Evaluation of binary classifiers

MATH-412 - Statistical Machine Learning

Sensitivity, precision and co.

		Predicted	
		"P"	"N"
Actual	P	TP	FN
	N	FP	TN

- **sensitivity**, true positive rate or **recall** $r_{TP} = \frac{|TP|}{|P|}$
- **specificity** or true negative rate $r_{TN} = \frac{|TN|}{|N|}$
- **false positive rate** (type I error) $\alpha = r_{FP} = \frac{|FP|}{|N|} = 1 - r_{TN}$
- **false negative rate** (type II error) $\beta = r_{FN} = \frac{|FN|}{|P|} = 1 - r_{TP}$
- **precision** $prec = \frac{|TP|}{|FP|+|TP|}$

Sensitivity, specificity, etc

	P	N
$\widehat{P}$	Sensitivity (TPR) $\frac{TP}{P}$	FPR $\frac{FP}{N}$
$\widehat{N}$	FNR $\frac{FN}{P}$	Specificity (TNR) $\frac{TN}{N}$

TPR	True Positive Rate
FPR	False Positive Rate
FNR	False Negative Rate
NPR	True Negative Rate

Precision, FDR, etc

	P	N
$\widehat{P}$	Precision (PPV) $\frac{TP}{\widehat{P}}$	FDR $\frac{FP}{\widehat{P}}$
$\widehat{N}$	FOR $\frac{FN}{\widehat{N}}$	NPV $\frac{TN}{\widehat{N}}$

PPV Positive Predictive Value

FDR False Discovery Rate

FOR False Omission Rate

NPV Negative Predictive Value

ROC curve

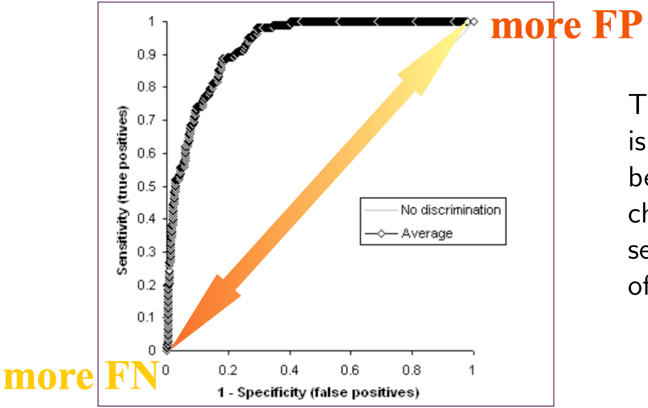

The **Receiver Operating Characteristic** is a representation of the trade-off between the Recall and the FPR as one changes the parameter controlling the sensitivity of the classifier, such as the offset b.

Evaluating ROC curves

Which one is the best ROC curve?

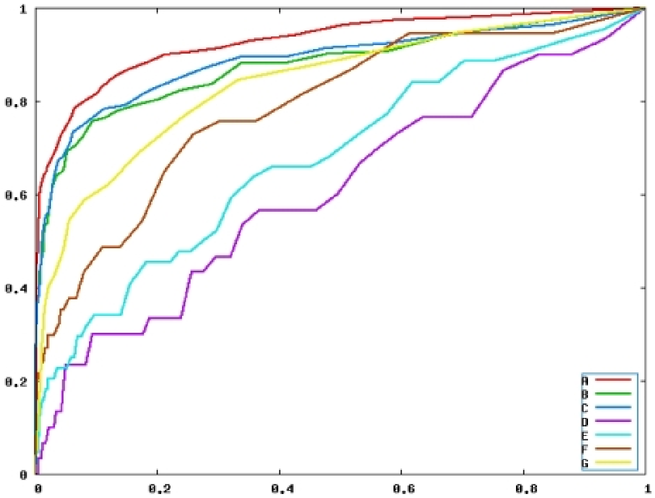

ROC curve and its convex hull

Convexity property of the ROC plot:
Given two points

- (α_0, β_0) for classifier f_0
- (α_1, β_1) for classifier f_1.

Consider the *randomized* classifier f_π which uses f_1 with probability π and f_0 with probability $1 - \pi$. Then

$$(\alpha_\pi, \beta_\pi) = \pi(\alpha_1, \beta_1) + (1 - \pi)(\alpha_0, \beta_0).$$

Attainable points in the ROC plane form a convex set.

The ROC convex hull is the convex envelope of the attainable points, whose upper/left boundary forms a concave Pareto front.

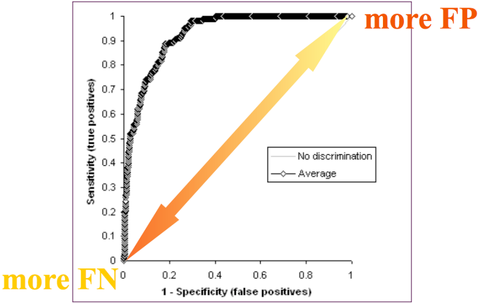

Performance measures derived from the ROC

A number of performance measures can be obtained from the ROC curve.
Let $\pi = \mathbb{P}(Y = 1)$. It is estimated by $\widehat{\pi} = \frac{|P|}{|P|+|N|}$.

- **Misclassification error**

$$\widehat{\mathcal{R}}^{\text{test}}_{\text{0-1}} = \widehat{\pi}\,(1 - r_{TP}) + (1 - \widehat{\pi})\,r_{FP}$$

- **Cost**

$$\widehat{\mathcal{R}}^{\text{test}}_{C_+,C_-} = \widehat{\pi}\,C_+\,(1 - r_{TP}) + (1 - \widehat{\pi})\,C_-\,r_{FP}$$

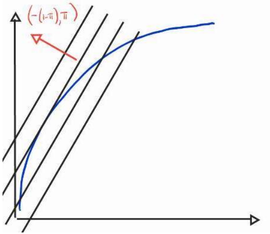

$\rightarrow$ For any cost, can change the threshold or other hyperparameter to minimize the cost.

- **Area under the curve (AUC)**
  Actually corresponds to the Mann-Whitney-Wilcoxon U-statistic aka Wilcoxon signed rank test statistic on the score distribution to distinguish class 1 and class 0.
- **Truncated AUC**

Precision - Recall curve