Fairness in ML

$\hat{Y}$ ← prediction of the ML classifier.

$X, Y, \hat{Y}$

Accuracy $\equiv P(Y = \hat{Y})$

Test data

X_1	Y_1	$\hat{Y}_1$
X_2	Y_2	$\hat{Y}_2$
X_3	Y_3	$\hat{Y}_3$
⋮	⋮	⋮
X_{10}	Y_{10}	$\hat{Y}_{10}$

$$\frac{\#(Y = \hat{Y})}{10}$$

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

← True Positive Rate.

X	Y	$\hat{Y}$
x_1	1	1
x_2	0	1
x_3	0	0
x_4	1	0
x_5	1	1
x_6	1	1
x_7	1	0
x_8	1	1
x_9	0	0
x_{10}	1	1

Accuracy $= P(\hat{y} = y) = \frac{7}{10} = 0.7$

$TPR = P(\hat{y} = 1 \mid y = 1) = \frac{5}{7}$

recall on +ve class

$FNR = P(\hat{y} = 0 \mid y = 1) = \frac{2}{7}$

$FPR = P(\hat{y} = 1 \mid y = 0) = \frac{1}{3}$

$$TNR = P(\hat{y}=0 \mid y=0) = \frac{2}{3}$$

recall on -ve class

$$P(y=1 \mid \hat{y}=1) = \frac{5}{6}$$

Precision (on +ve class)

$$P(y=0 \mid \hat{y}=0) = \frac{2}{4}$$

Precision on -ve class

---

f-measure for +ve class

$$= \frac{2}{\frac{1}{\text{recall}_{+ve}} + \frac{1}{\text{precision}_{+ve}}}$$

$$E[y \mid X=x]$$

X	Y	R.	$\hat{y}$ $t=0.5$
x_1	1.	0.80	1
x_2	0.	0.52	1

x_3	0.	0.47	0
x_4	1.	0.60	1
x_5	1..	0.65	1
x_6	1..	0.39	0
x_7	1.	0.49	0
x_8	1..	0.80	1
x_9	0	0.05	0
x_{10}	1.	0.40	0

Receiver Operating Characteristic Curve (ROC)

$t=0$ FPR (+ve class) $P(\hat{y}=1|y=0)$

TPR (+ve class) $P(\hat{y}=1|y=1)$

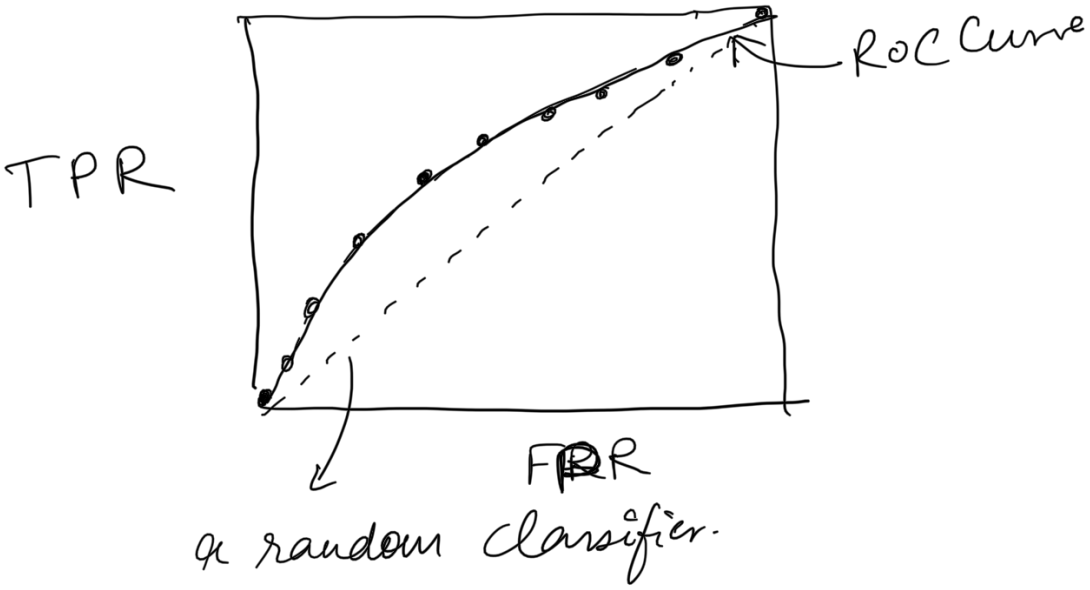

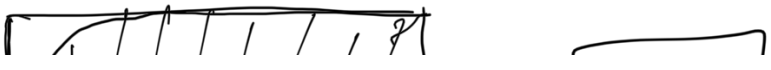

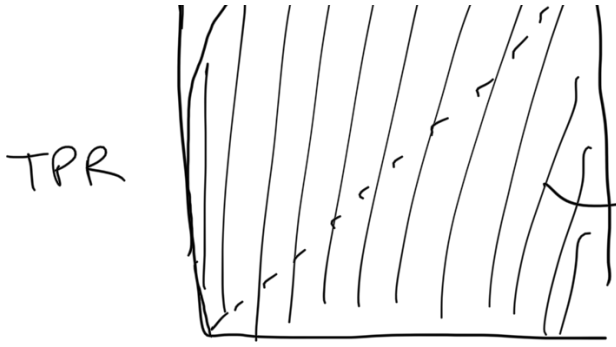

AUC

→ area under the ROC curve.

Sensitive attribute A [A - binary]

$\perp\!\!\!\perp$ → independence

$A \perp\!\!\!\perp B \quad P(A, B) = P(A)\,P(B)$

$A \perp\!\!\!\perp B \mid C \quad P(A, B \mid C) = P(A \mid C)\,P(B \mid C)$

Independence

A classifier is fair if:

$$P(\hat{y}=1 \mid A=a) = P(\hat{y}=1 \mid A=b)$$

Friday April 23

All Gradiance quizzes will be due

on May 9th at 11.59 PM

Gradiance 10, 11, 12

What threshold to use?

$t = 0.5$

Cost of classification

Profit