# Impossibilities in Computing Lecture 4 

Unit 2
ML and Society (Spring 2024)

Atri Rudra

## Pass phrase: Alexandra Chouldechova



# Alexandra Chouldechova 

Estella Loomis McCandless Assistant Professor of Statistics and Public Policy Heinz College, Carnegie Mellon University
Office: Hamburg Hall 2224
Email: achould(at)cmu.edu
Phone: 412-268-4414
CV Google Scholar

## Brief Academic Bio

Dr. Alexandra Chouldechova is the Estella Loomis McCandless Assistant Professor of Statistics and Public Policy at Carnegie Mellon University's Heinz College of Information Systems and Public Policy. Her research investigates questions of algorithmic fairness and accountability in datadriven decision-making systems, with a domain focus on criminal justice and human services. Her work has been supported through funding from organizations including the Hillman Foundation, the MacArthur Foundation, and the NSF Program on Fairness in Artificial Intelligence in Collaboration with Amazon. She is a member of the executive committee for the ACM Conference on Fairness, Accountability and Transparency (FAccT), and previously served as a Program Committee co-Chair for the conference.

Dr. Chouldechova is a 2020 Research Fellow with the Partnership on Al, where she is working on understanding factors that drive racial bias in algorithmic risk assessment tools being developed for use in pre-trial, parole and sentencing contexts. She is also a member of the Pittsburgh Task Force on Public Algorithms.

Dr. Chouldechova received her PhD in Statistics from Stanford University and an H.B.Sc. in Mathematical Statistics from the University of Toronto.

## Peer Survey score out by this afternoon



## [Combined] Peer Survey scores for Unit 1

Sorry for the delay but the "Average Normalized Score" based on peer survey for unit 1 submission is done. Please see the IP page for more details on how this score (which is the number $\rho$ mentioned on the IP page) will be used to compute the individual score of unit 1 for you.

For ML\&Soc students: this score will be released on Autolab by 3pm today. We'll post a comment on this post when this is done.

For Rage students: Dr. Muller will be communicating the scores to y'all.

All: Couple of followup comments:

- As a heads up, the scores will be released to Rage and ML\&Soc students on a slightly different time table-- thanks for your patience!
- If you are curious about how the peer survey ratings got converted into $\$ \backslash r h o \$$, please see below.

The "Average Normalized Score" is based on your survey scores. (If you have been in 331 or 442 with Matt Hertz then the stuff below should be very familiar to y'all.)

- Each student's "teamwork factor" is the percentage of the team's outcome attributable to that student (as judged by each member who completed the survey) and a student's "Average Normalized Score" is their average "teamwork factor" from their team-mates (including themselves)

Each "teamwork factor" is a number between 0 and the number of members of their team. A 0 means the team agrees that student did nothing; the max score means the team agrees that student did everything. This process begins by calculating how each student awarded the points. For each student's submission, we sum the marks they gave a teammate, divide that sum by the total marks they gave out, and then multiply it by the number of students on the team (to normalize scores).

- For example, consider the scores submitted by student A:
- A got 15 marks, B also got 15 marks, and C got 4 marks. So the total marks is 34 and the "teammate factor" scores are:
- $\mathrm{A}=15 / 34=0.44111$ * 3 (to normalize) $=1.32$
- There is a bit of rounding error in this explanation but not in how l'll use it to calculate your final grade
- $B=15 / 34=0.44111$ * $3=1.32$
- $C=4 / 34=0.1176$ * $3=0.35$


## Remember COMPAS?

## COMPAS (software)

From Wikipedia, the free encyclopedia
COMPAS, an acronym for Correctional Offender Management Profiling for Alternative Sanctions, is a case manag Equivant『) used by U.S. courts to assess the likelihood of a defendant becoming a recidivist. ${ }^{[1][2]}$

COMPAS has been used by the U.S. states of New York, Wisconsin, California, Florida's Broward County, and oth

## Contents [hide]

1 Risk Assessment
2 Critiques and legal rulings
3 Accuracy
4 Further reading
5 See also
6 References

## Risk Assessment [edit]



## Broward County

County in Florida

Broward County is a county in southeastern Florida, US. According to a 2018 census report, the county had a population of $1,951,260$, making it the second-most populous county in the state of Florida and the 17th-most populous county in the United States. The county seat is Fort Lauderdale. Wikipedia

## Incorporated cities: 24

Population: 1.936 million (2017)
Mayor: Mark D. Bogen

## 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

## A sample of their result

Black Defendants' Risk Scores


White Defendants' Risk Scores


## False Positives, False Negatives, and False Analyses: A Rejoinder to "Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And It's Biased Against Blacks."

## What is fair? Take-1



## What is fair? Take-2




## Very limited coverage

## Book (in progress) on fairness and ML

If you are interested in more details, we refer you to the book by Barocas, Hardt and Narayanan on this topic. As of the time of writing of these notes, this book is still not complete but it does contain a lot of the known material that we will not cover in this course.

Will we get the definition of fairness?


## What about bias?

## What is bias?

Another loaded term that we will use is the term bias. In particular, there are roughly three kinds of notions of bias that is relevant to these notes:

1. The first notion (which might be the least known) occurs in a dataset where there are certain specific collection of input variable values occur more than others. This essentially measure how far away from a truly random $\checkmark$ dataset the given dataset is. Note that this notion is bias is necessary for ML to work. If all the datapoints are completely random (i.e. both their input and target variable values are completely random), then there is no bias for a classifier to "exploit"-- in other words, one might as well just output a random label for prediction.
2. The second notion of bias is that of statistical bias $\mathcal{J}$, where in our setting this would mean that the binary classifier outcome does not reflect the distribution of the target variable. Such a classifier would be well calibrated ${ }^{〔}$, if this does not happen. One could consider a well-calibrated binary classifier to be fair in some sense. This will be one notion of fairness that will come up in the COMPAS story. (This is the notion of fairness used in the rejoinder to the ProPublica article).
3. The finally notion of bias is the colloquial use of the term $\rfloor$ that is mean to denote an outcome that is not fair. Most of the definitions of fairness in the literature deal with this notion of bias. And a couple of definition of this kind of fairness will also play a part in the COMPAS story (this is the notion of fairness used in the ProPublica article).


## Going beyond \# correctly classified points



## Binary classifier output



## "Confusion matrix"

## True Positive rate



False Negative Rate (FNR)


False Positive Rate (FPR)


## True Negative Rate (TNR)



## Exercise: which of these are true?

$T P R+F N R=1 ?$

$F P R+F N R=1 ?$


## $\mathrm{TPR}+\mathrm{FNR}=1$

$\mathrm{FPR}+\mathrm{TNR}=1$


## Positive Predictive Value (PPV)




## Back to fairness

## Protected/Sensitive attribute

To define group fairness, we have to well, define a group first. Towards this, we will use the notion of a protected attribute or sensitive attribute (we will use both terminology interchangeably): this will be a special attribute $R$ (which takes few pre-defined values i.e. is a categorical variable ['])-- and each choice of the value of $R$ defines a separate group. There is precedence in US law: grouping this way is used in the concept of protected class [ $\mathcal{B}$ in US anti-discrimination law-- i.e. one cannot discriminate on the basis of any protected class.
Coming back to the COMPAS example, we will use $R$ to denote the race and for simplicity we will assume the two values $R$ can take are $b$ (for black) and $w$ (for white). While clearly these are not the only racial classification, the results of ProPublica mentioned earlier focus on these two value of race and hence we concentrate on these two possibilities.
For the rest of the section, we will only consider groups corresponding to $R(x)=b$ and $R(x)=w$ (i.e. groups based on whether race of $x$ is black or white).

## Statistical parity

At a high level we would like the accuracy of binary classifier to be the same across groups. Since in real life false positive positives and false negatives have different costs, various instantiation of statistical parity definitions follows by asking that different notions of accuracy be the same across groups.

## Why statistical parity across groups?

## 

## 29 CFR § 1607.4 - Information on impact.

D. Adverse impact and the "four-fifths 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. Smaller differences in selection rate may nevertheless constitute

## Notes on ML and law

## Discrimination, Law and ML

This page will do a quick overview of anti-discrimination law and how it could/would interact with the ML pipeline.

A Under Construction
This page is still under construction. In particular, nothing here is final while this sign still remains here.

A Request
I know I am biased in favor of references that appear in the computer science literature. If you think I am missing a relevant reference (outside or even within CS), please email it to me.

## Anti-discrimination law

Rates for groups

$$
\begin{aligned}
& \text { W }
\end{aligned}
$$

FPR and FNR for groups


## PPV for groups




## Finally, the formal fairness definitions

## Equal FPR

We say a classifier fair with respect to FPR if

$$
F P R_{b}=F P R_{w}
$$

In the COMPAS context, a classifier is fair with respect to FPR if chances of a black and white defendants begin identified as reoffending when they actually did not end up reoffending are the same. This is one of the notions of fairness that ProPublica used.

## Equal FNR

We say a classifier fair with respect to FNR if

$$
F N R_{b}=F N R_{w}
$$

In the COMPAS context, a classifier is fair with respect to FNR if chances of a black and white defendants begin identified as not reoffending when they actually did end up reoffending are the same. This is one of the notions of fairness that ProPublica used.

## Why no Equal TPR and TNR?

You might be wondering why we did not define notions of fairness with respect to TPR and TNR? The answer is that because we already have! Do you see why?
Hint: We claim those follow from fairness with respect to FNR and FPR respectively.
In the COMPAS context, a classifier is fair (or does not have any statistical bias [ 6 ) if the chances of a black and white defendant being correctly identified as reoffending given that the classifier identified them as such are the same. This is the notion of fairness used in the rejoinder to the ProPublica article.

## Exercise!

## Exercise

For each notion of being fair with respect to FPR, FNR and well-calibrated, decide if it holds for the following instance (that we have seen before):



Break!

## Fairness definitions

## Equal FPR

We say a classifier fair with respect to FPR if

$$
F P R_{b}=F P R_{w}
$$

In the COMPAS context, a classifier is fair with respect to FPR if chances of a black and white defendants begin identified as reoffending when they actually did not end up reoffending are the same. This is one of the notions of fairness that ProPublica used.

## Equal FNR

We say a classifier fair with respect to FNR if

$$
F N R_{b}=F N R_{w}
$$

In the COMPAS context, a classifier is fair with respect to FNR if chances of a black and white defendants begin identified as not reoffending when they actually did end up reoffending are the same. This is one of the notions of fairness that ProPublica used.

## Well-calibrated

We say a classifier if well-calibrated if

$$
P P V_{b}=P P V_{w}
$$

In the COMPAS context, a classifier is fair (or does not have any statistical bias ${ }^{[17}$ ) if the chances of a black and white defendant being correctly identified as reoffending given that the classifier identified them as such are the same. This is the notion of fairness used in the rejoinder to the ProPublica article.

## Connecting back to COMPAS story

## ProPublica vs. its Rejoinder

First let us recap the notions of fairness used by the ProPublica article and its rejoinder. The ProPublica article used the fairness with respect to FPR and FNR as its notion of fairness while the rejoinder used well-calibrated as its notion of fairness. Here are the values of the corresponding rates take directly from the accompanying article $[\mathcal{B}$ to the original ProPublica article ("Low" and "High" correspond to $S=-1$ and $S=1$ while "Survived" and "Recidivated" correspond to $Y=-1$ and $Y=1$ resp.):

| All Defendants |  |  | Black Defendants |  |  | White Defendants |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Low | High |  | Low | High |  | Low | High |
| Survived | 2681 | 1282 | Survived | 990 | 805 | Survived | 1139 | 349 |
| Recidivated | 1216 | 2035 | Recidivated | 532 | 1369 | Recidivated | 461 | 505 |
| FP rate: 32.35 |  |  | FP rate: 44.85 |  |  | FP rate: 23.45 |  |  |
| FN rate: 37.40 |  |  | FN rate: 27.99 |  |  | FN rate: 47.72 |  |  |
| PPV: 0.61 |  |  | PPV: 0.63 |  |  | PPV: 0.59 |  |  |
| NPV: 0.69 |  |  | NPV: 0.65 |  |  | NPV: 0.71 |  |  |
| LR+: 1.94 |  |  | LR+: 1.61 |  |  | LR+: 2.23 |  |  |
| LR-: 0.55 |  |  | LR-: 0.51 |  |  | LR-: 0.62 |  |  |

By looking at the table above, it can be seen that they both are right. In particular, the COMPAS classifier is not fair with respect to either FPR (denoted by "FP rate" in the above table) not with respect to FNR (denoted by "FN rate" in the above table). On the other hand, COMPAS classifier seems well-calibrated since the PPV values are essentially same for both groups.


## Perhaps COMPAS can be improved?

## Digression: How do you measure recidivism

This is a good time to clarify/remind you that the recidivism rates being higher for blacks than whites does not imply that blacks necessarily reoffend at a higher rate the whites. Think about why this could be the case.

Hint: How would you measure whether someone reoffended or not?

## NO, you can't!

it is impossible for a binary classifier to satisfy all three notions of fairness (i.e. fairness with respect to FPR, FNR and being well-calibrated) unless the fraction of positives to the overall number of points is the same in both groups.

In the COMPAS dataset, the recidivism rate for blacks and whites are $50 \%$ and $39 \%$ respectively. Hence, the fact that COMPAS could not satisfy all three notions of fairness, is mathematically unavoidable.

The above kind of result is also known as impossibility theorem : see e.g this impossibility theorem for voting systems $\boldsymbol{\beta}$ for a more well-known such result.

$\mathrm{FNR}_{\mathrm{b}}=\mathrm{FNR}_{\mathrm{w}}=1 / 2$

$F P R_{b}=\operatorname{FPR}_{\mathrm{w}}=1 / 2$


Overall situation
What are $\mathrm{PPV}_{\mathrm{b}}$ and $\mathrm{PPV}_{\mathrm{w}}$ ?


$$
P P V_{b}=\frac{x}{x+z} \text { and } P P V_{w}=\frac{y}{y+u}
$$

$$
p_{b}=\frac{x+x}{z+z+x+x}=\frac{x}{x+z}=P P V_{b}
$$



$$
p_{w}=\frac{y+y}{y+y+u+u}=\frac{y}{y+u}=P P V_{w} .
$$

## Argue the general case!


b


W

## Argument for the general case (left as an exercise)

Argue the impossibility theorem for the more general case of

$$
F P R_{b}=F P R_{w}=p \text { and } F N R_{b}=F N R_{w}=q
$$

where $p$ and $q$ are arbitrary numbers between 0 and 1 (i.e. they need not even be the same let above both be equal to $\frac{1}{2}$.

## Exercise: What are $\mathrm{PPV}_{\mathrm{b}}$ and $\mathrm{PPV}_{\mathrm{w}}$ ?


b


W

$$
P P V_{b}=\frac{(1-p) \cdot x}{(1-p) \cdot x+q \cdot z} \quad \text { and } \quad P P V_{w}=\frac{(1-p) \cdot y}{(1-p) \cdot y+q \cdot u}
$$

Exercise: What are $\mathrm{p}_{\mathrm{b}}$ and $\mathrm{p}_{\mathrm{w}}$ ?


$$
\mathrm{p}_{\mathrm{b}}=\frac{x}{x+z}
$$



W

$$
\mathrm{p}_{\mathrm{w}}=\frac{y}{y+u}
$$

## Exercise: When is $\mathrm{PPV}_{\mathrm{b}}=\mathrm{PPV}_{\mathrm{w}}$ ?



$$
\mathrm{p}_{\mathrm{b}}=\frac{x}{x+z} \quad P P V_{b}=\frac{(1-p) \cdot x}{(1-p) \cdot x+q \cdot z} \quad \text { and } \quad P P V_{w}=\frac{(1-p) \cdot y}{(1-p) \cdot y+q \cdot u} . \quad \mathrm{P}_{\mathrm{w}}=\frac{y}{y+u}
$$

## A useful result

Let $a, b, c$ and $d$ be numbers such that $0<a<b$ and $0<c<d$. Then the following is true

$$
\frac{a}{b}=\frac{c}{d} \text { if and only if } \frac{a}{b-a}=\frac{c}{d-c} \text { and } \frac{a}{a+b}=\frac{c}{c+d}
$$



## Alternate conclusion from the impossibility result

We will present two equivalent descriptions of the impossibility result:

1. If we want all three forms of fairness (fairness with respect to FPR, fairness with respect to FNR and well-calibrated classifier), then that is only possible when the prevalence of the target variable (i.e. the fraction of points $x$ with $Y(X)=1$ ) is the same across groups.
2. Supposed we want to fairness with respect to FPR and FNR (we will see shortly why this is a desired outcome). Then if the prevalence of the target variable is not same across groups (i.e. the data itself is biased), then the binary classifier has to "correct" for the bias and hence, cannot be well-calibrated.

We note that the first interpretation is the literal interpretation of the impossibility result while the second interpretation is the equivalent contrapositive of the imporily theorem.

From a practical point of the view what the second interpretation implies is that (given that in real life) data is (almost always) biased, if we want fair outcomes (in the sense of fairness with respect to FPR and FNR), then the classifier has to "correct" for the bias. We will briefly come to this point in a bit.

## NO, you still can't!

A 2017 ITCS paper by Kleinberg, Mullainathan and Raghavan show that even with the above two relaxations, we cannot simultaneously satisfy all three notions of (appropriately defined) approximate notions of fairness (even for non-binary classification).

## Why should we care fairness wrt FNR and FPR

Why should we care about fairness with respect to FPR and FNR?
As alluded to earlier: there are good reasons to ask for fairness with respect to FPR and FNR. In particular, one strong motivation comes from the legal principle of disparate impact $]$. The basic principle is that the outcome of a decision maker should not impact one group under a protected class (e.g. blacks) more than another group in the same protected class (e.g. whites). Thus, if one group is disproportionately impacted by a decision process then that process can be said to be legally discriminatory.

We will see shortly that under a reasonable (but very simplified) model, difference in FPR and FNR lead to disproportionate impact.


## What if we only care about one fairness def?

## How do we incorporate a fairness notion

Here is the technical problem: given that we want a binary classifier such that it has equal (or approximately) close to equal FNR (and equal FPR), can be train a model that has the best accuracy subject to these fairness constraints?

It turns out that one can express these fairness constraints as "linear constraints" and we can use existing techniques from optimization to get a model with the required fairness constraints. Agarwal et al. show this to do this for a broad class of fairness constraints (i.e. even beyond fairness with respect to FNR and FPR) in a fairly generic way.

## Shortcomings of group fairness

## Fairness through blindness

This example is take from Dwork et al.
One common notion of fairness used is that of "fairness through blindness"-- the idea here is that the classifier explicitly does not use the sensitive attribute ( $R$ in our case is race). In particular, such classifiers explicitly excludes the sensitive attribute as one of the input variables. However, in practice e.g. it turns out that zip code is very good predictor of race. For the roots of why this is true, see this video:


## Fairness gerrymandering

## Fairness gerrymandering

The general idea of this example is from Dwork et al. though the specific version (and indeed the term fairness gerrymandering ) is from Kearns et al.
The basic idea behind this "attack" is that while a classifier's output is fair with respect with FPR and FNR for sat race and gender individually, they might no longer be fair when we combine race and gender. Before going into the details of an example, we would like to point out that this at a high level is the same issue as that of intersectionality $\widetilde{\sim}$ that was coined by Kimberlé Crenshaw ©. Here is a TED talk by Crenshaw on this (warning: there are some graphic violence scenes towards the end of the video):


## Consider this situation

> Man
> Woman

## FNR and FPR of various groups?

$$
\begin{aligned}
& \text { Man } \\
& \text { Woman } \\
& \begin{array}{l|llllll|lllllll} 
& \Delta & \Delta & \Delta & 0 & 0 & 0 & \Delta & \Delta & \Delta & 0 & 0 & 0 \\
\text { Black } & \Delta & \Delta & \Delta & 0 & 0 & 0 & \Delta & \Delta & \Delta & 0 & 0 & 0 \\
\hline & \text { White } & \Delta & \Delta & \Delta & 0 & 0 & 0 & \Delta & \Delta & \Delta & 0 & 0 & 0 \\
\hline & \Delta & \Delta & 0 & 0 & 0 & \Delta & \Delta & \Delta & 0 & 0 & 0 \\
\hline
\end{array}
\end{aligned}
$$

## Individual Fairness

## Individual fairness

We say that a classifier is fair if it treats two "similar" individuals "similarly." Note that this is the first notion of fairness that we started off in these notes.

## The natural followup questions are:

1. How do we determine how similar two individuals are.
2. How do we define what it means for a classifier to treat two individuals similarly.

## Shortcoming of individual fairness

Dwork et al. state that society will need to decide on what it means for two individuals to be similar and once this notion of similarity is well-defined it can be used to answer the first followup question above.

In my personal opinion, the above is a really strong assumption and is a shortcoming of the proposal of individual fairness. The main reason is that soliciting much simpler information from humans is hard-- trying to elicit a "true" distance between individuals for all human beings is not realistic.

The notion of individual fairness get get over the shortcomings of group fairness that we talked about-- see Dwork et al. for more details.

## In between individual and group fairness

Kearns et al. define a notion of fairness that potentially holds against (exponentially or even infinitely) many subgroups. The paper then shows how one can compute the optimal model with these fairness constraints. Please see the paper for more details.

