


ETHICAL ASPECTS IN MACHINE LEARNING

Advanced Machine Learning

- Machine learning methods are more and more applied in real-life application, especially for automated decision making:

Fairness in Machine Learning

- Credit scoring and insurance applications — Should the credit/insurance be granted to a certain person or not?
- Rating job applications — Machine learning models can help filter applications much more effectively than simple keyword-based approaches.



Learning goals

- Law — In legal systems around the world, algorithmic tools such as risk assessment instruments (RAI), are being used to supplement or replace the human judgment of judges, civil servants and police officers in many contexts.
- Economics — Automated trading systems buy and sell orders and automatically transmit the orders to market centers or exchanges.
- ...

- Know why fairness aspects are relevant for Machine Learning
- Get familiar with prevalent fairness criteria
- Know their disadvantages and their relationship




ETHICAL ASPECTS IN MACHINE LEARNING

- These are critical applications involving humans, which raises life various ethical issues in different dimensions: making:
 - **Accountability** – Can we make sure that the system is functioning as intended?
 - **Explainability/Transparency**: Is it evident or explainable why one specific decision was made rather than another?
 - **Fairness** – Does the system disadvantage specific individuals or groups?
 - **Privacy** – Is the information (data) on the basis of which the system was developed secure against external access?
 - **Security** – Is it possible to attack the systems, e.g. by "poisoning" the data so that undesirable effects occur?
- It should be noted that all of these aspects are intertwined in some way and becoming increasingly important from a legal perspective, e.g. due to the *European ethics guidelines for trustworthy AI*.



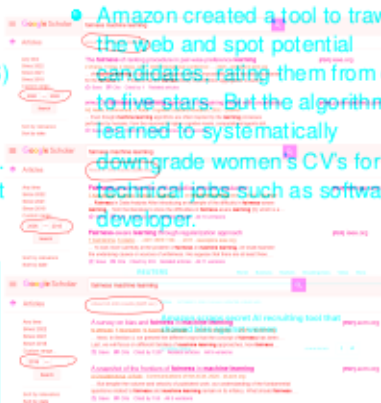
FAIRNESS IN MACHINE LEARNING: WHY BOTHER?

- These are critical applications involving humans, which raises
- In the recent past, there have been a number of automated decision making tools that have attracted attention for discriminatory behavior:
 - **Accountability** — Can we make sure that the system is functioning as intended?
 - **Correctional Offender Management Profiling for Alternative Sanctions (COMPAS)** is a tool used by U.S. courts to assess the likelihood of a defendant becoming a recidivist.
 - **Privacy** — Is the information (data) on the basis of which the system was developed secure against external access?
 - **Security** — Is it possible to attack the system, e.g. by "poisoning" the data so that undesirable effects occur?
 - Amazon created a tool to trawl the web and spot potential candidates, rating them from one to five stars. But the algorithm learned to systematically downgrade women's CV's for technical jobs such as software developer.
- It should be noted that all of these aspects are intertwined in some way and becoming increasingly important from a legal perspective, e.g. due to the *European ethics guidelines for trustworthy AI*.



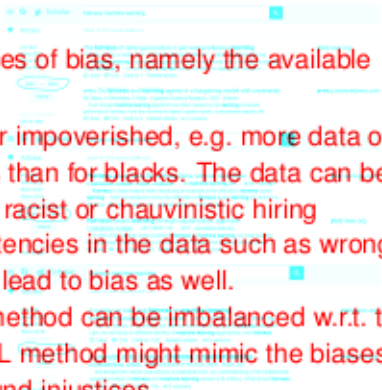
RESEARCH ON FAIRNESS: WHY BOTHER?

- The question of what fairness actually is goes back thousands of years to antiquity. Even back then, philosophers such as Aristotle asked themselves this question.
- The academic research on fairness started with the pioneering works in educational testing (Clearly, 1968) and economics (Becker 1957, Phelps 1972, Arrow 1973).
- In computer science, the research essentially started in the early 2000s and has recently attracted a lot of interest, which is of course due to the increasing use of machine learning models for automated decision making systems.



FAIRNESS IN MACHINE LEARNING: ROUGH OVERVIEW

- The question of what fairness actually is goes back thousands of years to the story of the blind men and an elephant.
- The goal of fairness in Machine Learning is, roughly speaking, to identify and mitigate or even prevent biases of any kind in the decision making based on ML methods along all aspects of the pipeline.
- The academic research on fairness started with the pioneering works in educational testing (Clearly, 1968) and economics (Becker 1957, Phelps 1972, Arrow 1973).
- In computer science, the research essentially started in the early 2000s and has recently attracted a lot of interest, which is of course due to the increasing use of machine learning for automated decision making systems.



FAIRNESS-AWARE BINARY CLASSIFICATION

OVERVIEW

- We will consider only the binary classification/prediction setting in this course, due to the following reasons:
 - 1 The majority of fairness-critical applications in real life are in fact binary classification/prediction tasks, e.g. credit application (granting vs. not granting a loan), job application (hire applicant or not), trading systems (buy or sell an order), ...
 - 2 Quantifying fairness based on a binary outcome variable is mathematically more convenient, while the multi-class variant would introduce additional terms in the fairness quantities. Moreover, multi-class problems can (and are) often addressed by reducing them into multiple binary classification tasks, e.g. one-vs-rest, one-vs-one or error-correcting-codes approaches.
 - 3 The principles for fairness-aware regression tasks are often just modifications of the ones for binary classification.
- There are essentially two sources of bias, namely the available data and the ML model itself:
 - Data can be biased or contain bias, e.g. due to wrong labels or simply noise in the data such as wrong labels or simply noise can lead to bias as well. The data can be biased, e.g. collected by a racist or chauvinistic hiring manager. Finally, inconsistencies in the data such as wrong labels or simply noise can lead to bias as well.
 - The prediction of the ML method can be imbalanced w.r.t. to the error. Moreover, the ML method might mimic the biases in the data and even compound injustices.



FAIRNESS-AWARE BINARY CLASSIFICATION: FORMAL SETTING



● We will consider only the binary classification/prediction setting in this course, due to the following reasons:

- \mathcal{X} is the input/feature/attribute space with $p = \dim(\mathcal{X})$,
- \mathcal{Y} the output/target/label space (for now $\mathcal{Y} = \{-1, 1\}$),
- the tuple $(\mathbf{x}^{(i)}, y^{(i)}) \in \mathcal{X} \times \mathcal{Y}$ is the i -th observation, (granting vs. not granting a loan), job application (hire applicant or not), trading systems (buy or sell an order), ...
- $\mathbf{x}_j = \begin{pmatrix} x_j^{(1)} \\ \vdots \\ x_j^{(n)} \end{pmatrix}$ the j -th feature vector.


So we have observed n objects, described by p features.

- We assume the observed data \mathcal{D} to be generated by a process that can be characterized by some probability distribution \mathbb{P} , defined on $\mathcal{X} \times \mathcal{Y}$.
- In particular, $((\mathbf{x}^{(1)}, y^{(1)}), \dots, (\mathbf{x}^{(n)}, y^{(n)}))$ are i.i.d. with $(\mathbf{x}^{(i)}, y^{(i)}) \sim \mathbb{P}$.
- We denote the random variables (vectors) following this distribution by lower case \mathbf{x} and y .

one-vs-one or error-correcting-codes approaches.

The ultimate goal for a machine learning model f is then loosely speaking "to predict y from \mathbf{x} ", which leads to a decision $f(\mathbf{x}) = \hat{y} \in \{-1, 1\}$.

Note that \hat{y} is a random variable (can be constant), as it is essentially a function of the random input \mathbf{x} .



DECISION THEORY 101

BINARY CLASSIFICATION: FORMAL SETTING



- In binary classification, we typically call one class "positive" and the other "negative".
- We are provided with a data set $\mathcal{D} = ((x^{(1)}, y^{(1)}), \dots, (x^{(n)}, y^{(n)})) \in (\mathcal{X} \times \mathcal{Y})^n$, where
 - \mathcal{X} is the input/feature/attribute space with $n = \dim(\mathcal{X})$
 - \mathcal{Y} is the output/target / label space (for now $\mathcal{Y} = \{-1, 1\}$),
 - the tuple $(x^{(i)}, y^{(i)}) \in \mathcal{X} \times \mathcal{Y}$ is the i -th observation.

- $x_j = (x_j^{(1)}, \dots, x_j^{(n)})^T$ the j -th feature vector.
- So we have observed n objects, described by n features.
- | | | | |
|--------------------|---|---------------------|---------------------|
| | | True Class y | |
| | | + | - |
| Decision \hat{y} | + | True Positive (TP) | False Positive (FP) |
| | - | False Negative (FN) | True Negative (TN) |
- We assume the observed data \mathcal{D} to be generated by a process that can be characterized by some probability distribution P_{xy} , defined on $\mathcal{X} \times \mathcal{Y}$.
 - In particular, $((x^{(1)}, y^{(1)}), \dots, (x^{(n)}, y^{(n)}))$ is i.i.d. with $(x^{(i)}, y^{(i)}) \sim P_{xy}$.


Here: We denote the random variables (vectors) following this distribution by lowercase x and y .

- True Positive (TP)** means that we decide for +1 for a given instance that is really a +1 (correct decision).
- False Positive (FP)** means that we decide for +1 for a given instance that is actually a -1 (incorrect decision).
- False Negative (FN)** means that we decide for -1 for a given instance that is actually a +1 (incorrect decision).
- True Negative (TN)** means that we decide for -1 for a given instance that is really a -1 (correct decision).

The ultimate goal for a machine learning model

is to find a function $f(x) = \hat{y} \in \{-1, 1\}$, which is correct decision.

Note that \hat{y} is a random variable (can be constant), as it is essentially a function of the random input x .



DECISION THEORY 101

The confusion matrix gives rise to common classification/decision criteria, which highlight different aspects of the decision making.

- In binary classification, we typically call one class "positive" and the other "negative".
- The positive class is the more important, often smaller one.
- The confusion matrix gives an overview over the errors as well as correct decisions in a tabulated form:

		True Class y		
		+	-	
Decision \hat{y}	+	TP	FP	$\rho_{PPV} = \frac{TP}{TP+FP}$
	-	FN	TN	$\rho_{NPV} = \frac{TN}{FN+TN}$
Decision		$\rho_{TPR} = \frac{TP}{TP+FN}$	$\rho_{FNR} = \frac{FP+TN}{FP+TN}$	$\rho_{ACC} = \frac{TP+TN}{TOTAL}$

- True positive rate ρ_{TPR} : for how many of the true 1s did we decide for 1?

Population counterpart: $P(\hat{y} = 1 | y = 1)$

- True Negative rate ρ_{TNR} : for how many of the true -1s did we decide for -1?

Population counterpart: $P(\hat{y} = -1 | y = -1)$

- Positive predictive value ρ_{PPV} : if we decide for 1, how likely is it a true 1?

False Positive (FP) means that we decide for +1 for a given instance that is actually a -1

Population counterpart: $P(y = -1 | \hat{y} = +1)$

- Negative predictive value ρ_{NPV} : if we decide for -1, how likely is it a true -1?

False Negative (FN) means that we decide for -1 for a given instance that is actually a +1

Population counterpart: $P(y = +1 | \hat{y} = -1)$

- Accuracy ρ_{ACC} : for how many instances did we decide correctly?

True Negative (TN) means that we decide for -1 for a given instance that is really a -1

Population counterpart: $P(\hat{y} = y)$



DECISION THEORY 101

The confusion matrix gives rise to common classification decision criteria, which highlight different aspects of the decision making



		True Class y			
		\pm	$=$		
Decision \hat{y}	\pm	TP	FP	$p_{PPV} \equiv \frac{TP}{TP+FP}$	$\frac{TP}{TP+FP}$
	$=$	FN	TN	$p_{NPV} \equiv \frac{TN}{FN+TN}$	$\frac{TN}{FN+TN}$
		$p_{TPR} \equiv \frac{TP}{TP+FN}$	$p_{TNR} \equiv \frac{TN}{FP+TN}$	$p_{ACC} \equiv \frac{TP+TN}{TOTAL}$	

- True positive rate p_{TPR} : for how many of the true 1s did we decide for 1?
- False positive rate $p_{FPR} = \frac{FP}{FP+TN}$: for how many of the true -1s did we decide for +1?
- Population counterpart: $P(\hat{y} = +1 | y = 1)$
- True Negative rate p_{TNR} : for how many of the true -1s did we decide for -1?
- Population counterpart: $P(\hat{y} = -1 | y = -1)$
- Population counterpart: $P(\hat{y} = -1 | y = -1)$
- False Negative rate $p_{FNR} = \frac{FN}{FN+TN}$: for how many of the true 1s did we decide for -1?
- Positive predictive value p_{PPV} : if we decide for 1, how likely is it a true 1?
- Population counterpart: $P(y = 1 | \hat{y} = 1)$
- Negative predictive value p_{NPV} : if we decide for -1, how likely is it a true -1?
- Population counterpart: $P(y = -1 | \hat{y} = -1)$
- Error rate $p_{err} = 1 - p_{acc}$: for how many instances did we decide incorrectly?
- Accuracy p_{acc} : for how many instances did we decide correctly?
- Population counterpart: $P(\hat{y} = y)$
- Population counterpart: $P(\hat{y} = y)$

SENSITIVE ATTRIBUTES/FEATURES

The confusion matrix gives rise to common classification/decision criteria, which highlight different aspects of the decision making

- The aspect of fairness usually arises due to the presence of sensitive attributes/features among the attributes/features $\mathbf{x}_1, \dots, \mathbf{x}_p$, e.g. age, gender, nationality, race, ... True Class y

- Note that we assume that the attribute/feature observations $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$ are random observations of the random vector $\mathbf{x} = (x_1, \dots, x_p)^T$ with distribution \mathbb{P}_x . Accordingly, the j -th attribute/feature vector

$\mathbf{x}_j = (x_j^{(1)}, \dots, x_j^{(n)})^T$ is a collection of random observations of the random variable x_j with distribution \mathbb{P}_{x_j} , which is the marginal distribution of \mathbb{P}_x for the j -th attribute/feature.

- False positive rate $p_{FPR} = \frac{FP}{FP+TN}$: for how many of the true -1s did we decide for +1?

- We introduce the random variable \mathbf{A} to capture all sensitive attributes/features, which typically has discrete values.

- The basic idea of fairness criteria introduced for machine learning

→ methods is to equalize different decision criteria or statistical quantities

- involving \mathbf{A} : Error $p_{err} = 1 - p_{ACC}$ for how many instances did we decide incorrectly?

→ This goes back to Anne Cleary in the 1960s who studied group differences in educational testing.
 → Population counterpart: $\mathbb{P}(y \neq \hat{y})$



REMOVING SENSITIVE FEATURES

- A straightforward (and also naive) approach is to simply ignore or remove all sensitive features at prediction time. This approach is often called *fairness through unawareness*. However, in many cases other non-sensitive features are slightly correlated with the sensitive one(s). For example:
 - gender with hobbies or interests (job application),
 - race and zip code (law systems),
 - nationality and location id (credit application).
- Thus, an ML model trained on data including the sensitive features might combine the corresponding correlated non-sensitive features to make essentially the same decision, as it still seeks to maximize accuracy.
- The basic idea of fairness criteria introduced for machine learning methods is to equalize different decision criteria or statistical quantities involving \mathbf{A} .

This goes back to Anne Cleary in the 1960s who studied group differences in educational testing.



INDEPENDENCE AS A FAIRNESS CRITERION

- A quite natural fairness criterion is given by ensuring (stochastic) independence between the decision \hat{y} and the sensitive attributes/features \mathbf{A} .
This is often called fairness through unawareness. However, in many cases other non-sensitive features are slightly correlated with the sensitive one(s).


- For example, this is equivalent to ensuring an equal "acceptance rate" among all possible realizations $\mathbf{a}, \tilde{\mathbf{a}}$ of \mathbf{A} :

- race and zip code (law systems),
- nationality and location (credit application),
- ...

$$\mathbb{P}(\hat{y} = 1 \mid \mathbf{A} = \mathbf{a}) = \mathbb{P}(\hat{y} = 1 \mid \mathbf{A} = \tilde{\mathbf{a}})$$

- This criterion is also known as statistical/demographic parity, group fairness, equal positive rates or Darlington's fourth criterion.
- Thus, an ML model trained on data including the sensitive features might combine the corresponding correlated non-sensitive features to make a prediction.
- One can relax the criterion by introducing a fixed tolerance parameter $\epsilon > 0$ and only require that for all possible realizations $\mathbf{a}, \tilde{\mathbf{a}}$ of \mathbf{A} it holds that

$$\left| \mathbb{P}(\hat{y} = 1 \mid \mathbf{A} = \mathbf{a}) - \mathbb{P}(\hat{y} = 1 \mid \mathbf{A} = \tilde{\mathbf{a}}) \right| \leq \epsilon$$



DOWNSIDERS OF INDEPENDENCE AS A FAIRNESS CRITERION

- A quite natural fairness criterion is given by ensuring (stochastic) independence between the decision \hat{y} and the sensitive attributes/features \mathbf{A} .
- Independence does not take into account that the outcome y might be correlated with \mathbf{A} , which means that the different realizations of \mathbf{A} have different underlying distributions for y .
- Not considering this dependency can lead to decisions which are fair through the lens of the independence criterion, but not for the groups themselves.

$$\mathbb{P}(\hat{y} = 1 \mid \mathbf{A} = \mathbf{a}) = \mathbb{P}(\hat{y} = 1 \mid \mathbf{A} = \bar{\mathbf{a}})$$

- Moreover, independence does not rule out the possibility of unfair practices. For example, consider a job hiring process involving different groups of people. Assume that we
- make thoughtful and good decisions in one specific group with accepting people from that group with a rate $p \in (0, 1)$.
- make poor and bad decisions in all other groups with the same acceptance rate $p \in (0, 1)$, respectively.

$$|\mathbb{P}(\hat{y} = 1 \mid \mathbf{A} = \mathbf{a}) - \mathbb{P}(\hat{y} = 1 \mid \mathbf{A} = \bar{\mathbf{a}})| \leq \epsilon$$



ACHIEVING INDEPENDENCE VIA REPRESENTATION LEARNING AS A FAIRNESS CRITERION

- Independence does not take into account that the outcome y might be correlated with \mathbf{A} , which means that the different realizations of \mathbf{A} have different underlying distributions for y .
- One common idea to satisfy the independence criterion is by finding a "fair representation" \mathbf{Z} of the data \mathbf{x} , i.e., one such that $\mathbf{Z} \perp \mathbf{A}$ holds. Then, the ML method f uses \mathbf{Z} instead of \mathbf{x} for the decision: $\hat{y} = f(\mathbf{Z})$



- Moreover, independence does not rule out the possibility of unfair practices. For example, consider a job hiring process involving different groups of people. Assume that we the representation are formulated:
 - Information about \mathbf{x} should be preserved \Leftrightarrow Mutual information between \mathbf{x} and \mathbf{Z} is high.
 - make thoughtful and good decisions in one specific group with accepting people from that group with a rate $p \in (0, 1)$.
 - The sensitive attributes/features \mathbf{A} are obfuscated \Leftrightarrow Mutual information between \mathbf{A} and \mathbf{Z} is low.
 - make poor and bad decisions in all other groups with the same acceptance rate $p \in (0, 1)$, respectively.
 - Accuracy of the model f using \mathbf{Z} is (still) high \Leftrightarrow Mutual information between y and \mathbf{Z} is high.



SEPARATION AS A FAIRNESS CRITERION

REPRESENTATION LEARNING

- As we discussed above, the independence criterion does not take correlation between y and \mathbf{A} into account. As an alternative fairness criterion one can consider separation, which ensures (stochastic) independence between the decision \hat{y} and the sensitive attributes \mathbf{A} given a representation \mathbf{Z} of the data \mathbf{x} , i.e., one such that $\mathbf{Z} \perp\!\!\!\perp \mathbf{A} \mid y$.
- This is equivalent to equalize the (population) error rates for all possible realizations $\mathbf{a}, \bar{\mathbf{a}}$ of \mathbf{A} :

$$\mathbb{P}(\hat{y} = 1 \mid y = -1, \mathbf{A} = \mathbf{a}) = \mathbb{P}(\hat{y} = 1 \mid y = -1, \mathbf{A} = \bar{\mathbf{a}})$$

- The idea goes back to Zemel et al. (2013), where (equal false positive rates) the representation are formulated

$$\mathbb{P}(\hat{y} = -1 \mid y = 1, \mathbf{A} = \mathbf{a}) = \mathbb{P}(\hat{y} = -1 \mid y = 1, \mathbf{A} = \bar{\mathbf{a}})$$

- Information about \mathbf{x} should be preserved \Leftrightarrow Mutual information between \mathbf{x} and \mathbf{Z} is high
 - The sensitive attributes/features \mathbf{A} are obfuscated \Leftrightarrow Mutual information between \mathbf{A} and \mathbf{Z} is low
 - Accuracy of the model f using \mathbf{Z} is (still) high \Leftrightarrow Mutual information between y and \mathbf{Z} is high
- The idea is that all realizations of \mathbf{A} experience the same FPR and FNR.
 - This criterion is also known as equalized odds, avoiding disparate mistreatment, equalized error rates or conditional procedure accuracy.
 - This is a posthoc criterion, as it is not known at the time of the decision whether the current instance is positive or negative. Only in hindsight the positive and negative instances can be collected and compared with the decisions made.




INTERLUDE ROC SPACE LESS CRITERION

- For comparing classifiers, we characterize them by their TPR and FPR values and plot them in a coordinate system.
- We could also use two different ROC metrics (decision criteria) which define a trade-off, for instance, TPR and PPV.



$$\hat{y} \perp\!\!\!\perp A \mid y$$

- This is equivalent to equalize the (population) error rates for all possible realizations a, \bar{a} of A :



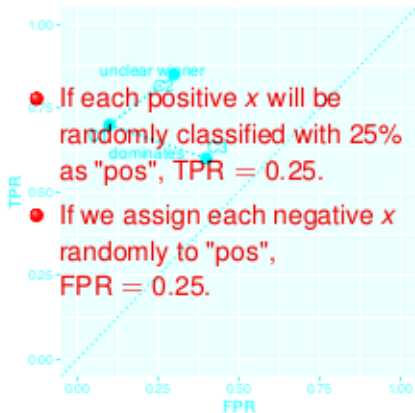
- The idea is that all realizations of A experience the same FPR and FNR.
- This criterion is also known as equalized odds, avoiding disparate mistreatment, equalized error rates or conditional procedure accuracy.
- This is a posthoc criterion, as it is not known at the time of the decision whether the current instance is positive or negative. Only in hindsight the positive and negative instances can be collected and compared with the decisions made.

$$TPR = \frac{TP}{TP + FN}$$


$$FPR = \frac{FP}{FP + TN}$$

INTERLUDE: ROC SPACE

- The best classifier lies on the top-left corner, where FPR equals 0 and TPR is maximal.
- We could also use two different ROC metrics (decision criteria) which define a trade-off, for instance, TPR and PPV.
- The diagonal is worst as it corresponds to a classifier producing random labels (with different proportions).



- If each positive x will be randomly classified with 25% as "pos", $TPR = 0.25$.
- If we assign each negative x randomly to "pos", $FPR = 0.25$.



INTERLUDE: ROC CURVES FOR SCORING CLASSIFIERS

- The best classifier lies on the top-left corner, where FPR equals 0 and TPR is maximal.
- Many binary classification methods use a score (function) $s: \mathcal{X} \rightarrow \mathbb{R}$ and a threshold value c to make the prediction (decision):
- The diagonal is worst as it corresponds to a classifier producing random labels (with different proportions).

$$f(\mathbf{x}) = 2 \cdot \mathbb{I}_{[s(\mathbf{x}) \geq c]} - 1.$$

- The choice of threshold affects the TPR and FPR, so it is interesting to examine the effects of different thresholds on these.
- If each positive x will be randomly classified with 25% as "pos", TPR = 0.25.
- If we assign each negative x randomly to "pos", FPR = 0.25.



INTERLUDE: ROC CURVES FOR SCORING CLASSIFIERS / 2

To draw a ROC curve:

Many binary classification methods use a score (function) $s: \mathcal{X} \rightarrow \mathbb{R}$ (W.l.o.g. $s: \mathcal{X} \rightarrow [0, 1]$) and a threshold value c to make the prediction (decision):

- Rank test observations on decreasing score. $(s(x) \geq c) \rightarrow -1$.
- Start with $c = 1$, so we start in $(0, 0)$; we predict everything as negative.
The choice of threshold affects the TPR and FPR, so it is interesting to examine the effects of different thresholds on these.
- Iterate through all possible thresholds c and proceed for each observation x as follows:
 - If x is positive, move TPR $1/n_+$ up, as we have one TP more.
 - If x is negative, move FPR $1/n_-$ right, as we have one FP more.

(n_+ : number of positives, n_- : number of negatives.)

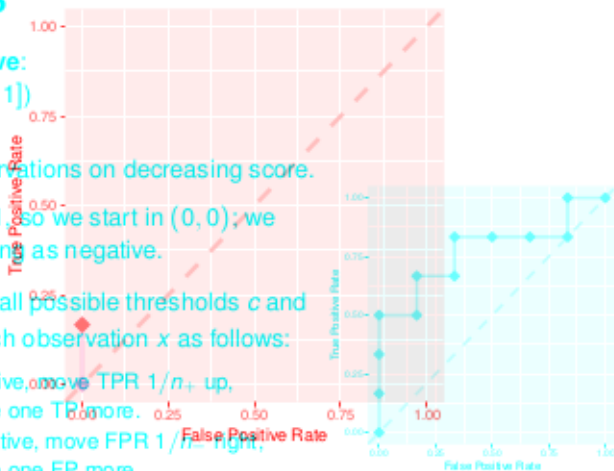


DRAWING ROC CURVES EXAMPLE CLASSIFIERS

To draw a ROC curve:
(With θ 's: $\chi \rightarrow [0, 1]$)

- Rank test observations on decreasing score.
- Start with $c = 1$, so we start in $(0, 0)$; we predict everything as negative.
- Iterate through all possible thresholds c and proceed for each observation x as follows:
 - If x is positive, move TPR $1/n_+$ up, as we have one TP more.
 - If x is negative, move FPR $1/n_-$ right, as we have one FP more.

#	Truth	Score
1	Pos	0.95
2	Pos	0.85
3	Pos	0.69
4	Neg	0.65
5	Pos	0.59
6	Neg	0.52
7	Pos	0.51
8	Neg	0.39
9	Neg	0.28
10	Neg	0.18
11	Pos	0.15
12	Neg	0.05



$c = 0.9$ (number of positives, n_+ : number of negatives.)


→ TPR = 0.167

→ FPR = 0



DRAWING ROC CURVES: EXAMPLE

#	Truth	Score
1	Pos	0.95
2	Pos	0.86
3	Pos	0.69
4	Neg	0.65
5	Pos	0.59
6	Neg	0.52
7	Pos	0.51
8	Neg	0.39
9	Neg	0.28
10	Neg	0.18
11	Pos	0.15
12	Neg	0.06



$$c = 0.85$$


$$\rightarrow \text{TPR} = 0.333$$

$$\rightarrow \text{FPR} = 0$$



DRAWING ROC CURVES: EXAMPLE

#	Truth	Score
1	Pos	0.95
2	Pos	0.86
3	Pos	0.69
4	Neg	0.65
5	Pos	0.59
6	Neg	0.52
7	Pos	0.51
8	Neg	0.39
9	Neg	0.28
10	Neg	0.18
11	Pos	0.15
12	Neg	0.06



$$c = 0.66$$


$$\rightarrow \text{TPR} = 0.533$$

$$\rightarrow \text{FPR} = 0$$



DRAWING ROC CURVES: EXAMPLE


#	Truth	Score
1	Pos	0.95
2	Pos	0.86
3	Pos	0.69
4	Neg	0.65
5	Pos	0.59
6	Neg	0.52
7	Pos	0.51
8	Neg	0.39
9	Neg	0.28
10	Neg	0.18
11	Pos	0.15
12	Neg	0.06



$$c = 0.66$$


$$\rightarrow \text{TPR} = 0.5$$

$$\rightarrow \text{FPR} = 0.167$$



DRAWING ROC CURVES: EXAMPLE


#	Truth	Score
1	Pos	0.95
2	Pos	0.86
3	Pos	0.69
4	Neg	0.65
5	Pos	0.59
6	Neg	0.52
7	Pos	0.51
8	Neg	0.39
9	Neg	0.28
10	Neg	0.18
11	Pos	0.15
12	Neg	0.06



$c = 0.55$


→ TPR = 0.667

→ FPR = 0.167



DRAWING ROC CURVES: EXAMPLE


#	Truth	Score
1	Pos	0.95
2	Pos	0.86
3	Pos	0.69
4	Neg	0.65
5	Pos	0.59
6	Neg	0.52
7	Pos	0.51
8	Neg	0.39
9	Neg	0.28
10	Neg	0.18
11	Pos	0.15
12	Neg	0.06



$$c = 0.35$$


$$\rightarrow \text{TPR} = 0.833$$

$$\rightarrow \text{FPR} = 0.567$$



DRAWING ROC CURVES: EXAMPLE

#	Truth	Score
1	Pos	0.95
2	Pos	0.86
3	Pos	0.69
4	Neg	0.65
5	Pos	0.59
6	Neg	0.52
7	Pos	0.51
8	Neg	0.39
9	Neg	0.28
10	Neg	0.18
11	Pos	0.15
12	Neg	0.06



$$c = 0.3$$

$$\rightarrow \text{TPR} = 0.833$$

$$\rightarrow \text{FPR} = 0.5$$



SEPARATION AND ROC CURVES

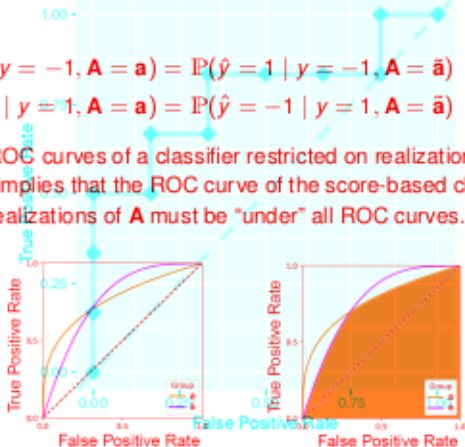
- Separation, i.e.,

#	Truth	Score
1	Pos	0.95
2	Pos	0.86
3	Pos	0.63
4	Neg	0.65
5	Neg	0.75
6	Neg	0.57
7	Pos	0.51
8	Neg	0.39
9	Neg	0.28
10	Neg	0.18
11	Pos	0.15
12	Neg	0.06

$$P(\hat{y} = 1 | y = -1, \mathbf{A} = \mathbf{a}) = P(\hat{y} = 1 | y = -1, \mathbf{A} = \bar{\mathbf{a}}) \quad (\text{equal FPRs})$$

$$P(\hat{y} = -1 | y = 1, \mathbf{A} = \mathbf{a}) = P(\hat{y} = -1 | y = 1, \mathbf{A} = \bar{\mathbf{a}}) \quad (\text{equal FNRs})$$


- means that all ROC curves of a classifier restricted on realizations of \mathbf{A} should be the same. This implies that the ROC curve of the score-based classifier conditional on realizations of \mathbf{A} must be "under" all ROC curves.



- $c = 0$
- $\rightarrow \text{TPR} = 1$
- $\rightarrow \text{FPR} = 1$

- In practice, we should never obtain a classifier below the diagonal.

- Inverting the predicted labels ($-1 \mapsto 1$ and $1 \mapsto -1$) will result in a reflection at the diagonal $\Rightarrow \text{TPR}_{\text{now}} = 1 - \text{TPR}$ and $\text{FPR}_{\text{now}} = 1 - \text{FPR}$.

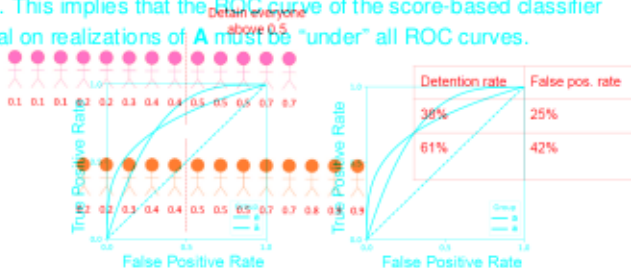


DOWNSIDERS OF SEPARATION AS A FAIRNESS CRITERION

- Separation, i.e.,

- Consider two groups of people: blue and orange. We are interested to decide whether we should detain (positive class) a person and use a scoring classifier with scores in $[0, 1]$ and a threshold $c = 0.5$.

means that all ROC curves of a classifier restricted on realizations of A should be the same. This implies that the ROC curve of the score-based classifier conditional on realizations of A must be "under" all ROC curves.




- The classifier is not satisfying separation as FPR and FNR are not the same among the two groups.
- Inverting the predicted labels ($-1 \mapsto 1$ and $1 \mapsto -1$) will result in a reflection at the diagonal $\Rightarrow TPR_{\text{now}} = 1 - TPR$ and $FPR_{\text{now}} = 1 - FPR$.



Downsides of Separation as a Fairness Criterion / 2

- In order to achieve separation we would need to arrest more low risk individuals in the orange group.
- Consider two groups of people: blue and orange. We are interested to decide whether we should detain (positive class) a person and use a scoring classifier with scores in $[0, 1]$ and a threshold $c = 0.5$.



- Thus, as with achieving independence, separation can also lead to undesirable outcomes.
- Unreasonable outcomes: satisfying separation as FPR and FNR are not the same among the two groups.



SUFFICIENCY AS A FAIRNESS CRITERION

- Another idea to specify a fairness criterion for score-based classifiers is
- that the corresponding score random variable $S = s(x)$ already subsumes the sensitive attributes A for the prediction:



- This is equivalent to require that the fraction of positive instances assigned some score s is the same for all possible realizations a, \tilde{a} of A :




$$P(y = 1 | S = s, A = a) = P(y = 1 | S = s, A = \tilde{a})$$

- This criterion is also known as Cleary's model, conditional use accuracy or calibration within groups.
- This is an a priori guarantee: The decision maker sees the score value and knows based on this what the frequency of positives is.

SUFFICIENCY AND CALIBRATION CRITERION

- Sufficiency is very closely related to the concept of calibration of a probabilistic classifier, i.e., a classifier such that $s = \mathbb{P}(y = 1 | \mathbf{S} = s, \mathbf{A})$ for the prediction.
- Another idea to specify a fairness criterion for score-based classifiers is that the corresponding score random variable $\mathbf{S} = s(\mathbf{x})$ already satisfies the sufficiency attribute \mathbf{A} for the prediction:

$$y \perp\!\!\!\perp \mathbf{A} | \mathbf{S}$$



- This is equivalent to require that the fraction of positive instances assigned some score s is the same for all possible realizations $\mathbf{a}, \tilde{\mathbf{a}}$ of \mathbf{A} :
 - Note that:
 - 1 This condition means that the set of all instances assigned a score value s also account for a proportion s of positive instances: $\mathbb{P}(y = 1 | \mathbf{S} = s, \mathbf{A} = \mathbf{a}) = \mathbb{P}(y = 1 | \mathbf{S} = s, \mathbf{A} = \tilde{\mathbf{a}})$
 - 2 It is a condition over all features and in particular on the sensitive ones. Consequently, it does not mean that at the level of a single value of \mathbf{A} a score of s corresponds to a probability s of a positive outcome.
- This criterion is also known as Cleary's model, conditional use accuracy or calibration within groups.
- The notion of calibration can be specified also on the group level, that is, a probabilistic classifier is called *calibrated on the group level* if for all $s \in [0, 1]$ and all possible realizations \mathbf{a} of \mathbf{A} :

$$\mathbb{P}(y = 1 | \mathbf{S} = s, \mathbf{A} = \mathbf{a}) = s.$$

SUFFICIENCY AND CALIBRATION / 2

- Sufficiency is very closely related to the concept of calibration of a probabilistic classifier, i.e., a classifier such that $s : \mathcal{X} \rightarrow [0, 1]$ is more specific, by a probabilistic classifier is called *calibrated* if for all $s \in [0, 1]$
$$P(y = 1 | \mathbf{S} = s, \mathbf{A} = \mathbf{a}) = P(y = 1 | \mathbf{S} = s, \mathbf{A} = \bar{\mathbf{a}}) \quad (\text{sufficiency})$$

vs.
$$P(y = 1 | \mathbf{S} = s) = s \quad (\text{calibration on individual level})$$
- If a probabilistic classifier is *calibrated on the group level*, then it also satisfies sufficiency.

- If a probabilistic classifier f satisfies sufficiency, then we can find a function $C : [0, 1] \rightarrow [0, 1]$ such that f based on $C(\mathbf{S})$ (instead of \mathbf{S}) is calibrated on the group level.
- Note that: this condition means that the set of all instances assigned a score value s also account for a proportion s of positive instances.
- Sufficiency is only slightly weaker, but it is fair to say that both properties are essentially equivalent. It is a condition over all features and in particular on the sensitive ones. Consequently, it does not mean that at the level of a single value of \mathbf{A} a score of s corresponds to a probability s of a positive outcome.

- The notion of calibration can be specified also on the group level, that is, a probabilistic classifier is called *calibrated on the group level* if for all $s \in [0, 1]$ and all possible realizations \mathbf{a} of \mathbf{A} :

$$P(y = 1 | \mathbf{S} = s, \mathbf{A} = \mathbf{a}) = s.$$



DOWNSIDERS OF SUFFICIENCY AS A FAIRNESS CRITERION


$$P(y = 1 | S = s, A = a) = P(y = 1 | S = s, A = \bar{a}) \quad (\text{sufficiency})$$

- Consider a group of blue people and assume we are interested in deciding whether we should detain (positive class) a person and use a scoring classifier with scores in $[0, 1]$ and a threshold $c = 0.5$. Suppose we know the true probability that a person will reoffend and the scores are equal to these.
- If a probabilistic classifier f satisfies sufficiency, then we can find a function $C : [0, 1] \rightarrow [0, 1]$ such that f based on $C(S)$ (instead of S) is calibrated on the group level.
- Sufficiency is only slightly weaker, but it is fair to say that both properties are essentially equivalent.
- Assume that there are two groups among the blue people:



Downsides of Sufficiency as a Fairness Criterion / 2

- Consider a group of blue people and assume we are interested in deciding whether we should detain (positive class) a person and use a scoring classifier with scores in $[0, 1]$ and a threshold $c = 0.5$. Suppose we know the true probability that a person will reoffend and the scores are equal to these.



If we calibrate the classifier, we have no detentions any more!




- Assume that there are two groups among the blue people:



RELATIONSHIPS BETWEEN THE FAIRNESS CRITERIA

- We have considered three fairness criteria:



$$\hat{y} \perp\!\!\!\perp A$$

(Independence)

$$\hat{y} \perp\!\!\!\perp A \mid y$$

(Separation)

$$y \perp\!\!\!\perp A \mid S$$

(Sufficiency)

If we calibrate the classifier, we have no detections any more!



- A tempting question is how these criteria relate to each other.
- (Informal) Theorem.** Any two of these criteria are mutually exclusive in general.

 0.4 0.4 0.4 0.4 0.4 0.3 0.4 0.4 0.4 0.4 0.4 0.4 0.4
- As a consequence, we cannot impose multiple of these criteria as hard constraints on the classifier.
- A possible solution to this issue is to consider relaxed version of these criteria as constraints.

FINAL REMARKS BETWEEN THE FAIRNESS CRITERIA

- Fairness is a challenging issue as also philosophers and social scientists have been trying to define it for decades.

- Due to the increased use of ML methods in automated decision making, there is a need to think about fairness in more detail.

- Fairness criteria such as independence, sufficiency and separation are a statistical objective way to incorporate fairness aspects into ML methods.

- However, on their own they are neither equivalent to a “proof of fairness” nor are they perfect objective functions for this purpose.

- **(Informal) Theorem.** Any two of these criteria are mutually exclusive in general.

- In summary, there are three ways to tackle the question: “how to satisfy fairness criteria?”

- As a consequence, we cannot impose multiple of these criteria as hard constraints on the classifier.

- A possible solution to this issue is to consider relaxed version of these criteria as constraints.

- 1 Pre-processing phase: Adjust the feature space to be uncorrelated with the sensitive attribute.
- 2 Training phase: Build the constraint into the optimization process for the classifier.
- 3 Post-processing phase: Adjust a learned classifier so that it is uncorrelated to the sensitive attribute.



FINAL REMARKS

- Fairness is a challenging issue as also philosophers and social scientists have been trying to define it for decades.
- Due to the increased use of ML methods in automated decision making there is a need to think about fairness in more detail.
- Fairness criteria such as independence, sufficiency and separation are a statistical objective way to incorporate fairness aspects into ML methods. However, on their own they are neither equivalent to a “proof of fairness” nor are they perfect objective functions for this purpose.
- In summary, there are three ways to tackle the question: “how to satisfy fairness criteria?”
 - 1 Pre-processing phase: Adjust the feature space to be uncorrelated with the sensitive attribute.
 - 2 Training phase: Build the constraint into the optimization process for the classifier.
 - 3 Post-processing phase: Adjust a learned classifier so that it is uncorrelated to the sensitive attribute.