Y} = 1 | A = b\}}{P\{\hat{Y} = 1 | A = a\}}\right) \geq \frac{p}{100}$$





Post-processing

- Deal with output predictions of the model
 - Useful in black-box models or if we **don't have access to the train** pipeline → NO retraining
 - Find a proper threshold using the output for each group
 - Require A to be available in testing → compliance risk

`aif360.algorithms.postprocessing.EqOddsPostprocessing`

```
class aif360.algorithms.postprocessing.EqOddsPostprocessing(unprivileged_groups, privileged_groups, seed=None)
[source]
```

Equalized odds postprocessing is a post-processing technique that solves a linear program to find probabilities with which to change output labels to optimize equalized odds [8] [9].



`aif360.algorithms.postprocessing.RejectOptionClassification`

```
class aif360.algorithms.postprocessing.RejectOptionClassification(unprivileged_groups, privileged_groups,
low_class_thresh=0.01, high_class_thresh=0.99, num_class_thresh=100, num_ROC_margin=50, metric_name='Statistical parity
difference', metric_ub=0.05, metric_lb=-0.05)
[source]
```

Reject option classification is a postprocessing technique that gives favorable outcomes to unprivileged groups and unfavorable outcomes to privileged groups in a confidence band around the decision boundary with the highest uncertainty [10].



More prominent approaches

Causality

Domain-specific
Images
Text
Graphs

Discriminatory Transfer
Multitask Fairness

XAI
Interpretability

Game theoretical
approaches





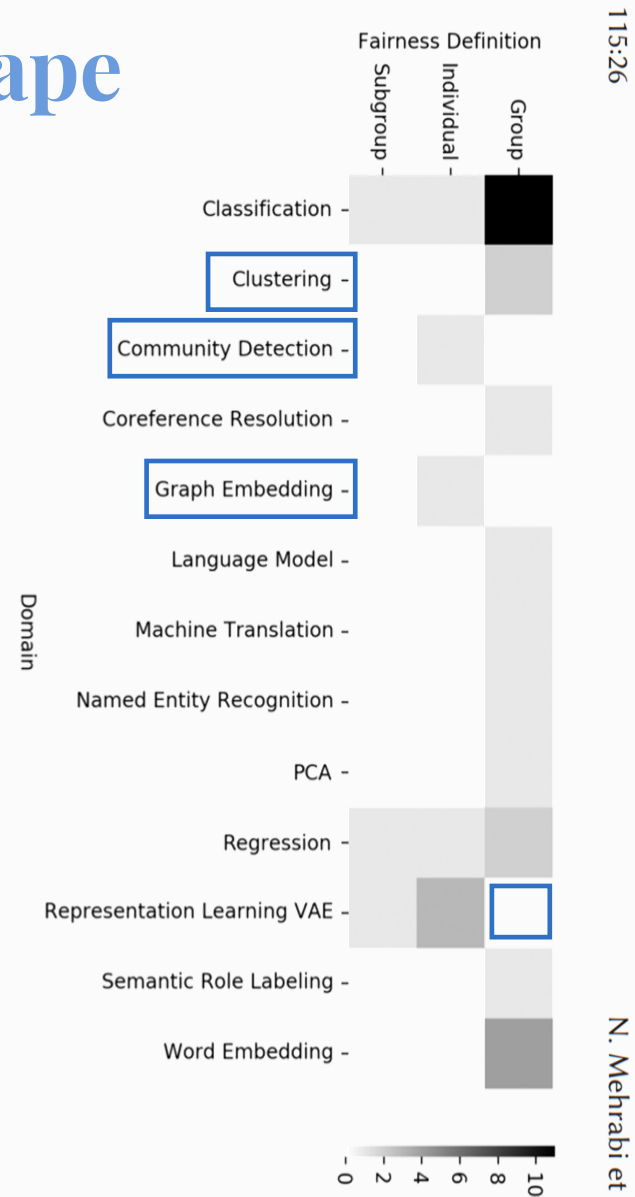
Current situation

Quick view on graphs & causality

Current landscape

Table 2. List of Papers Targeting and Talking about Bias and Fairness in Different Areas

Area	Reference(s)
Classification	[25, 49, 57, 63, 69, 73, 75, 78, 85, 102, 118, 143, 150, 151, 155]
Regression	[1, 14]
PCA	[133]
Community detection	[101]
Clustering	[8, 31]
Graph embedding	[22]
Causal inference	[82, 95, 111, 112, 123, 156, 160, 161]
Variational auto encoders	[5, 42, 96, 108]
Adversarial learning	[90, 152]
Word embedding	[20, 58, 165] [23, 162]
Coreference resolution	[130, 164]
Language model	[21]
Sentence embedding	[99]
Machine translation	[52]
Semantic role labeling	[163]
Named Entity Recognition	[100]



115:26

N. Mehrabi et al.




Graphs & Fairness


What fairness need? <i>Defining – detecting – imposing – apply</i>	How can Graphs help?
Capture Individual similarity	<ul style="list-style-type: none"> – Natural node pairwise distance – Structural similarity – Role similarity – Graph Representation Learning (<i>for Nodes & Edges & Graphs</i>)
Capture Group Structure-Behavior	<ul style="list-style-type: none"> – Community detection – Inherent data structure in graphs – Structural Analysis (e.g., Laplacian)
Capture deeper relationships between data	<ul style="list-style-type: none"> – Node – Edge - classification – Missing link prediction – Message passing – Information Flow – Rewiring – Changing graph structure
Different label bias problems	<ul style="list-style-type: none"> – Semi-Supervised Learning <i>i.e., help with labels we cannot see</i>
Causality	<ul style="list-style-type: none"> – Strong theory behind graphs – GNN \rightarrow SCM
Applied to social problems	<ul style="list-style-type: none"> – Network is the natural structure of data – Also, everything can be modeled as a graph
XAI	<ul style="list-style-type: none"> – Interpretable by design – Friendly straightforward graph explanations – Great XAI graph-based



Causality



- Previous definitions relies on **Joint probabilities of (X,Y,S,A)**
 - Reactive vision: take everything as given about the world as it is → Observational
- Can we capture social context? **Let's use causal models**
 - How changes in variables propagate in a system, be it natural, engineered or social
 - What should we do when there's no direct effect?



Exploit Structural Causal Model properties to look for biases Neal, B. (2020)

Definition 4.2 (Structural Causal Model (SCM)) *A structural causal model is a tuple of the following sets:*

1. A set of endogenous variables V
2. A set of exogenous variables U
3. A set of functions f , one to generate each endogenous variable as a function of other variables

$$B := f_B(A, U_B)$$

$$M : C := f_C(A, B, U_C)$$

$$D := f_D(A, C, U_D)$$





Figure 4.8: Graph for the structural equations in Equation 4.24.



J. Pearl, 2009 Causality: Models, Reasoning and Inference, 2nd ed. New York, NY, USA: Cambridge University Press,
 Neal, B. (2020). Introduction to causal inference from a ML perspective. *Book (draft)*. [https://www.bradyneal.com/Introduction to Causal Inference-Dec17 2020-Neal.pdf](https://www.bradyneal.com/Introduction%20to%20Causal%20Inference-Dec17%2020-Neal.pdf)
 Kusner, M. J., Loftus, J. R., Russell, C., & Silva, R. (2017). Counterfactual fairness.
 Loftus, J. R., Russell, C., Kusner, M. J., & Silva, R. (2018). Causal reasoning for algorithmic fairness
 Makhlof, K., Zhioua, S., & Palamidessi, C. (2020). Survey on Causal-based Machine Learning Fairness Notions. arXiv preprint arXiv:2010.09553.
 Kilbertus, N., Rojas-Carulla, M., Parascandolo, G., Hardt, M., Janzing, D., & Schölkopf, B. (2017). Avoiding discrimination through causal reasoning
 Zhang, J., & Bareinboim, E. (2018, April). Fairness in decision-making—the causal explanation formula. In Thirty-Second AAAI
 Wu, Y. (2020). Achieving Causal Fairness in Machine Learning
 S. Chiappa. 2019, Path-specific counterfactual fairness. Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)
 Chiappa, S., & Isaac, W. S. (2018,). A causal bayesian networks viewpoint on fairness. In IFIP International Summer School on Privacy and Identity Management
 Fairness – Moritz Hardt – Part 2 – MLS2020 – <https://www.youtube.com/watch?v=9oNVFQ9lIPc&t=1449s>

Causal fairness
 criteria and
 path-specific
 effects



Causality: examples

- **Counterfactual fairness:**
 - Outcome probability in factual world = the counterfactual world
 - How would the world have to be different for a desirable output to occur?
 - *What would have happened if I were different?*
- **Causal Representation Learning**
- **Algorithmic Recourse**
 - → Causality +XAI → explanations + recommendations
 - **Actionable** feedback about how to change the outcomes of ML models
 - *“To have your loan approved, you would need to increase your income by \$10,000 per year”*

*“**Counterfactuals** explain complex models with the **use of examples...**
...while **recourse** tries to **find actions** that leads to a better outcome”* Annabelle Redelmeier

	Counterfactuals	Recourse
Optimization function	Loss function	Cost function
Algorithm solves for...	Vectors/Individuals (x)	Actions (δ)
Ultimate goal	Explain a model	Solve for actions to achieve “recourse”





Libraries

Libraries

IBM Research Trusted AI

AI Fairness 360





Datasets

Benchmarking datasets

- Big amount of tabular dataset in all domains




- Every dataset may have intrinsic bias

Images

Text

School Effectiveness	[66]	15362	9	Ethnicity, Gender	R
Heart Disease	[90]	303	75	Age, Gender	MC, R
German Credit	[85]	1K	20	Age, Gender/Marital-Stat	MC
Census/Adult Income	[112]	48842	14	Age, Ethnicity, Gender, Native-Country	BC
Contraceptive Method Choice	[121]	1473	9	Age, Religion	MC
Law School Admission	[187]	21792	5	Ethnicity, Gender	R
Arrhythmia	[70]	452	279	Age, Gender	MC
Communities & crime	[169]	1994	128	Ethnicity	R
Wine Quality	[154]	4898	13	Color	MC, R
Heritage Health	[146]	≈60K	≈20	Age, Gender	MC, R
Stop, Question & Frisk	[45]	84868	≈100	Age, Ethnicity, Gender	BC, MC
Bank Marketing	[142]	45211	17-20	Age	BC
Diabetes US	[181]	101768	55	Age, Ethnicity	BC, MC
Student Performance	[38]	649	33	Age, Gender	R
CelebA Faces	[122]	≈200K	40	Gender Skin-Paleness, Youth	BC
xAPI Students Perf.	[6]	480	16	Gender, Nationality, Native-Country	MC
Chicago Faces	[127]	597	5	Ethnicity, Gender	MC
Credit Card Default	[195]	30K	24	Age, Gender	BC
COMPAS	[119]	11758	36	Age, Ethnicity, Gender	BC, MC
MovieLens	[77]	100K	≈20	Age, Gender	R
Drug Consumption	[54]	1885	32	Age, Ethnicity, Gender, Country	MC
Student Academics Perf.	[87]	300	22	Caste, Gender	MC
NLSY	[148]	≈10K		Birth-date, Ethnicity, Gender	BC, MC, R
Diversity in Faces	[140]	1 M	47	Age, Gender	MC, R



Pilot Parliaments Benchmark

Retiring Adult: New Datasets for Fair Machine Learning

Frances Ding*
UC Berkeley

Moritz Hardt*
UC Berkeley

John Miller*
UC Berkeley

Ludwig Schmidt*
Toyota Research Institute

Quy, T. L., Roy, A., Iosifidis, V., & Ntoutsis, E. (2021). A survey on datasets for fairness-aware machine learning. arXiv

Oneto, L. (2020). Learning fair models and representations. *Intelligenza Artificiale*, 14(1), 125-152

Barocas, S., Hardt, M., & Narayanan, A. (2017). *Fairness in machine learning*. Nips tutorial, 1, 2017

Majumder, S., Chakraborty, J., Bai, G. R., Stolee, K. T., & Menzies, T. (2021). Fair Enough: Searching for Sufficient Measures of Fairness. preprint arXiv:2110.13029.

<http://gendershades.org/overview.html> - <https://nips.cc/media/neurips-2021/Slides/26854.pdf>





History and conceptual point of view

What should we learn from the past fairness research?
What other conceptual concerns should we consider?

Fairness beginning: 60's & 70's

Shout out to pioneers

1966	1968	1971	1971	1973	1976
Guion	Cleary	Thorndike	Darlington	Cole	Peterson and Novick

- 60's: start to quantify bias
- 70's: From unfairness to Fairness
 - FP & FN rates
 - Fair use of the test, rather than the scores themselves
- Mid 70's: halt 😞, Why?
 - **No analyses to unequivocally indicate fairness**
 - **No clear procedures to avoid unfairness**
 - **Disagreement in views of fairness** view between professionals and general public

*“Fairness actually obscure the fundamental problem, which is to find some rational basis for **providing compensatory treatment for the disadvantaged**” (Melvin R Novick et al. 1976)*
- Rediscovered by ML around 13 year ago (Calders et al. 2009)

What should we learn?

- DON'T reinvent the wheel
- DON'T forget actual objective
 - compensatory treatment to disadvantaged
- DON'T get stacked in discussions far from real-world problems
- DON'T be far from **practical needs** of society, politics & law
- Work in political and law implication
- Relating fairness debates to ethical theories and value systems
- ML Fairness community should be more aware of our own implicit cultural biases




Fair ML and law

“Careful attention should be paid to **legal and public concerns about fairness**. The experiences of the test fairness field suggest that in the coming years, **courts may start ruling on the fairness of ML models**. Therefore, **if technical definitions of fairness stay too far from the public’s perceptions of fairness**, then the **political will to use scientific contributions in advance of public policy may be difficult to obtain**”

Hutchinson, B., & Mitchell, M. 2019.

50 years of test (un) fairness: Lessons for machine learning. FAccT 2019



Other cultural and conceptual challenges

Even we are looking for bias, we are **inducing bias**

CONTEXT MATTERS
Quantitative techniques
+ policy-level questions

Make methods flexible to **adapt to each situation, context and use**

PUBLIC'S NOTION OF FAIRNESS
Explicitly connect fairness criteria to different **socio-cultural and philosophical values**

Try to **unify fairness** definition and framework

Politics and law **implication**

Remind: Fairness and unfairness are related but different concepts

Make Fair ML research **accessible** to general public, other researchers

From equality to equity
Give each one the resources that each one need to reach to the same point

Example of conceptual bias: Why groups should be treated as discrete categories?

- Most definitions of protected attribute-group relies on **categoric division** → **implicit cultural bias & unstable social construct**
- Other possibility: intersectional modelling → **Protected attribute as continuous variables**
 - Quantify fairness along one dimension (e.g., age) conditioned on another dimension (e.g., skin tone)

e.g., Use Computer vision clustering of skin tones instead of pre-defined ethnics






Wrapping up

Conclusion

- **Don't feel overwhelmed** by the big amount methods and measures!
 - Method depends on task, and technical context
 - Definitions and metrics depends on the context
 - Development and relationship of the measures with ethics → Now you choose context – experts – social and ethical analysis
- More work needed in **ethical-cultural aspect**
 - Equity → Considering individual resources
 - Continual protected attributes
 - Social-Law-Political needs close relationship
- **Technical takeaways**
 - Beyond observational → **Causality**
 - Deep structural data relationship → **Graphs**




Bibliography

More references in each slide

- M. Hardt, E. Price, N. Srebro, Equality of opportunity in supervised learning, Advances in Neural Information Processing Systems (2016).
- Cynthia Dwork, et al. 2012. Fairness Through Awareness. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference
- Alexandra Chouldechova. 2016. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. Big Data.
- Verma, J. Rubin, Fairness definitions explained, IEEE/ACM International Workshop on Software Fairness (2018) 1–7.
- Richard Berka, Hoda Heidaric, Shahin Jabbaric, Michael Kearns, and Aaron Roth. 2017. Fairness in Criminal Justice Risk Assessments: The State of the Art.
- Alexandra Chouldechova. 2016. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. Big Data (2016)
- Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2017. Inherent Trade-Offs in the Fair Determination of Risk Scores. In ITCS
- Ben Hutchinson and Margaret Mitchell. 2019. 50 Years of Test (Un)fairness: Lessons for Machine Learning.
- M.J. Kusner, J. Loftus, C. Russell and R. Silva, Counterfactual fairness, In Neural Information Processing Systems, (2017)
- Barocas, S., Hardt, M., & Narayanan, A. (2017). Fairness in machine learning. Nips tutorial, 1, 2017
- Shira Mitchell. 2018. Reflection on quantitative fairness. Web Book
- Majumder, S., Chakraborty, J., Bai, G. R., Stolee, K. T., & Menzies, T. (2021). Fair Enough: Searching for Sufficient Measures of Fairness. arXiv preprint arXiv:2110.13029.
- L. Oneto, Learning fair models and representations, Intelligenza Artificiale 14 (1) (2020) 151–178.
- Castelnovo, A., Crupi, R., Greco, G., & Regoli, D. (2021). The zoo of Fairness metrics in Machine Learning. arXiv
- Franco, D., Navarin, N., Donini, M., Anguita, D., & Oneto, L. (2022). Deep fair models for complex data: Graphs labeling and explainable face recognition. Neurocomputing, 470
- A.F. Winfield, K. Michael, J. Pitt, V. Evers, Machine ethics: the design and governance of ethical ai and autonomous systems, Proceedings of the IEEE 107 (2019) 509–517
- D. Gunning, Explainable artificial intelligence (xai), Defense Advanced Research Projects Agency (DARPA), nd Web 2 (2)
- Majumder, S., Chakraborty, J., Bai, G. R., Stolee, K. T., & Menzies, T. (2021). Fair Enough: Searching for Sufficient Measures of Fairness. preprint arXiv:2110.13029.
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. Calif. L. Rev., 104, 671
- Lim Swee Kiat. Retrieved December 2021. Machines go Wrong. <https://machinesgonewrong.com/fairness/>
- Manuel Gomez Rodriguez et al. (2020). Human-Centric Machine Learning Feedback loops, Human-AI Collaboration and Strategic Behavior [Link]. Web
- Corbett-Davies & Goel. (2018). The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning
- Mehrabi, N., et al. (2021). A survey on bias and fairness in machine learning. ACM Computing Surveys (CSUR), 54(6), 1-35
- 2017. CS 294: Fairness in Machine Learning. <https://fairmlclass.github.io> (2017). Online; accessed February 2018





Talk in the scope of the project:

Achieving Fair, Accountable and Transparent Machine Learning Models through Graph Theory and Causality

Thesis in Progress by PhD Student Adrián Arnaiz Rodríguez

PhD Nuria Oliver

PhD Francisco Escolano

PhD Manuel Gómez Rodríguez



ellis
ALICANTE unit 

Thank you!

Q's & feedback?

adrian@ellisalicante.org



@arnaiztech



AdrianArnaiz

