

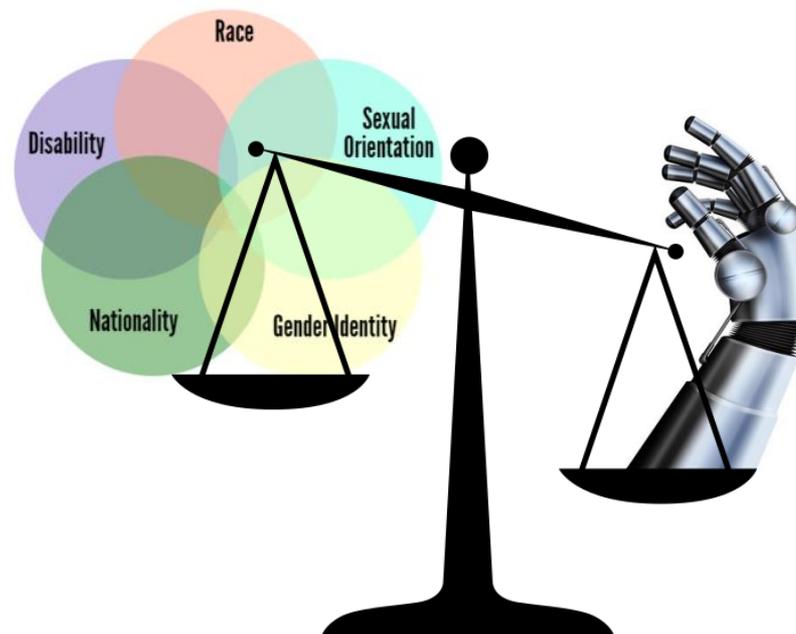
Differential Fairness for Machine Learning and Artificial Intelligence Systems: Unbiased Decisions with Biased Data

James Foulds

Department of Information Systems
University of Maryland, Baltimore County

(Joint work with Rashidul Islam, Kamrun Keya, Shimei Pan)

Artificial Intelligence Maryland Meetup, Emerging Technology Centers, Nov 14, 2018



Machine Learning

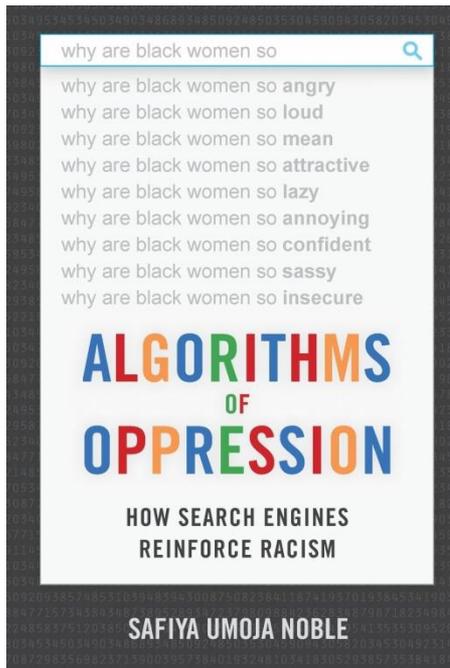
- **Machine learning algorithms**, which make predictions based on data, are having an increasing impact on our daily lives.
- Example: **credit scoring** - predicting whether you will repay or default on a loan

# Late Payments	% of available credit used	Previous defaults?	Employed?	...	Repay Loan?
Feature vector X					Class label Y

- Given the feature vector, the algorithm learns to predict the class label
- The models are “trained” on many labeled feature vectors
- This is called **classification**, an instance of **supervised machine learning**

Fairness in Machine Learning

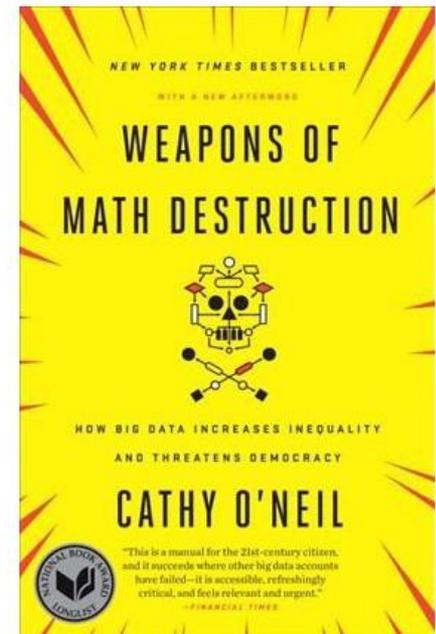
- There is growing awareness that **biases inherent in data** can lead the behavior of machine learning algorithms to **discriminate against certain populations.**



Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights

Executive Office of the President

May 2016



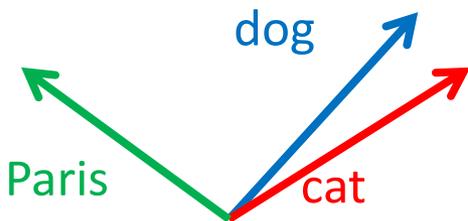
Bias in Predicting Future Criminals

- Correctional Offender Management Profiling for Alternative Sanctions (**COMPAS**), by Northpointe company
 - An algorithmic system for predicting risk of re-offending in criminal justice
 - Used for sentencing decisions across the U.S.
- ProPublica study (Angwin et al., 2016):
 - **COMPAS almost twice as likely to incorrectly predict re-offending for African Americans** than for white people. Similarly much more likely to incorrectly predict that white people would not re-offend than for African Americans
 - Northpointe disputes the findings

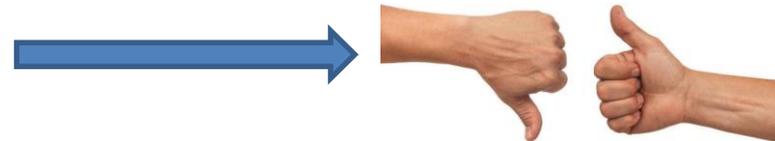
	WHITE	AFRICAN AMERICAN
Labeled Higher Risk, But Didn't Re-Offend	23.5%	44.9%
Labeled Lower Risk, Yet Did Re-Offend	47.7%	28.0%

Illustrative Example: Sentiment Analysis

- An example from “How to make a racist AI without really trying” by Rob Speer
- Application: sentiment analysis
 - Predict whether the sentiment expressed in a text is positive or negative

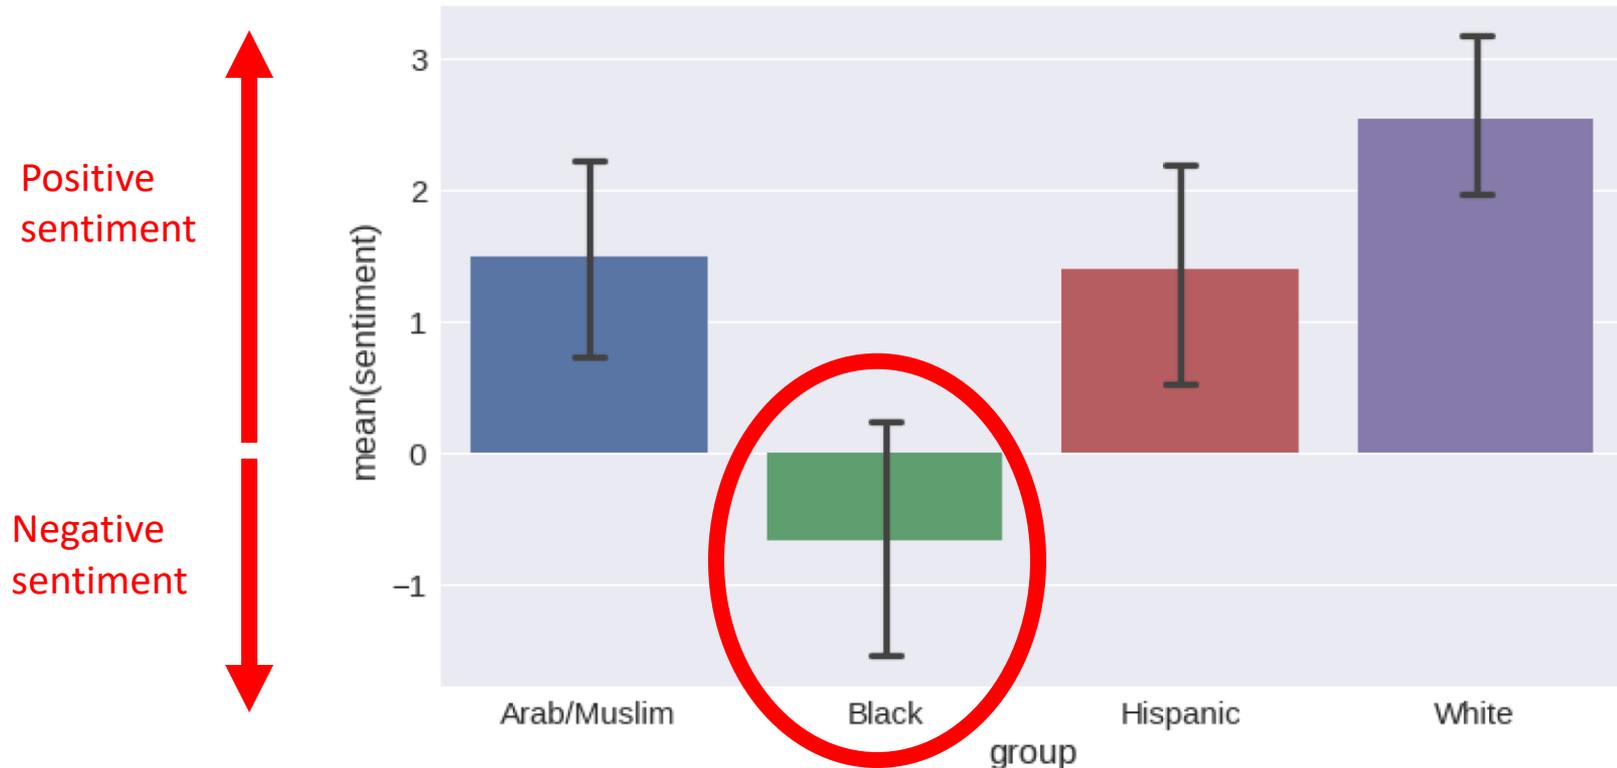


dog: (0.11, -1.5, 2.7, ...)
cat: (0.15, -1.2, 3.2, ...)
Paris: (4.5, 0.3, -2.1, ...)



“How to Make a Racist AI Without Really Trying”

- **Sentiment of stereotypical names for different race groups**
(bar plot with 95% confidence interval of means shown)



Amazon scraps secret AI recruiting tool that showed bias against women

Jeffrey Dastin

SAN FRANCISCO (Reuters) - Amazon.com Inc's ([AMZN.O](#)) machine-learning specialists uncovered a big problem: their new recruiting engine did not like women.

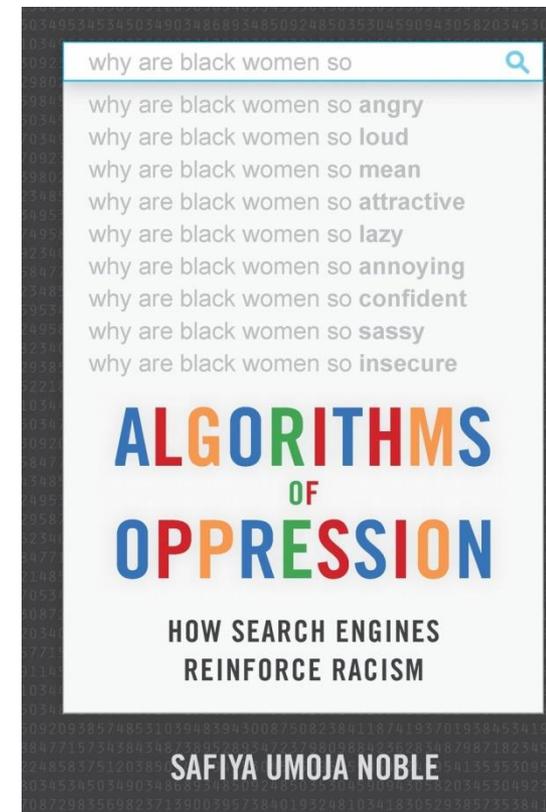


The team had been building computer programs since 2014 to review job applicants' resumes with the aim of mechanizing the search for top talent, five people familiar with the effort told Reuters.

Sources of Bias in Data

(cf. *Barocas and Selbst (2016)*)

- Data encodes **societal prejudices**
 - e.g. racism/sexism in social media data
- Data encodes **societal (dis)advantages**
 - college admissions, criminal justice
- **Less data** for minorities
- **Collection bias**
 - data from smartphones, automobiles,...
- **Intentional prejudice. Digital redlining, masking**
 - St. George's Hospital Med School encoded its existing race/gender-biased decision-making for admissions interviews in an algorithm (Lowry & McPherson, 1988)
- **Proxy variables**
 - (e.g. zip code highly correlated with race, leading classifier to unintentionally consider race)



Considerations

- **Fairness is a highly complicated socio-technical-political-legal construct**



- **Harms of representation** vs **harms of outcome** (cf. Kate Crawford, Bolukbasi et al. (2016))
- Differences between **equality** and **fairness** (Starmans and Sheskin, 2017). How to balance these?
- Whether (and how) to model underlying **differences between populations** (Simoiu et al., 2017)
- Whether to aim to correct **biases in society** as well as biases in data (**fair affirmative action**) (Dwork et al., 2012)

The Machine Learning / AI Community's Response to Fairness

- A recent explosion of research (since circa 2016)
- **Publication venues** dedicated to fairness and related issues
 - Fairness, Accountability and Transparency in ML (FAT/ML) Workshop
 - ACM FAT*
 - AAAI/ACM Conference on AI, Ethics & Society
- **Mathematical definitions, algorithms** for enforcing and measuring fairness



AAAI / ACM conference on
**ARTIFICIAL INTELLIGENCE,
ETHICS, AND SOCIETY**

Fairness and the Law: Adverse Impact Analysis

- Title VII, other anti-discrimination laws prohibit employers from intentional discrimination against employees with respect to protected characteristics
 - gender, race, color, national origin, religion
- Uniform Guidelines for Employee Selection Procedures (Equal Employment Opportunity Commission)

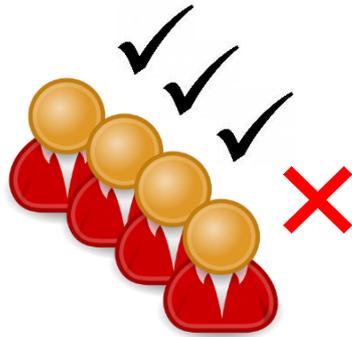
The “four-fifths rule” (a.k.a. 80% rule)

*“A **selection rate for any race, sex, or ethnic group which is less than four-fifths (4/5) (or eighty percent) of the rate for the group with the highest rate** will generally be regarded by the Federal enforcement agencies as evidence of adverse impact, while a greater than four-fifths rate will generally not be regarded by Federal enforcement agencies as evidence of adverse impact.”*

-Code of Federal Regulations 29 Part 1607 (1978)

Fairness and the Law: Adverse Impact Analysis

$$Pr(\text{hire}|\text{group A}) < 0.8 \times Pr(\text{hire}|\text{group B}) ?$$



If so, there is evidence of adverse impact

Fairness and Intersectionality

- **Intersectionality:**

systems of oppression built into society lead to **systematic disadvantages** along **intersecting dimensions**

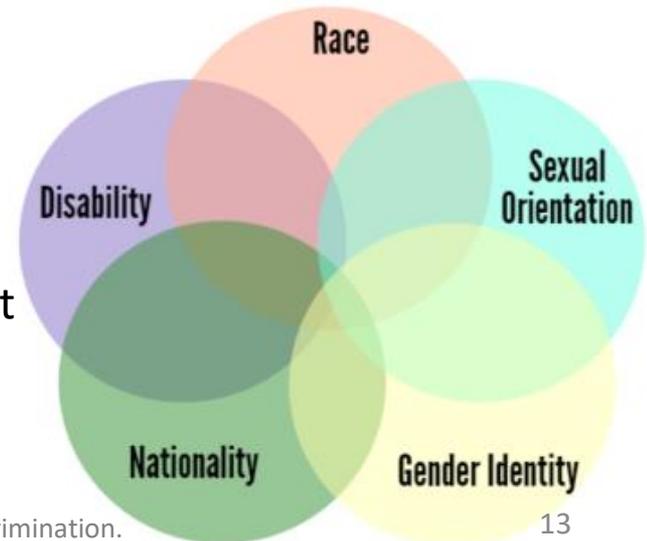
- gender, race, nationality, sexual orientation, disability status, socioeconomic class, ...

versus

- **Infra-marginality:**

attributes used by algorithm may have **different distributions**, depending on the **protected attributes**.

Algorithm should behave differently for each group, is biased if it is more inequitable than the data suggest



K. Crenshaw. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. U. Chi. Legal F., pages 139–167, 1989.

C. Simoiu, S. Corbett-Davies, S. Goel, et al. The problem of infra-marginality in outcome tests for discrimination. The Annals of Applied Statistics, 11(3):1193–1216, 2017.

Fairness and Intersectionality

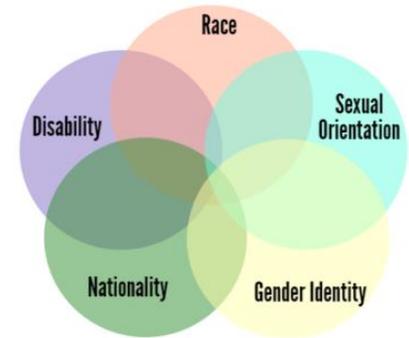
- *Subgroup fairness* (Kearns et al., 2018)
 - Aims to prevent “*fairness gerrymandering*” a.k.a. *subset targeting*, by protecting specified subgroups

Definition 2.1 (Statistical Parity (SP) Subgroup Fairness). Fix any classifier D , distribution \mathcal{P} , collection of group indicators \mathcal{G} , and parameter $\gamma \in [0, 1]$. For each $g \in \mathcal{G}$, define

$$\alpha_{SP}(g, \mathcal{P}) = \Pr_{\mathcal{P}}[g(x) = 1] \quad \text{and} \quad \beta_{SP}(g, D, \mathcal{P}) = |\text{SP}(D) - \text{SP}(D, g)|,$$

where $\text{SP}(D) = \Pr_{\mathcal{P}, D}[D(X) = 1]$ and $\text{SP}(D, g) = \Pr_{\mathcal{P}, D}[D(X) = 1 | g(x) = 1]$ denote the overall acceptance rate of D and the acceptance rate of D on group g respectively. We say that D satisfies γ -statistical parity (SP) Fairness with respect to \mathcal{P} and \mathcal{G} if for every $g \in \mathcal{G}$

$$\alpha_{SP}(g, \mathcal{P}) \beta_{SP}(g, D, \mathcal{P}) \leq \gamma.$$



- punts on small groups (in order to prove generalization)

See also *multicalibration*, a similar definition but for calibration of probabilities (Hebert-Johnson et al., 2018)

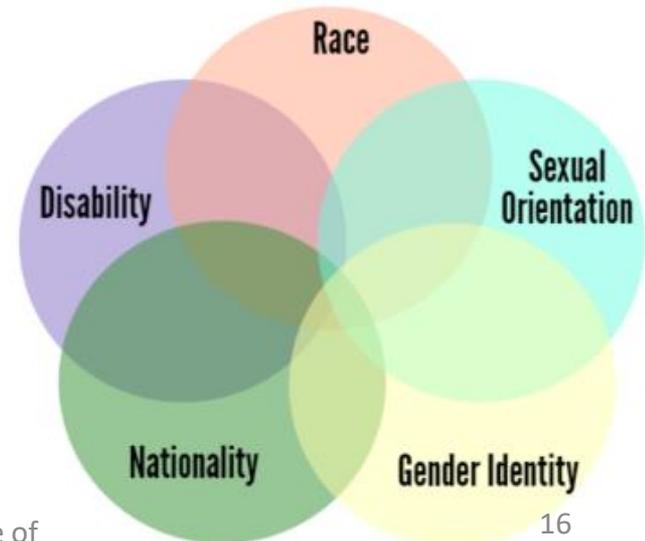
Our contributions

- We address fairness in machine learning from an **intersectional perspective**
 - Fairness definitions that respect intersectionality
 - Address untrusted vendor scenario
 - Also provide a more politically conservative option
 - Propose methods to address uncertainty with multiple protected attributes

Fairness and Intersectionality

We argue that an **intersectional definition of fairness** should satisfy:

- **Multiple protected attributes** should be considered
- **All** of the **intersecting values** of the protected attributes, e.g. *black women*, should be protected
 - We should still ensure that the individual protected attributes are protected overall, e.g. *women* are protected
- Systematic differences, due to structural oppression, are **rectified, rather than codified.**



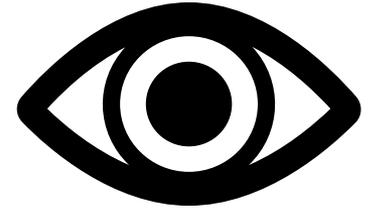
Differential Fairness (DF)

We propose a fairness definition with the following properties:

- **Measures the fairness cost of algorithms and data**
 - Can measure difference in fairness between algorithms and data: **bias amplification**
- **Privacy and economic guarantees**
 - Privacy perspective provides an **interpretation** of definition, based on **differential privacy**
- Implements **intersectionality**: e.g. fairness for (gender, race) provably ensures fairness for gender and for race separately

Essentially, differential fairness extends the 80% rule to multiple protected attributes and outcomes, and provides a privacy interpretation

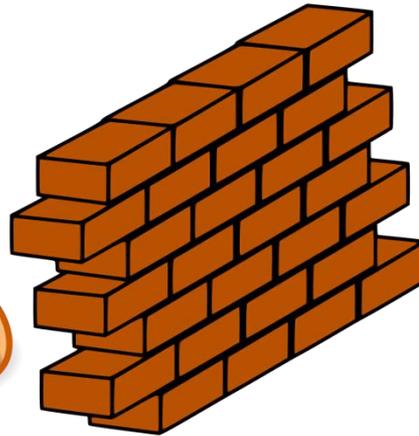
Fairness and Privacy: the Untrusted Vendor



Vendor
(may be untrusted)

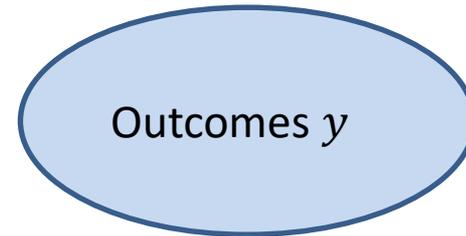


Individuals' data x_i ,
including protected
attribute(s)



Fair algorithm

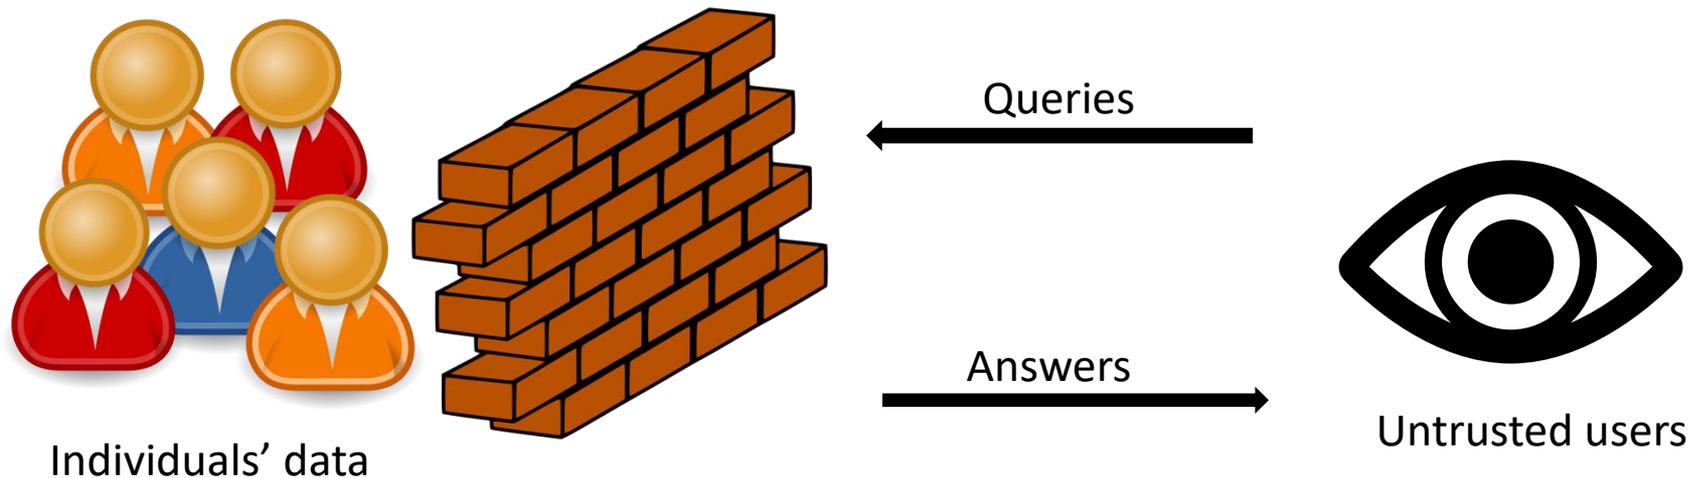
$$M(x)$$



Outcomes y

The user of the algorithm's outputs (the *vendor*) may discriminate, e.g. in retaliation for a fairness correction (Dwork et al., 2012)

Interlude: Differential Privacy (Dwork et al., 2006)

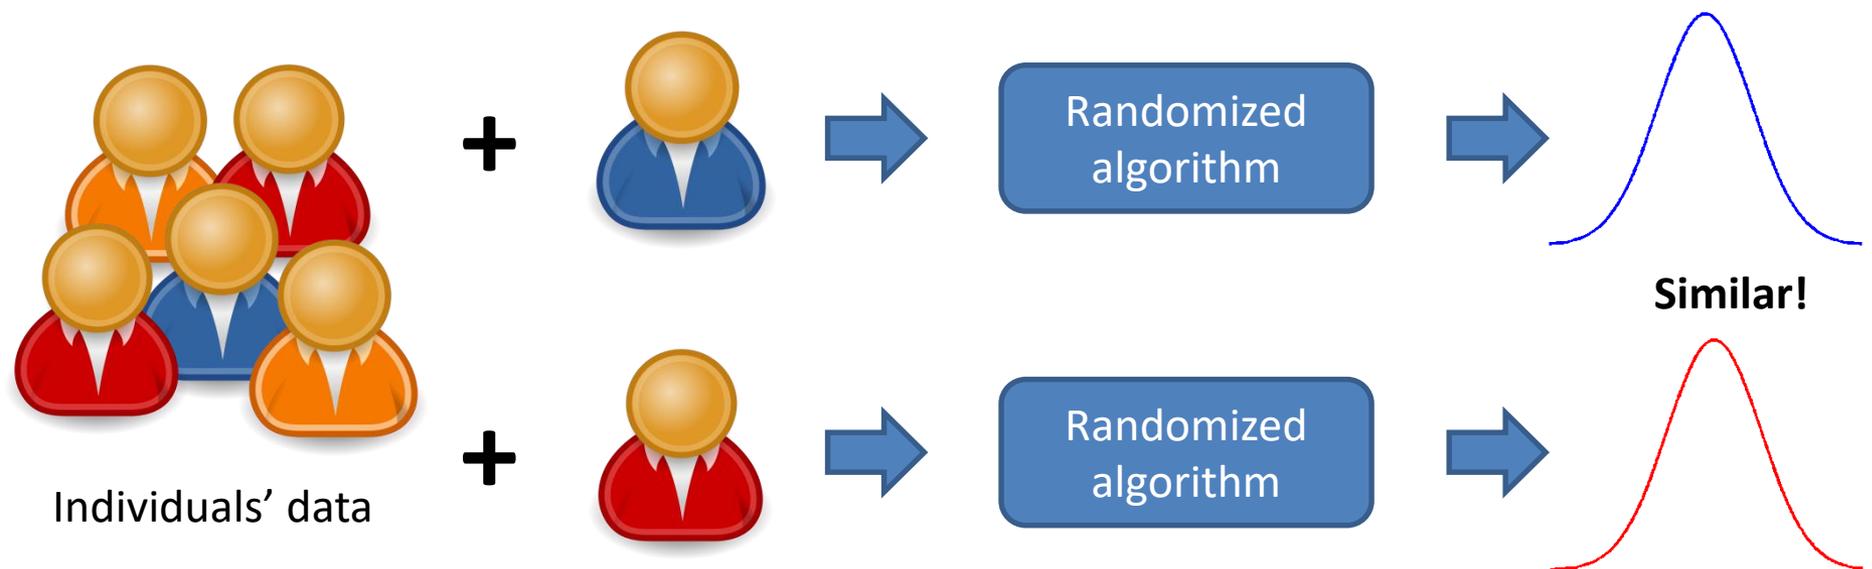


Privacy-preserving interface: randomized algorithms

- **DP is a promise:**
 - “If you add your data to the database, you will not be affected much”

Interlude: Differential Privacy (Dwork et al., 2006)

- Consider randomized algorithm $\mathcal{M}(\mathbf{X})$
- DP guarantees that the likely output of $\mathcal{M}(\mathbf{X})$ is not greatly affected by **any one data point**
- In particular, the **distribution over the outputs** of the algorithm will not change too much

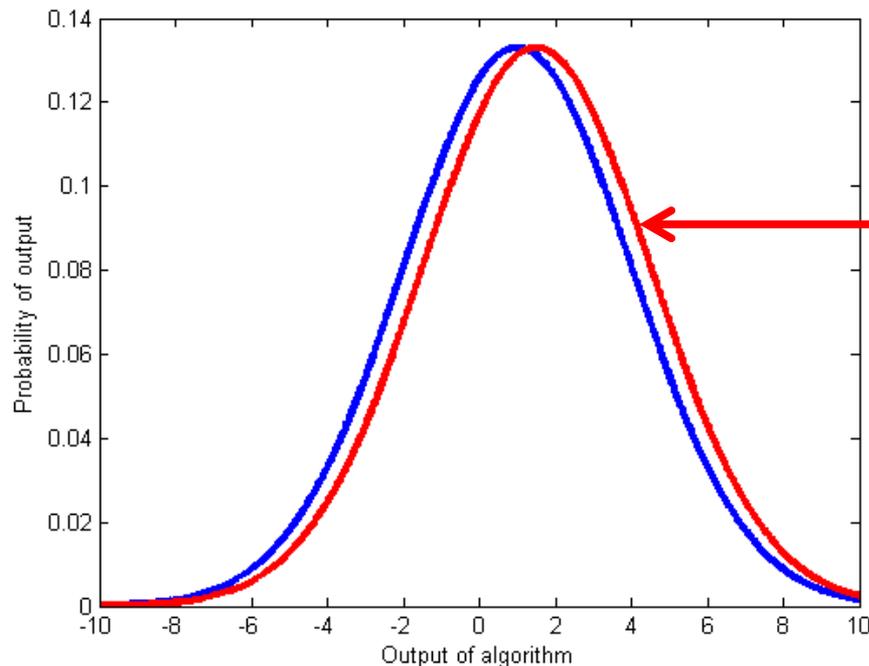


Interlude: Differential Privacy (Dwork et al., 2006)

Definition: $\mathcal{M}(\mathbf{X})$ is ϵ -differentially private if

$$\frac{\Pr(\mathcal{M}(\mathbf{X}) \in \mathcal{S})}{\Pr(\mathcal{M}(\mathbf{X}') \in \mathcal{S})} \leq e^\epsilon$$

for all outcomes \mathcal{S} , and pairs of databases \mathbf{X} , \mathbf{X}' differing in a single element.



Interlude: Differential Privacy (Dwork et al., 2006)

- Interpreting ϵ : an **economic guarantee**
 - If the user contributes their data, their **expected utility** due to the outcome of the mechanism will be **harmed by at most a factor of** $\exp(\epsilon) \approx 1 + \epsilon$ (for small ϵ)
 - Example: **randomized response** (tell the truth if coin flip = heads, otherwise answer with the flip of another coin) has $\epsilon = \ln(3) \approx 1.0986$
- In practice, differential privacy is implemented by adding noise from the Laplace distribution (two exponential distributions, back to back).
 - Noise from the Gaussian distribution preserves a slightly relaxed definition

Differential Privacy vs the 80% Rule

Definition: $\mathcal{M}(\mathbf{X})$ is ϵ -differentially private if

$$e^{-\epsilon} \leq \frac{\Pr(\mathcal{M}(\mathbf{X}) \in \mathcal{S})}{\Pr(\mathcal{M}(\mathbf{X}') \in \mathcal{S})} \leq e^{\epsilon}$$

for all outcomes \mathcal{S} , and pairs of databases \mathbf{X} , \mathbf{X}' differing in a single element.

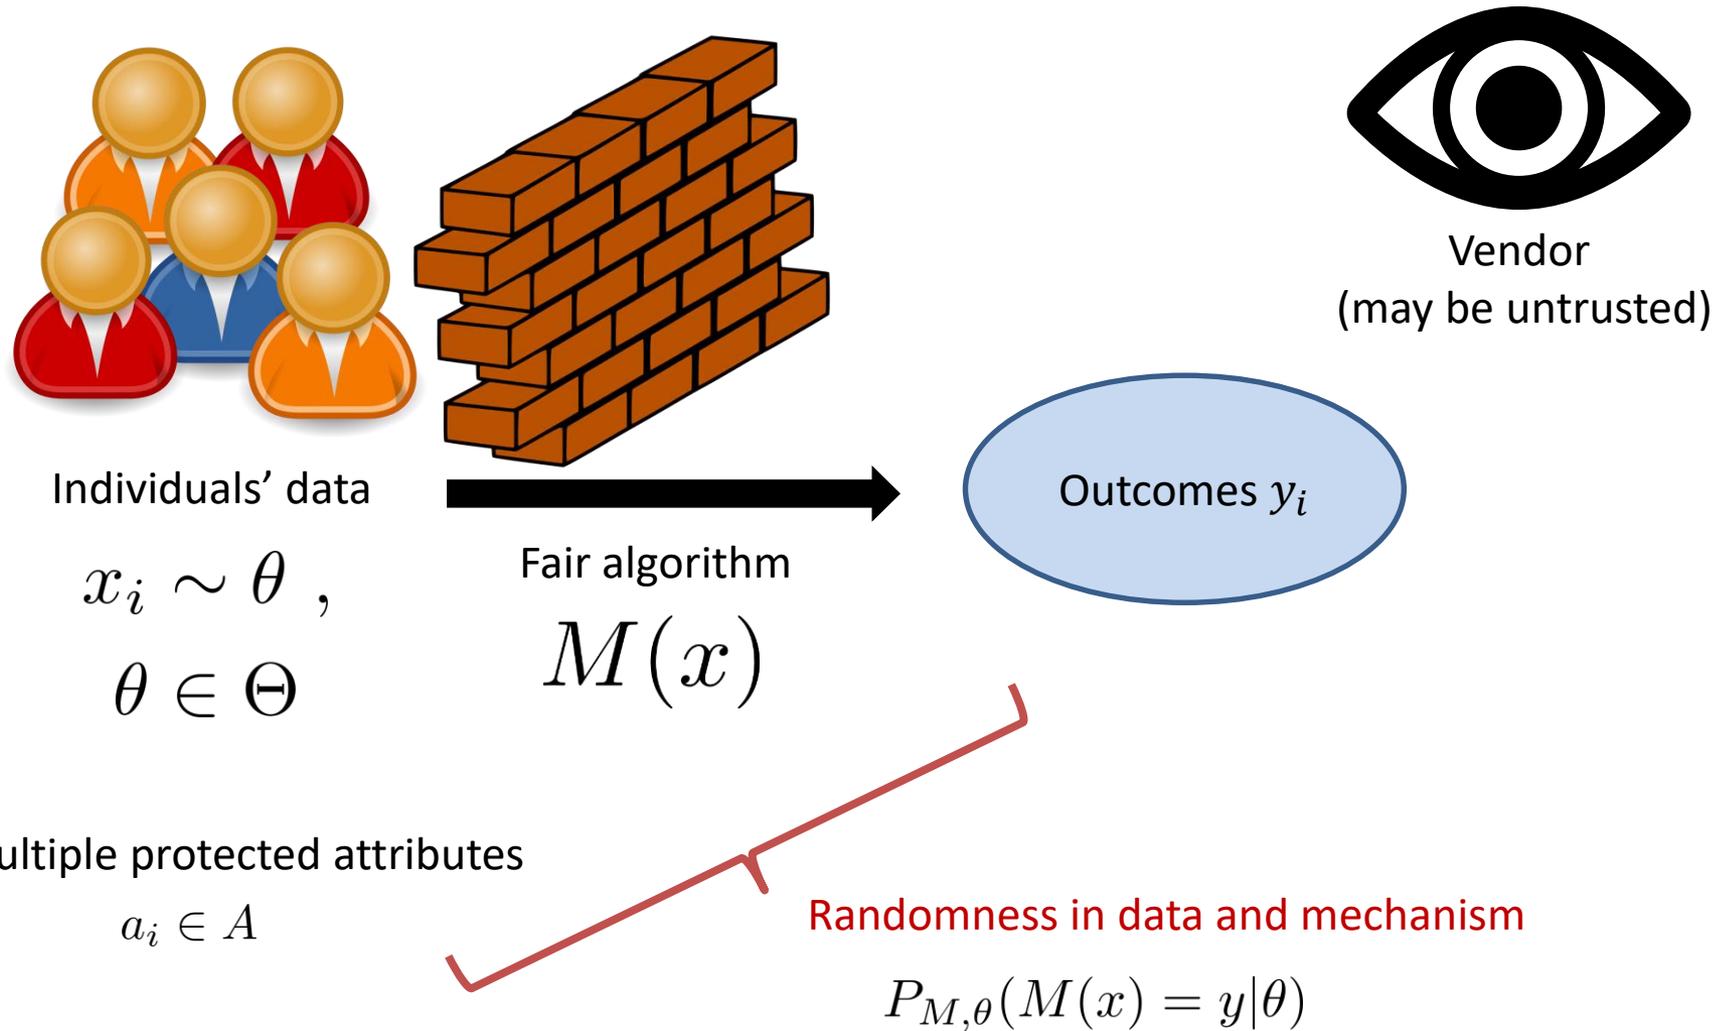
Follows from taking the reciprocal. We want ratios close to 1

- 80% rule: Evidence of unfairness if:

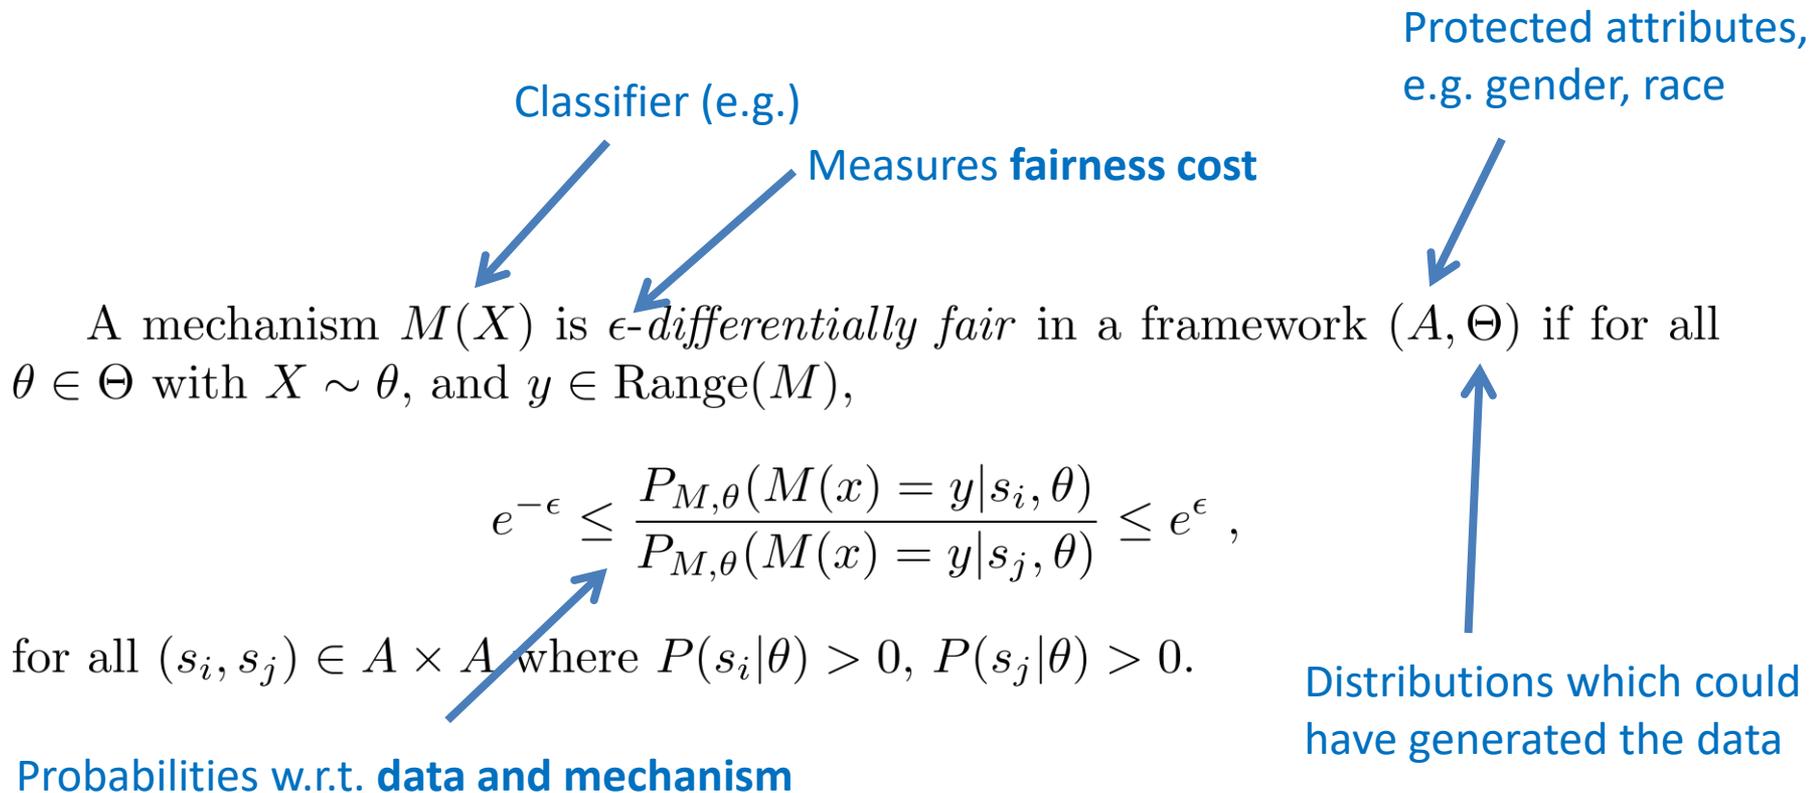
$$\frac{\Pr(\text{hire}|\text{group A})}{\Pr(\text{hire}|\text{group B})} < 0.8$$

The ratio determines the degree of disparate impact between groups. Like differential privacy, we want to bound a ratio to be somewhere near 1

Scenario for Differential Fairness

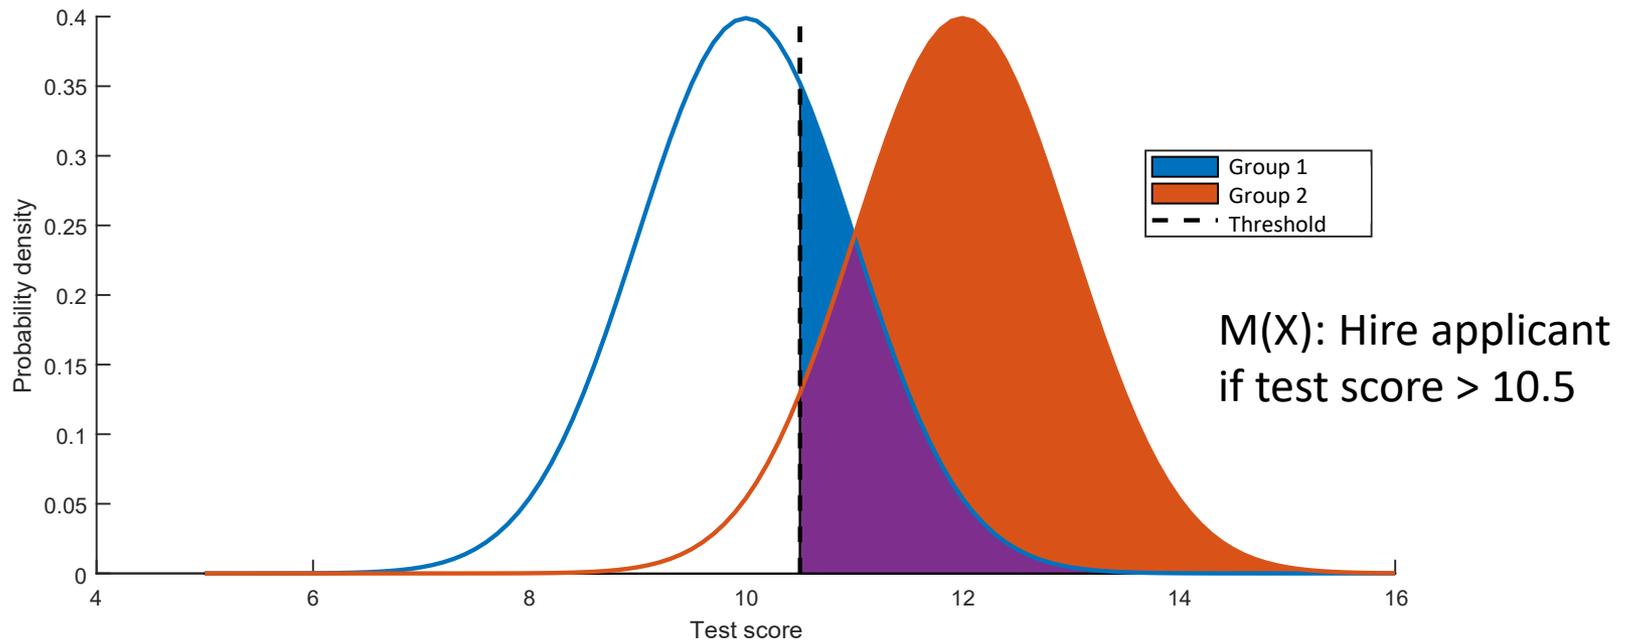


Our Proposed Fairness Definition: Differential Fairness (DF)



Key idea: ratios of probabilities of outcomes bounded for any pair of values of protected attributes

Differential Fairness Example

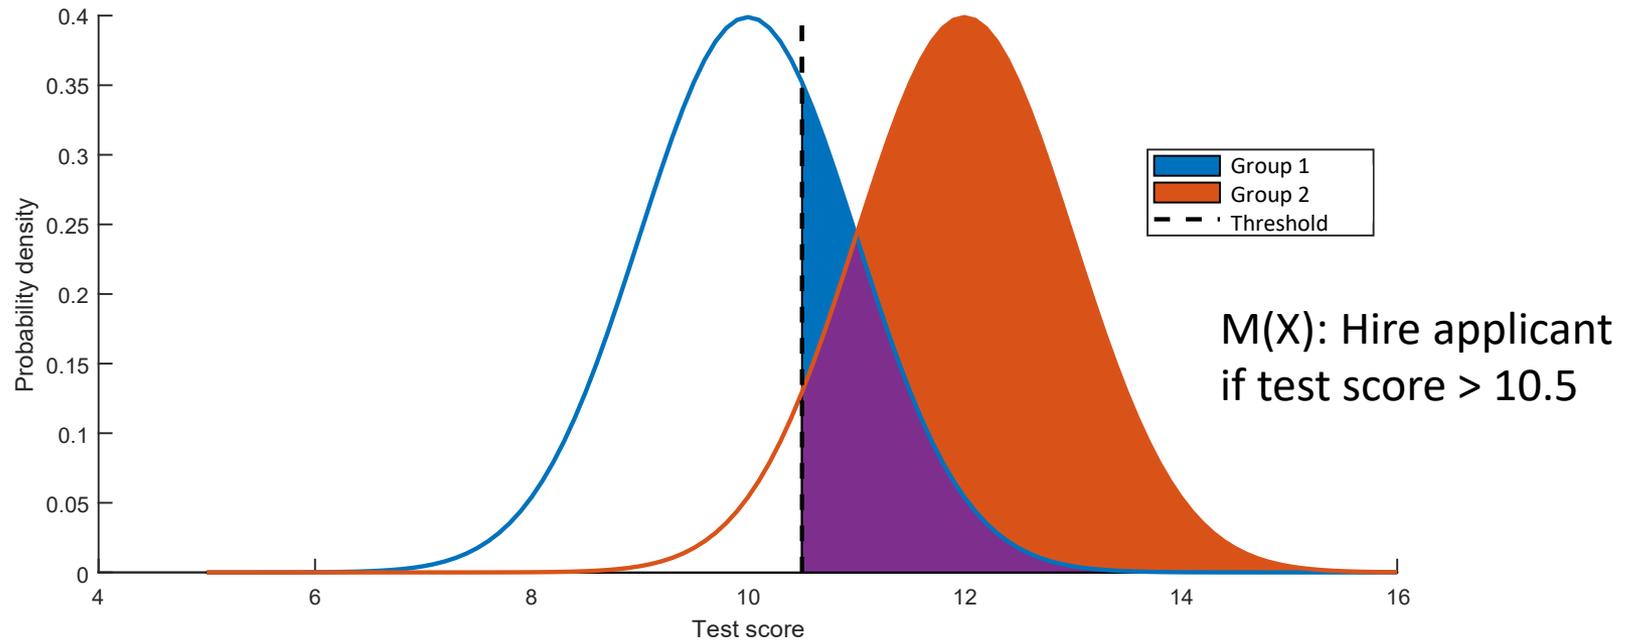


Scenario: Given an applicant's score on a standardized test, an applicant is hired if their test score is greater than a threshold t . Here, $t = 10.5$. Each group of applicant has a different distribution over scores:

$$\text{Group 1: } N(X; \mu_1 = 10, \sigma = 1)$$

$$\text{Group 2: } N(X; \mu_2 = 12, \sigma = 1)$$

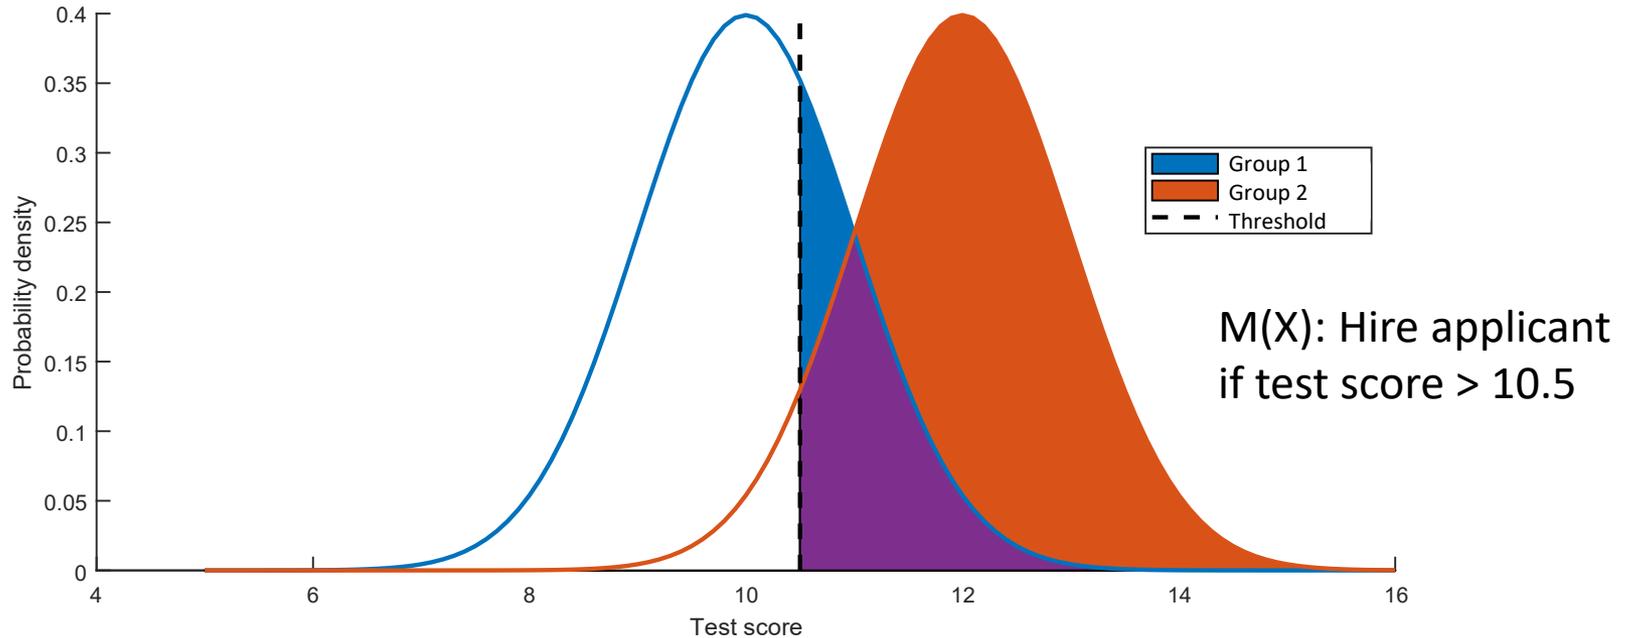
Differential Fairness Example



Probability of Hiring Outcome Given Group

		Group	
		1	2
Outcome	yes	0.3085	0.9332
	no	0.6915	0.0668

Differential Fairness Example



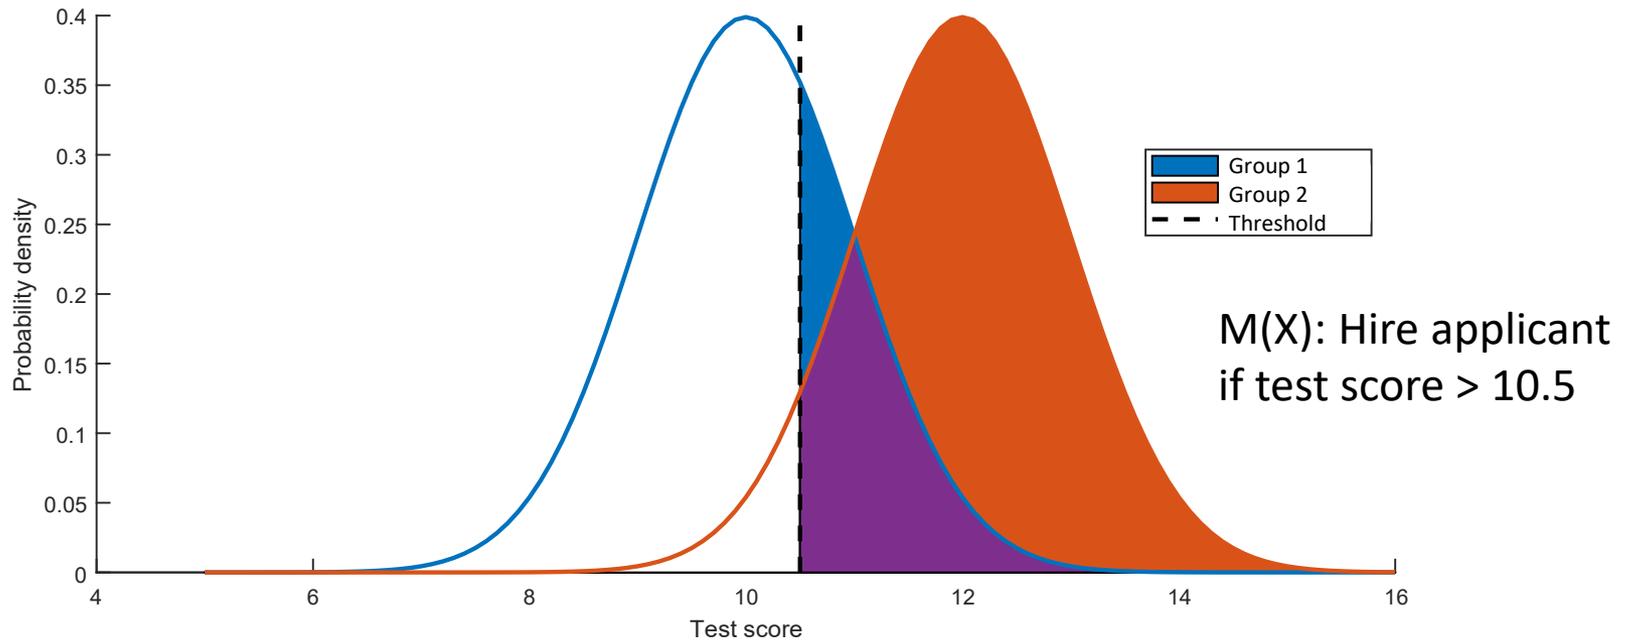
Probability of Hiring Outcome Given Group

		Group	
		1	2
Outcome	yes	0.3085	0.9332
	no	0.6915	0.0668

Log Ratios of Probabilities

y	s_i	s_j	$\log \frac{P_{M,\theta}(M(X)=y s_i,\theta)}{P_{M,\theta}(M(X)=y s_j,\theta)}$
no	1	2	2.337
	2	1	-2.337
yes	1	2	-1.107
	2	1	1.107

Differential Fairness Example



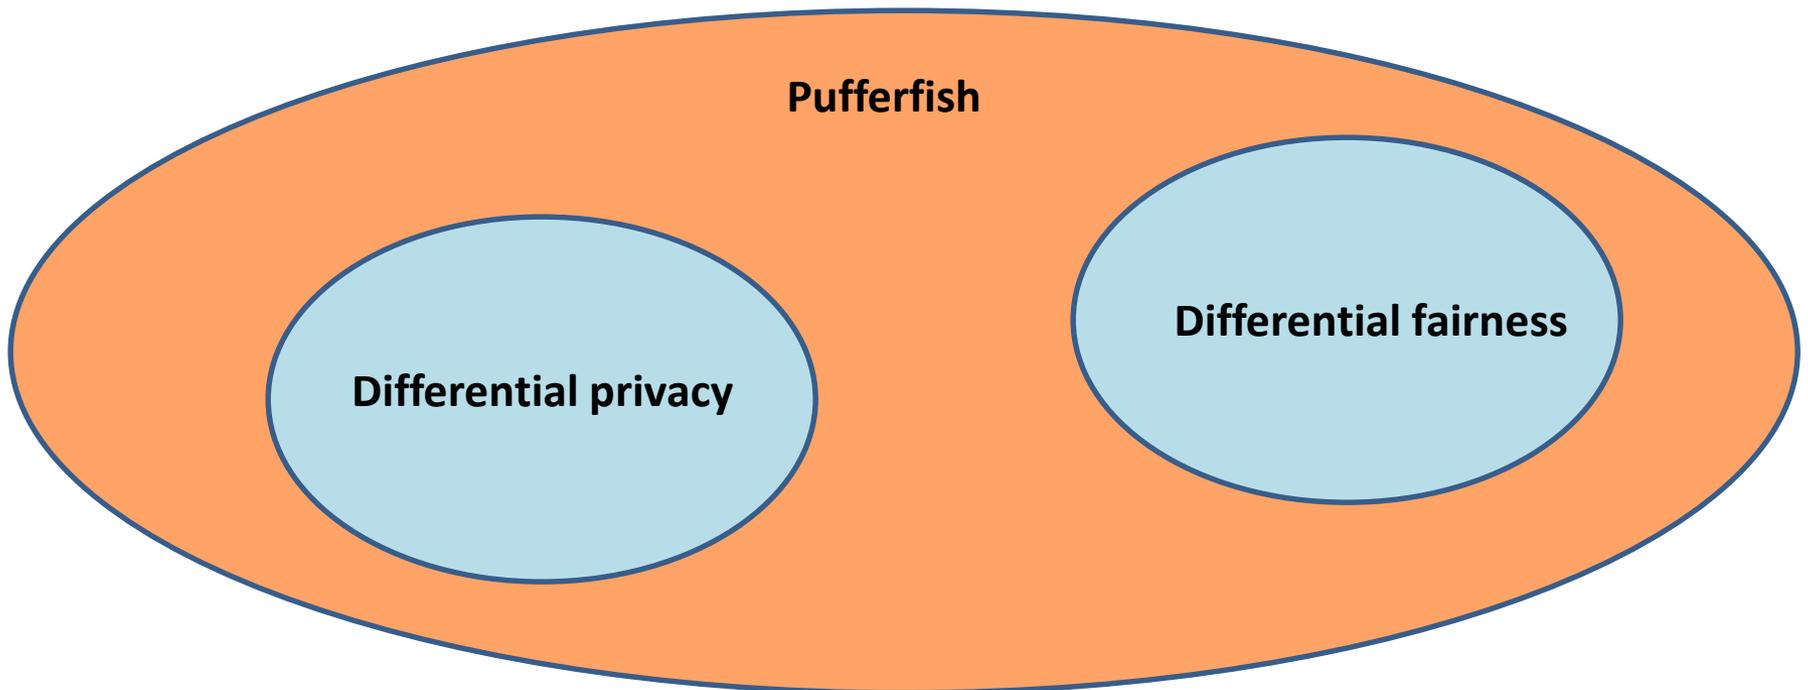
Probability of Hiring Outcome Given Group			
		Group	
		1	2
Outcome	yes	0.3085	0.9332
	no	0.6915	0.0668

Log Ratios of Probabilities			
y	s_i	s_j	$\log \frac{P_{M,\theta}(M(X)=y s_i,\theta)}{P_{M,\theta}(M(X)=y s_j,\theta)}$
no	1	2	2.337
	2	1	-2.337
yes	1	2	-1.107
	2	1	1.107

Find the worst case: $\epsilon = 2.337$

Differential Fairness and Differential Privacy

- **Differential fairness** and **differential privacy** are special cases of a more general framework, **Pufferfish** (Kifer and Machanavajjhala, 2014)



Interpreting ϵ : Bayesian Privacy

- Untrusted vendor/adversary can learn very little about the protected attributes of the instance, relative to their prior beliefs, assuming their prior beliefs are in Θ :

$$e^{-\epsilon} \frac{P(s_i|\theta)}{P(s_j|\theta)} \leq \frac{P(s_i|M(x) = y, \theta)}{P(s_j|M(x) = y, \theta)} \leq e^{\epsilon} \frac{P(s_i|\theta)}{P(s_j|\theta)}$$

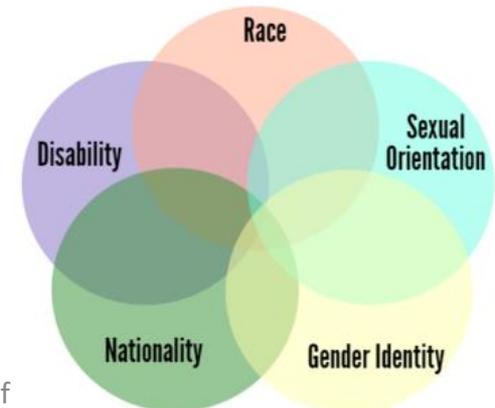
- E.g., if a loan was given to an individual, the vendor or adversary's Bayesian posterior beliefs about their race and gender will not be substantially changed
- This can **prevent subsequent discrimination**, e.g. in retaliation for a correction against bias.

Intersectionality Property of DF: Fairness with Multiple Protected Attributes

- **Intersectionality theory:** gender is not the only dimension upon which power structures in society impose systems of oppression and marginalization.
 - **The intersection of a number of aspects must be considered,** including race, sexual orientation, class, and disability status

Theorem: Let M be an ϵ -differentially fair mechanism in (A, Θ) , $A = S_1 \times S_2 \times \dots \times S_p$, and let $D = S_a \times \dots \times S_k$ be the Cartesian product of a nonempty proper subset of the variables included in A . Then M is 2ϵ -differentially fair in (D, Θ) .

E.g., if M is differentially fair in (race, gender, nationality), it is differentially fair to a similar degree in gender alone (within a factor of 2)



Measuring Bias in Data

- **Can measure bias in a dataset**

Special case of differential fairness, in which the **algorithm is the data distribution**

Empirical differential fairness (EDF) of a labeled dataset:

Corresponds to verifying that for any y, s_i, s_j , we have

$$e^{-\epsilon} \leq \frac{N_{y,s_i}}{N_{s_i}} \frac{N_{s_j}}{N_{y,s_j}} \leq e^{\epsilon}$$

- Also applies to a probabilistic model of the data



Measuring Bias Amplification

- We can measure the extent to which an algorithm **increases the bias over the original data**
- Calculate differential fairness of data, ϵ_1
- Calculate differential fairness of algorithm, ϵ_2
- Bias amplification: $\epsilon_2 - \epsilon_1$

This is a more politically conservative fairness definition: implements infra-marginality

Uncertainty in Measuring Intersectional Fairness

- Intersectionality suggests all intersections of protected groups are important for fairness
- However, more protected attributes means **less data at their intersections**

Protected attributes	gender	gender, nationality	gender, nationality, race
Median # instances	14,719	5,195	172
Minimum # instances	9,216	963	5

- With little data, estimated frequencies are unreliable.

Proposed solution

- **Predict behavior of algorithm** on each intersection using **probabilistic models** of outcome given protected attributes,

$$\Pr(y|s, \theta)$$

- Any **probabilistic classifier** can be used.
 - Naïve Bayes, logistic regression, deep neural networks...
 - We propose a hierarchical extension to logistic regression
 - Gaussian “noise” around logistic regression’s prediction allows deviations from this, if given enough data to justify it
- We recommend **Bayesian** models, to account for uncertainty

Overall approach

- Bayesian estimation of differential fairness

Input: Dev. set \mathcal{D} , mechanism $M(x)$, protected attrs A

Output: $\hat{\epsilon}_{data}$, $\hat{\epsilon}_{M(x)}$, posterior uncertainty boxplots

Apply $M(x)$ to $x_i \in \mathcal{D}$ to obtain mechanism labels y'_i ;

Fit Bayesian classifier $p_1(y|s, \bar{\theta}_1)$ on $\mathcal{D}_s = \{(s_i, y_i)\}$;

Fit Bayesian classifier $p_2(y'|s, \bar{\theta}_2)$ on $\mathcal{D}'_s = \{(s_i, y'_i)\}$;

Estimate $\hat{\epsilon}_{data}$ via DF, posterior predictive $p_1(y|s)$;

Estimate $\hat{\epsilon}_{M(x)}$ via DF, posterior predictive $p_2(y'|s)$;

Plot posterior uncertainty in ϵ_{data} , $\epsilon_{M(x)}$, $\epsilon_{M(x)} - \epsilon_{data}$;

Experiments: US Census Data

- Used the Adult dataset from the UCI repository
- Binary classification problem: does an individual earn \geq \$50,000 per year?
 - Can be a proxy for e.g., whether to approve housing application
 - 14 attributes on work, relationships, demographics
 - Training set: 32,561 instances, Test set: 16,281 instances
- We select **protected attributes** = *race, gender, nationality*

Experiments: US Census Data

- Predictive accuracy of $\Pr(y|s, \theta)$ models on test set

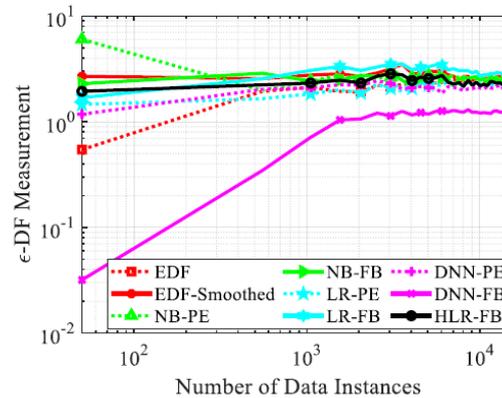
Models	Actual-labeled test set (full training set)		$M(x)$ -reabeled test set (held-out training subset)		Actual-labeled test set (20% of training set)		$M(x)$ -reabeled test set (20% of the training subset)	
	PE	FB	PE	FB	PE	FB	PE	FB
EDF	-0.437	-0.436	-0.358	-0.358	-0.435	-0.441	-0.436	-0.368
NB	-0.433	-0.433	-0.365	-0.351	-0.433	-0.431	-0.362	-0.351
LR	-0.442	-0.43	-0.382	-0.35	-0.443	-0.43	-0.377	-0.349
DNN	-0.436	-0.434	-0.366	-0.361	-0.441	-0.44	-0.366	-0.363
HLR	X	-0.432	X	-0.353	X	-0.433	X	-0.35

Comparison of predictive performance of intersectional fairness models with respect to average negative cross-entropy per intersection on the test set (higher is better), Adult dataset. Here, PE = point estimate, FB = fully Bayesian estimate.

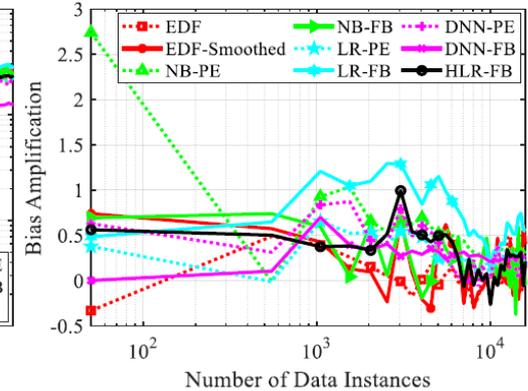
- Probabilistic models beat empirical frequencies
- Bayesian models beat point estimates
- Our proposed Bayesian hierarchical logistic regression model works best
- These differences are magnified in the small-data regime

Measuring the “variance of bias”

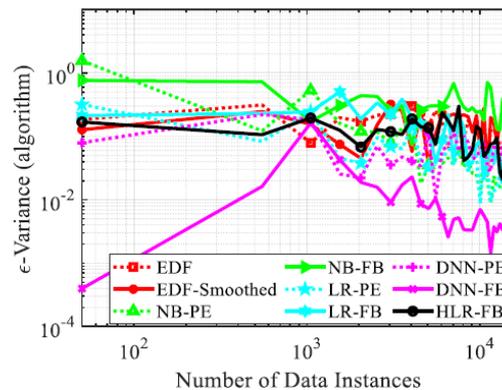
- Vary amount of data via bootstrap sampling
- Measure ϵ , bias amplification, their variances



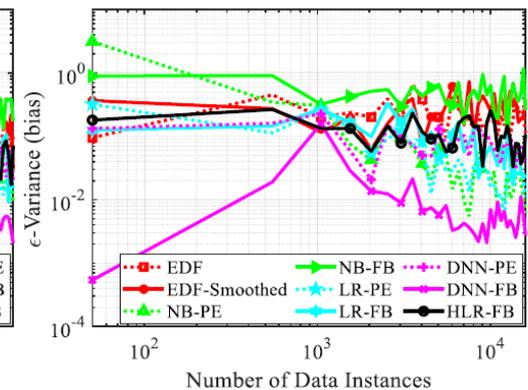
(a)



(b)



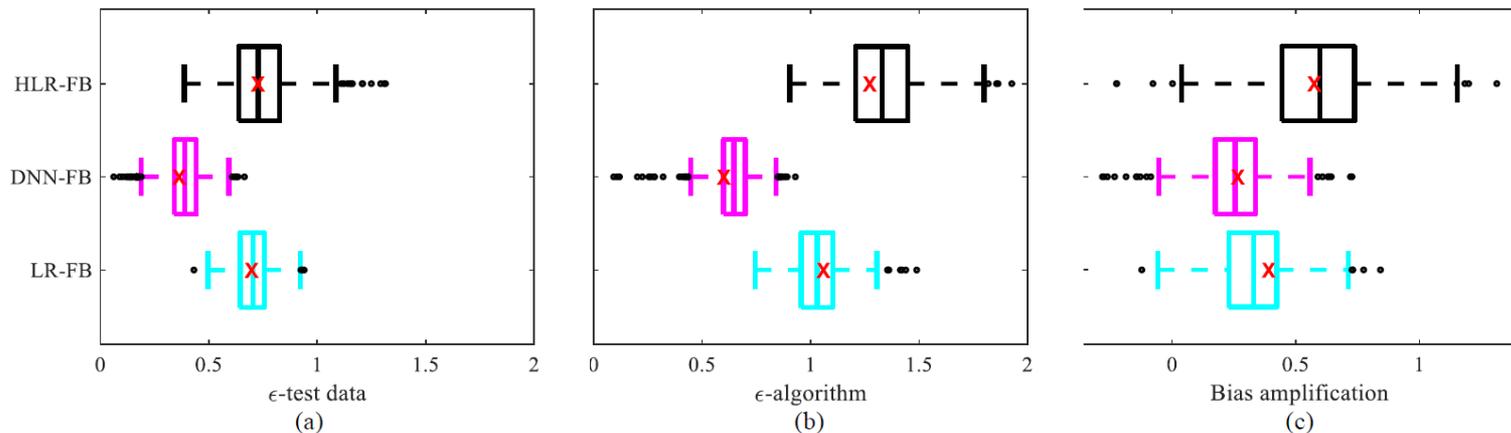
(c)



(d)

Case Study: COMPAS dataset

- Measured differential fairness, bias amplification of COMPAS recidivism predictor



- 80% rule requires $\epsilon < -\log(0.8) = 0.2231$
- All models predict that the bias exceeds this

Conclusion

“The rise of big-data optimism is here, and if ever there were a time when politicians, industry leaders, and academics were enamored with artificial intelligence as a superior approach to sense-making, it is now.

This should be a wake-up call for people living in the margins, and people aligned with them, to engage in thinking through the interventions we need.”

-Safiya Umoja Noble. *Algorithms of Oppression: How Search Engines Reinforce Racism.* New York University Press, 2018

References

- **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. ProPublica, May, 23, 2016.
- Barocas, S., & Selbst, A. D. Big data's disparate impact. Cal. L. Rev., 104, pp. 671-732, 2016.
- Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In Advances in Neural Information Processing Systems, 2016.
- **K. Crenshaw**. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. U. Chi. Legal F., pages 139–167, 1989.
- **C. Dwork**, M. Hardt, T. Pitassi, O. Reingold, & R. Zemel. Fairness through awareness. In Proceedings of the 3rd innovations in theoretical computer science conference (pp. 214-226). ACM, 2012.
- **C. Dwork**, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In Theory of Cryptography, pages 265–284. Springer, 2006.
- Executive Office of the President. Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights. White House, Executive Office of the President, 2016.
- M. Hardt, E. Price, N. Srebro, et al. Equality of opportunity in supervised learning. In Advances in NIPS, pages 3315–3323, 2016.
- D. Kifer and A. Machanavajjhala. Pufferfish: A framework for mathematical privacy definitions. ACM Trans. on Database Systems, 39(1):3, 2014.
- M. J. Kusner, J. Loftus, C. Russell, & R. Silva. Counterfactual fairness. In Advances in Neural Information Processing Systems (pp. 4069-4079), 2017.
- **S. Lowry** & G. Macpherson, A Blot on the Profession, BRIT. MED. J. 296, pp. 657-658 (1988).
- **S. U. Noble**. Algorithms of Oppression: How Search Engines Reinforce Racism. NYU Press, 2018.
- **C. Simoiu**, S. Corbett-Davies, S. Goel, et al. The problem of infra-marginality in outcome tests for discrimination. The Annals of Applied Statistics, 11(3):1193–1216, 2017.
- R. Speer. How to Make a Racist AI Without Really Trying. ConceptNet blog.2017.
<http://blog.conceptnet.io/posts/2017/how-to-make-a-racist-ai-without-really-trying/>
- **C. Starmans**, M. Sheskin, and P. Bloom. Why people prefer unequal societies. Nature Human Behaviour, 1(4):0082, 2017.
- **Zhao, J.**, Wang, T., Yatskar, M., Ordonez, V., & Chang, K. W. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. EMNLP, 2017.

(9/16 female first authors, indicated in bold)

Thank you!

- Contact information:
 - James Foulds
Assistant professor
Department of Information Systems
UMBC
Email: jfoulds@umbc.edu
Webpage: <http://jfoulds.informationssystemsystems.umbc.edu>
- A pre-print of our work is online at arxiv.org:
*J. R. Foulds and S. Pan. **An Intersectional Definition of Fairness.** ArXiv preprint arXiv:1807.08362 [CS.LG], 2018.*

(Pre-print on Bayesian fairness modeling coming in a few days)

Hierarchical Extension to Logistic Regression

- Assumed generative process:
 - $\sigma_2 \sim \text{Exponential}(\lambda)$
 - $\beta_i \sim \text{Normal}(\mu, \sigma_1), \quad c \sim \text{Normal}(\mu, \sigma_1)$
 - $\gamma_j \sim \text{Normal}(\beta^\top \vec{s}_j + c, \sigma_2), \quad P(y = 1 | s_j) = \sigma(\gamma_j)$



Gaussian deviation from prediction of logistic regression, in logit domain

Fairness Definitions: Pros and Cons

Definition	Pros	Cons
Fairness through unawareness	Simple	Defeated by proxy variables
Demographic parity <i>-outcome distributions to be equal for each protected category</i>	Appealing for civil rights	Does not consider infra-marginality. May harm accuracy. Can be abused by subset targeting
Equalized odds/Equality of opportunity <i>-parity for error rates</i>	Rewards accurate classification	Incompatible with calibrated probabilities. Weak on civil rights
Individual fairness <i>-similar individuals get similar outcomes</i>	Privacy guarantees, protects vs subset targeting	Must define “fair” distance measure. No generalization
Counterfactual fairness <i>-parity of outcomes under a causal model</i>	Addresses infra-marginality	Requires accurate causal model, inference. Cannot use descendants of A
Differential fairness <i>-our definition, addresses privacy and intersectionality</i>	Measurement. Privacy guarantees. Civil rights. Intersectionality. Lightweight	Similar to demographic parity, but can mitigate subset targeting.

Proof of Intersectionality Theorem

Proof. Let M be an ϵ -differentially fair mechanism in (C, Θ) , $C = S_1 \times S_2 \times \dots \times S_p$, and let $D = S_a \times \dots \times S_k$ be the Cartesian product of a nonempty proper subset of the variables included in C . We want to show that M is 2ϵ -differentially fair in (D, Θ) .

Define $E = S_1 \times \dots \times S_{a-1} \times S_{a+1} \dots \times S_{k-1} \times S_{k+1} \times \dots \times S_p$ as the Cartesian product of the variables included in C but not in D . Let $\theta \in \Theta$. Then for any outcome w , and $(s_i, s_j) \in D \times D$ where $P(s_i|\theta) > 0$ and $P(s_j|\theta) > 0$, let q be any value of E where $P(q, s_i|\theta) > 0$. (Such a q must exist, since $P(s_i|\theta) = \int_E P(E = e, s_i|\theta)de > 0$.) We have

$$\begin{aligned}
 \frac{P_{M,\theta}(M(X) = w|D = s_i, \theta)}{P_{M,\theta}(M(X) = w|D = s_j, \theta)} &= \frac{\int_E P_{M,\theta}(M(X) = w, E = e|s_i, \theta)de}{\int_E P_{M,\theta}(M(X) = w, E = e|s_j, \theta)de} \\
 &= \frac{\int_E P_{M,\theta}(M(X) = w|E = e, s_i, \theta)P_{M,\theta}(E = e|s_i, \theta)de}{\int_E P_{M,\theta}(M(X) = w|E = e, s_j, \theta)P_{M,\theta}(E = e|s_j, \theta)de} \\
 &= \frac{\int_E P_{M,\theta}(M(X) = w|E = e, s_i, \theta)P_\theta(E = e|s_i, \theta)de}{\int_E P_{M,\theta}(M(X) = w|E = e, s_j, \theta)P_\theta(E = e|s_j, \theta)de} \\
 &\quad \times \frac{P_{M,\theta}(M(X) = w|q, s_i, \theta)}{P_{M,\theta}(M(X) = w|q, s_i, \theta)} \\
 &= \frac{\int_E \frac{P_{M,\theta}(M(X)=w|E=e,s_i,\theta)}{P_{M,\theta}(M(X)=w|q,s_i,\theta)} P_\theta(E = e|s_i, \theta)de}{\int_E \frac{P_{M,\theta}(M(X)=w|E=e,s_j,\theta)}{P_{M,\theta}(M(X)=w|q,s_i,\theta)} P_\theta(E = e|s_j, \theta)de} \\
 &\leq \frac{\int_E e^\epsilon P_\theta(E = e|s_i, \theta)de}{\int_E e^{-\epsilon} P_\theta(E = e|s_j, \theta)de} && \text{(note: must be careful to verify case} \\
 &= e^{2\epsilon} . && \text{where probabilities are 0 here)}
 \end{aligned}$$

Reversing s_i and s_j and taking the reciprocal shows the other inequality. \square