1/28/14

1 Evaluating Binary Classifiers

y_{gt} \ $\hat{y}$	0	1	
0	TN	FP	$\rightarrow \Sigma = \#\text{Neg-in-GT}$
1	FN	TP	$\rightarrow \Sigma = \#\text{Pos-in-GT}$

TP: True Positive
GT was "positive" (+1)
you called it "positive" (+1)

TN: True Negative
GT was "negative" (class 0)
you called it negative

FP: False Positive
FN: False Negative

$\rightarrow$ TP-Rate (aka 'Recall')

$$= \frac{\#TP}{\#\text{Pos-in-GT}} = \frac{\#TP}{\#TP + \#FN}$$

what % of cancer patients did you identify / find?

$\rightarrow$ FP-Rate

$$= \frac{\#FP}{\#\text{Neg-in-GT}} = \frac{\#FP}{\#FP + \#TN}$$

How frequently do you "hallucinate" cancer?

$\rightarrow$ Note: $Acc = \dfrac{\#TP + \#FP}{\#TP + \#TN + \#FP + \#FN}$

DB: Mistake; should be #TP + #TN

→ Let's examine some "dumb" classifiers

→ $\hat{y}(x) = 1 \quad \forall x$ "Everybody has cancer!"

$$\text{TP-Rate} = \frac{\#TP}{\#\text{Pos-in-GT}} = \underset{(100\%)}{1}$$ Excellent! We found all cancers!

$$\text{FP-Rate} = \frac{\#FP}{\#\text{Neg-in-GT}} = \underset{(100\%)}{1}$$ Ooops!

→ $\hat{y}(x) = 0 \quad \forall x$ "No-body has cancer"

FP-Rate = 0 "Yay, no mistakes"
TP-Rate = 0 "Useless"

→ In general
say $\hat{y}(x) = \text{sign}[\text{Confidence/Score}(x) \geq \underset{t}{\text{threshold}}]$

e.g say 10-NN of x = 3 - Pos, 7 Neg

$$\Rightarrow \quad \text{Positive Score}(x) = \frac{3}{10} = 0.3$$

Compare this to a threshold t to decide if your confidence $\geq t$

As threshold t is varied from $0 \to 1$, you get a curve
at $t = 1$, $\hat{y}(x) = 0 \quad \forall x$
at $t = 0$ $\hat{y}(x) = 1 \quad \forall x$

ROC - Curve

Reciever Operating Characteristics
[Name comes from old-school radio people]

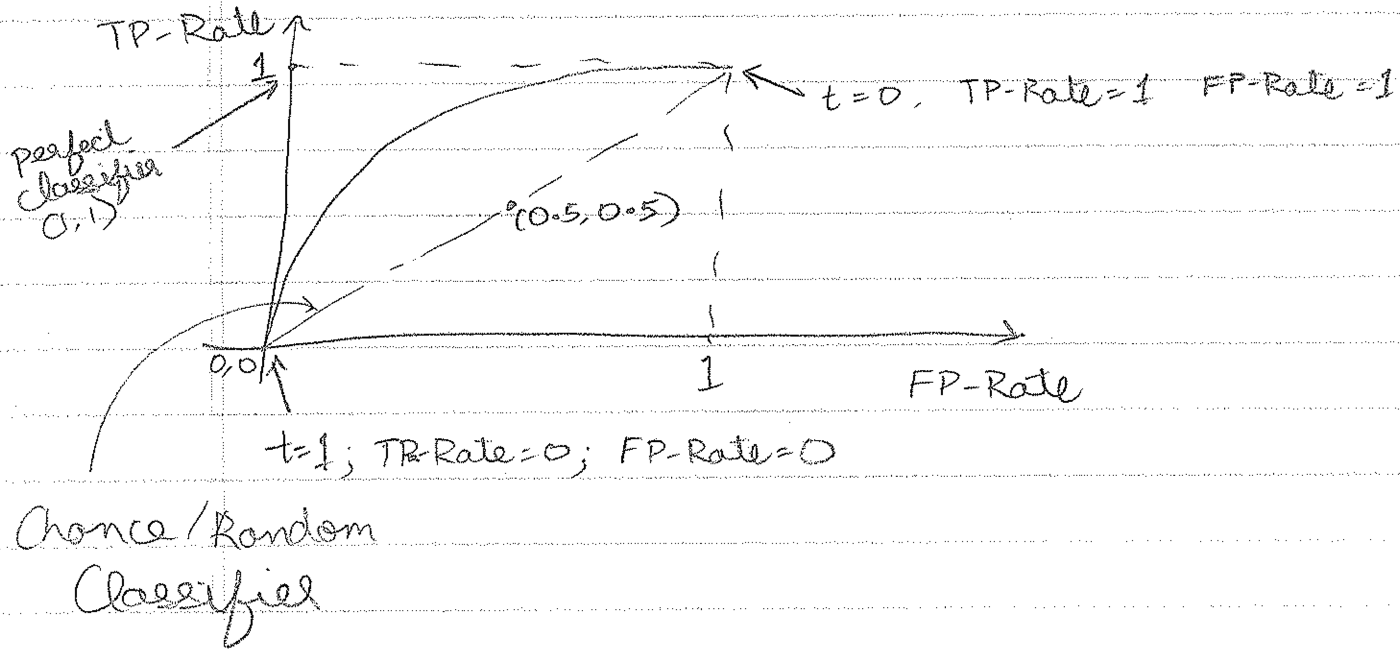

→ Side Note: There is a similar plot called
Precision - Recall Curve

where

$$\text{Precision} = \frac{\# TP}{\# \text{Pos-Predicted}}$$

$$= \frac{\# TP}{\# TP + \# FP}$$

} Out of cases we called cancer, how many actually had cancer?

② k-NN

Let $N_k(\vec{x}) = \{$ indices of k-NN of $\vec{x}$ in $D\}$

k-NN predictor

→ Regression $\hat{y} = g(x) = \frac{1}{K} \sum_{i \in N_k(x)} y_i$

Predict unweighted average of neighbours

→ Classification $\hat{y} = g(\vec{x}) = \underset{\text{class } c}{\text{argmax}} \; \underset{i \in N_k(\vec{x})}{\#}(y_i = c)$

$$= \underbrace{\underset{c}{\text{argmax}} \sum_{i \in N_k(\vec{x})} I(y_i = c)}_{\text{unweighted majority vote}}$$

③ Distances

→ Most common

Euclidean Distance / L_2-norm of difference

$\vec{x}, \vec{z} \in \mathbb{R}^d$

$$d(\vec{x}, \vec{z}) = \left[\sum_{i=1}^{d} (x_i - z_i)^2 \right]^{1/2}$$

→ Let's generalize this in 2 ways

① Mahalanobis

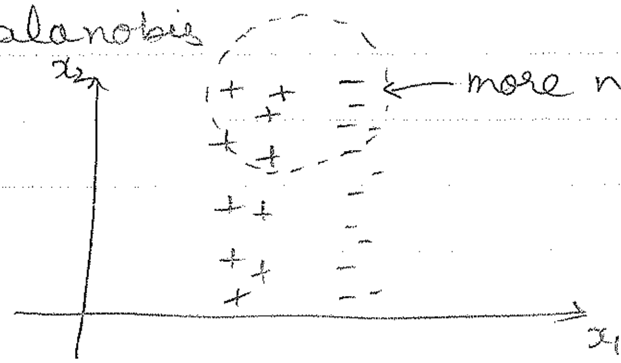

← more neg points than pos.

New definition

$$d^2(\vec{x}, \vec{z}) = 10(x_1 - z_1)^2 + (x_2 - z_2)^2$$

↑ deviations in dim 1 should be penalized more

In general $d^2(\vec{x}, \vec{z}) = \sum_{i=1}^{d} \sigma_i^2 (x_i - z_i)^2$

$$= (\vec{x} - \vec{z})^T \begin{bmatrix} \sigma_1^2 & & 0 \\ & \ddots & \\ 0 & & \sigma_d^2 \end{bmatrix} (\vec{x} - \vec{z})$$

More generally,

$$d^2(\vec{x}, \vec{z}) = (\vec{x} - \vec{z})^T A (\vec{x} - \vec{z})$$

Set $A = I_{d \times d}$ ⇒ Euc. dist

Note $A \succeq 0$

positive semi-definite

Definition: $A = A^T$ symetric
& $\vec{x}^T A \vec{x} \geq 0 \quad \forall \vec{x} \in \mathbb{R}^d$

$\longleftrightarrow$

Other generalization

Minkowski-distance / L_p-norm of difference

$$d(\vec{x}, \vec{z}) = \left[\sum_{i=1}^{d} |x_i - z_i|^p \right]^{1/p}$$

$p = 2 =$ Euc. dist

$p = 1 \equiv$ Manhattan distance

$$= \sum_{i=1}^{d} |x_i - z_i|$$

$p \to \infty \equiv$ Max-distance

$$= \max_i |x_i - z_i| \qquad 1 \leq i \leq d$$

Why? Simple proof.

$$\lim_{p \to \infty} \left[\sum_{i=1}^{d} |x_i - d_i|^p \right]^{1/p}$$

Let j = index of max-difference
$= \underset{i=1,\ldots,d}{\operatorname{argmax}} |x_i - z_i|$
[For simplicity, assume unique argmax]

$$= \lim_{p \to \infty} \left[|x_j - z_j|^p + \sum_{i \neq j} |x_i - z_i|^p \right]^{1/p}$$

$$= \lim_{p \to \infty} |x_j - z_j|^{p/p} \left[1 + \underbrace{\sum_{i \neq j} \left(\frac{|x_i - z_i|}{|x_j - z_j|} \right)^p}_{(<1)^p \to 0 \text{ as } p \to \infty} \right]^{1/p}$$

$$\left[1 + \sum_{i \neq j} (\to 0) \right]$$

$$= |x_j - z_j|$$

Similarly $p = 0$

$d(\vec{x}, \vec{z}) = \#\text{dims where } x_i \neq z_i$

Level Sets