## DATA MINING 1 Classification Model Evaluation

Dino Pedreschi, Riccardo Guidotti

Revisited slides from Lecture Notes for Chapter 3 "Introduction to Data Mining", 2nd Edition by
Tan, Steinbach, Karpatne, Kumar

## What is Classification?



Model Evaluation

## Model Evaluation

- Metrics for Performance Evaluation
- How to evaluate the performance of a model?
- Methods for Performance Evaluation
- How to obtain reliable estimates?
- Methods for Model Comparison
- How to compare the relative performance among competing models?


## Model Evaluation

- Metrics for Performance Evaluation
- How to evaluate the performance of a model?
- Methods for Performance Evaluation
- How to obtain reliable estimates?
- Methods for Model Comparison
- How to compare the relative performance among competing models?


## Problem Setting

- Let suppose we have a vector y of actual/real class labels, i.e.,
-y y [0001110101011100]
- Let name $y^{\prime}$ the vector returned by a trained model f, i.e.,
-y' = [0011100101110000]


## Metrics for Performance Evaluation

- Focus on the predictive capability of a model
- Rather than how fast it takes to classify or build models, scalability, etc.
- Confusion Matrix:

|  | PREDICTED CLASS |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL <br> CLASS |  | Class=Yes | Class=No |
|  | Class=Yes | a | b |
|  | Class=No | c | d |

$$
\begin{aligned}
& \mathrm{a}: \mathrm{TP} \text { (true positive) } \\
& \text { b: FN (false negative) } \\
& \mathrm{c}: \mathrm{FP} \text { (false positive) } \\
& \text { d: TN (true negative) }
\end{aligned}
$$

## Metrics for Performance Evaluation

$$
\begin{aligned}
& \text { •y }=[0001110101011100] \\
& \cdot y^{\prime}=\left[\begin{array}{lllllll}
0 & 1 & 110 & 0 & 101110000
\end{array}\right] \\
& \text { TN FP } \\
& \text { FN } \\
& \text { TP }
\end{aligned}
$$

## Metrics for Performance Evaluation...

|  | PREDICTED CLASS |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL <br> CLASS | Class=Yes | Class=No |  |
|  | Class=No | a <br> (TP) | b <br> (FN) |
|  | c <br> (FP) | d <br> (TN) |  |

Most widely-used metric:

$$
\text { Accuracy }=\frac{a+d}{a+b+c+d}=\frac{T P+T N}{T P+T N+F P+F N}
$$

## Limitation of Accuracy

- Consider a 2-class problem
- Number of Class 0 examples $=9990$
- Number of Class 1 examples $=10$
- If model predicts everything to be class 0 , accuracy is $9990 / 10000=99.9 \%$
- Accuracy is misleading because model does not detect any class 1 example


## Cost-Sensitive Measures

$$
\begin{aligned}
& \operatorname{Precision}(\mathrm{p})=\frac{T P}{T P+F P} \\
& \operatorname{Recall}(\mathrm{r})=\frac{T P}{T P+F N} \\
& \text { F-measure }(\mathrm{F})=\frac{2 r p}{r+p}=\frac{2 T P}{2 T P+F N+F P}
\end{aligned}
$$

(1. Precision is biased towards $C($ Yes $\mid$ Yes $) \& C(Y e s \mid N o)$
@ Recall is biased towards $\mathrm{C}(\mathrm{Yes} \mid \mathrm{Yes}) \& \mathrm{C}(\mathrm{No} \mid \mathrm{Yes})$
[ F -measure is biased towards all except $\mathrm{C}(\mathrm{No} \mid \mathrm{No})$

$$
\text { Weighted Accuracy }=\frac{w_{1} a+w_{4} d}{w_{1} a+w_{2} b+w_{3} c+w_{4} d}
$$

## Cost Matrix

|  | PREDICTED CLASS |  |  |
| :---: | :---: | :--- | :--- |
| ACTUAL <br> CLASS | Class=Yes | C(Yes\|Yes) | Class=Yes |
|  | Class=No\|Yes) |  |  |
|  | Class=No | C(Yes\|No) | C(No\|No) |

$\mathrm{C}(\mathrm{i} \mid \mathrm{j})$ : Cost of misclassifying class j example as class i

## Computing Cost of Classification

| Cost <br> Matrix | PREDICTED CLASS |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL <br> CLASS | $\mathrm{C}(\mathrm{i} \mid \mathrm{j})$ | + | - |
|  | $\boldsymbol{+}$ | -1 | 100 |
|  | - | 1 | 0 |


| ${\text { Model } M_{1}}^{\|c\|}$ PREDICTED CLASS |  |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL <br> CLASS | + | + | - |
|  | $\boldsymbol{+}$ | 150 | 40 |
|  | 60 | 250 |  |


| ${\text { Model } M_{2}}^{2}$ | PREDICTED CLASS |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL |  | + | - |
|  | + | 250 | 45 |
|  | - | 5 | 200 |

Accuracy $=80 \%$
Accuracy = 90\%
Cost $=3910$
Cost $=4255$

## Cost vs Accuracy

| Count | PREDICTED CLASS |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL <br> CLASS |  | Class=Yes | Class=No |
|  | Class=No | c | d |


| Cost | PREDICTED CLASS |  |  |
| :--- | :---: | :---: | :---: |
| ACTUAL <br> CLASS | Class=Yes | Class=Yes | Class=No |
|  | Class=No | q | p |

Accuracy is proportional to cost if

1. $\mathrm{C}(\mathrm{Yes} \mid \mathrm{No})=\mathrm{C}(\mathrm{No} \mid \mathrm{Yes})=\mathrm{q}$
2. $C($ Yes $\mid$ Yes $)=C(N o \mid N o)=p$

$$
N=a+b+c+d
$$

Accuracy $=(a+d) / \mathrm{N}$

$$
\begin{aligned}
\text { Cost } & =p(a+d)+q(b+c) \\
& =p(a+d)+q(N-a-d) \\
& =q N-(q-p)(a+d) \\
& =N[q-(q-p) \times \text { Accuracy }]
\end{aligned}
$$

## Binary vs Multiclass Evaluation

|  | PREDICTED CLASS |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL <br> CLASS |  | Class=Yes | Class=No |
|  | Class=Yes | TP | FN |
|  | Class=No | FP | TN |


|  | PREDICTED CLASS |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
| ACTUAL <br> CLASS |  | Class=A | Class=B | Class=C |
|  | Class=A | TP-A |  |  |
|  | Class=B |  | TP-B |  |
|  | Class=C |  |  | TP-C |

[^0]Accuracy $=\#$ correct $/ N=(T P-A+T P-B+T P-C) / N$

## Multiclass Evaluation

|  | PREDICTED CLASS |  |  |  |
| :---: | :--- | :---: | :---: | :---: |
| ACTUAL <br> CLASS | Class=A | TP-A | a | b |
|  | Class=B | c | TP-B | d |
|  | Class=C | e | f | TP-C |
|  |  |  |  |  |

$\operatorname{Precision}(\mathrm{p})=\frac{T P}{T P+F P}$
$\operatorname{Recall}(\mathrm{r})=\frac{T P}{T P+F N}$
F-measure $(\mathrm{F})=\frac{2 r p}{r+p}=\frac{2 T P}{2 T P+F N+F P}$

| A | PREDICTED CLASS |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL <br> CLASS | Class=A | Class=A | Class=NotA |
|  | Class=NotA | $c+e$ | TP-B $+\mathrm{TP}-\mathrm{C}$ <br> $+d+f$ |


| B | PREDICTED CLASS |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL <br> CLASS |  | Class=B | Class=Not B |
|  | Class=B | TP-B | $c+d$ |
|  | Class=Not B | $a+f$ | TP-A + TP-C <br> $+b+e$ |


| C | PREDICTED CLASS |  |  |
| :---: | :---: | :---: | :---: |
| ACTUAL <br> CLASS |  | Class=C | Class=Not C |
|  | Class=C | TP-C | $\mathrm{e}+\mathrm{f}$ |
|  | Class=Not C | $\mathrm{b}+\mathrm{d}$ | TP-A $+\mathrm{TP}-\mathrm{B}$ <br> $+\mathrm{a}+\mathrm{c}$ |

## Model Evaluation

- Metrics for Performance Evaluation
- How to evaluate the performance of a model?
- Methods for Performance Evaluation
- How to obtain reliable estimates?
- Methods for Model Comparison
- How to compare the relative performance among competing models?


## Methods for Evaluation



## Parameter Tuning

- It is important that the test data is not used in any way to create the classifier
- Some learning schemes operate in two stages:
- Stage 1: builds the basic structure
- Stage 2: optimizes parameter settings
- The test data can't be used for parameter tuning!
- Proper procedure uses three sets:
- training data,
- validation data,
- test data
- Validation data is used to optimize parameters
- Once evaluation is complete, all the data can be used to build the final classifier
- Generally, the larger the training data the better the classifier
- The larger the test data the more accurate the error estimate


## Methods for Performance Evaluation

- How to obtain a reliable estimate of performance?
- Performance of a model may depend on other factors besides the learning algorithm:
- Class distribution
- Cost of misclassification
- Size of training and test sets


## Learning Curve



- Learning curve shows how accuracy changes with varying sample size
- Requires a sampling schedule for creating learning curve:

Effect of small sample size:

- Bias in the estimate
- Variance of estimate

1. How much a classification model benefits from adding more training data?
2. Does the model suffer from a variance error or a bias error?

## Methods of Estimation

- Holdout
- Reserve 2/3 for training and 1/3 for testing
- Random subsampling
- Repeated holdout
- Cross validation
- Partition data into $k$ disjoint subsets
- k-fold: train on $k-1$ partitions, test on the remaining one
- Leave-one-out: k=n
- Stratified sampling
- oversampling vs undersampling
- Bootstrap
- Sampling with replacement


## Holdout

- The holdout method reserves a certain amount for testing and uses the remainder for training
- Usually, one third for testing, the rest for training.
- Typical quantities are 60\%-40\%, 66\%-34\%, 70\%-30\%.
- For small or "unbalanced" datasets, samples might not be representative
- For instance, few or none instances of some classes
- Stratified sample
- Balancing the data
- Make sure that each class is represented with approximately equal proportions in both subsets


## Repeated Holdout

- Holdout estimate can be made more reliable by repeating the process with different subsamples
- In each iteration, a certain proportion is randomly selected for training (possibly with stratification)
- The error rates on the different iterations are averaged to yield an overall error rate
- This is called the repeated holdout method
- Still not optimum: the different test sets overlap


## Cross Vaildation



Run 2
Training Set

- Avoids overlapping test sets
- First step: data is split into k subsets of equal size
- Second step: each subset in turn is used for testing and the remainder for training
- This is called k-fold cross-validation
- Often the subsets are stratified before cross-validation is performed
- The error estimates are averaged to yield an overall error estimate
- Even better: repeated stratified cross-validation E.g. ten-fold cross-validation is repeated ten times and results are averaged (reduces the variance)


## Data Partitioning



Train the model for final testing


Train the model for parameter selection


## Test

- Test the model
- Compare different models once parameters have been selected



## Test

Cross Validation (check potential dataset bias)

## Evaluation: Training, Validation, Tests



## Cross Validation with Time



## Model Evaluation

- Metrics for Performance Evaluation
- How to evaluate the performance of a model?
- Methods for Performance Evaluation
- How to obtain reliable estimates?
- Methods for Model Comparison
- How to compare the relative performance among competing models?


## ROC (Receiver Operating Characteristic)

- Developed in 1950s for signal detection theory to analyze noisy signals
- Characterize the trade-off between positive hits and false alarms
- ROC curve plots TP (on the $y$-axis) against FP (on the $x$-axis)
- Performance of each classifier represented as a point on the ROC curve
- changing the threshold of algorithm, sample distribution or cost matrix changes the location of the point


## Receiver Operating Characteristic Curve

- It illustrates the ability of a binary classifier as its discrimination threshold THR is varied.
- The ROC curve is created by plotting the true positive rate (TPR) against the false positive rate (FPR) at various THR.
- The TPR = TP / (TP + FN) is also known as sensitivity, recall or probability of detection.

- The FPR = FP / (TN + FP) is also known as probability of false alarm and can be calculated as ( 1 - specificity).


## ROC Curve

(TP,FP):

- $(0,0)$ : declare everything to be negative class
- $(1,1)$ : declare everything to be positive class
- $(0,1)$ : ideal
- Diagonal line:
- Random guessing
- Below diagonal line:
- prediction is opposite of the true class



## Using ROC for Model Comparison


$\square$ No model consistently outperform the other

- $M_{1}$ is better for small FPR
- $M_{2}$ is better for large FPR
$\square$ Area Under the ROC curve
- Ideal: Area = 1
- Random: Area $=0.5$


## How to Construct the ROC curve

| Instance | $\mathrm{P}(+\mid \mathrm{A})$ | True Class |
| :---: | :---: | :---: |
| 1 | 0.95 | + |
| 2 | 0.93 | + |
| 3 | 0.87 | - |
| 4 | 0.85 | - |
| 5 | 0.85 | - |
| 6 | 0.85 | + |
| 7 | 0.76 | - |
| 8 | 0.53 | + |
| 9 | 0.43 | - |
| 10 | 0.25 | + |

- Use classifier that produces posterior probability for each test instance $P(+\mid A)$
- Sort the instances according to $\mathrm{P}(+\mid \mathrm{A})$ in decreasing order
- Apply threshold at each unique value of $P(+\mid A)$
- Count the number of TP, FP, TN, FN at each threshold
- TP rate, $\mathrm{TPR}=\mathrm{TP} /(\mathrm{TP}+\mathrm{FN})$
- FP rate, $\mathrm{FPR}=\mathrm{FP} /(\mathrm{FP}+\mathrm{TN})$


## How to Construct the ROC curve



## How to Construct the ROC curve



## How to Construct the ROC curve

|  |  |  |  |  |  |  |  |  |  |  |  |  | Inst. | $\mathrm{P}(+\mid$ ) | True |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Class | + | - | + | - | - | - | + | - | + | + |  | 1 | 0.95 | + |
| Thresh | ld >= | 0.25 | 0.43 | 0.53 | 0.76 | 0.85 | 0.85 | 0.85 | 0.87 | 0.93 | 0.95 | 1.00 | 2 | 0.93 | + |
|  | TP |  |  |  |  |  |  |  |  |  |  |  | 3 | 0.87 | - |
|  | 7 | 5 | 4 |  |  |  |  |  |  |  |  |  | 4 | 0.85 | - |
|  | FP | 5 | 5 |  |  |  |  |  |  |  |  |  | 5 | 0.85 | - |
|  | TN | 0 | 0 |  |  |  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |  |  |  | 6 | 0.85 | + |
|  | FN | 0 | 1 |  |  |  |  |  |  |  |  |  | 7 | 0.76 | - |
| $\longrightarrow$ | TPR | 1 | 0.8 |  |  |  |  |  |  |  |  |  | 8 | 0.53 | + |
| $\longrightarrow$ | FPR | 1 | 1 |  |  |  |  |  |  |  |  |  | 9 | 0.43 | - |
|  |  |  |  |  |  |  |  |  |  |  |  |  | 10 | 0.25 | + |

## How to Construct the ROC curve

|  |  |  |  |  |  |  |  |  |  |  |  |  | Inst. | $\mathrm{P}(+\mid A)$ | True |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Class | + | - | + | - | - | - | + | - | + | + |  | 1 | 0.95 | + |
| Thres | ld >= | 0.25 | 0.43 | 0.53 | 0.76 | 0.85 | 0.85 | 0.85 | 0.87 | 0.93 | 0.95 | 1.00 | 2 | 0.93 | + |
|  |  |  |  |  |  |  |  |  |  |  |  |  | 3 | 0.87 | - |
|  | TP | 5 | 4 | 4 |  |  |  |  |  |  |  |  | 4 | 0.85 | - |
|  | FP | 5 | 5 | 4 |  |  |  |  |  |  |  |  | 5 | 0.85 |  |
|  | TN | 0 | 0 | 1 |  |  |  |  |  |  |  |  | 6 |  |  |
|  | FN | 0 | 1 | 1 |  |  |  |  |  |  |  |  |  |  |  |
| $\longrightarrow$ | TPR | 1 | 0.8 | 0.8 |  |  |  |  |  |  |  |  | 8 | 0.53 | + |
| $\longrightarrow$ | FPR | 1 | 1 | 0.8 |  |  |  |  |  |  |  |  | 9 | 0.43 | - |
|  |  |  |  |  |  |  |  |  |  |  |  |  | 10 | 0.25 | + |

## How to Construct the ROC curve



## How to Construct the ROC curve



## How to Construct the ROC curve



## How to Construct the ROC curve



## How to Construct the ROC curve



## How to Construct the ROC curve



## How to Construct the ROC curve



## How to Construct the ROC curve

|  | Class | + | - | + | - | - | - | + | - | + | + |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Threshold >= |  | 0.25 | 0.43 | 0.53 | 0.76 | 0.85 | 0.85 | 0.85 | 0.87 | 0.93 | 0.95 | 1.00 |
|  | TP | 5 | 4 | 4 | 3 | 3 | 3 | 3 | 2 | 2 | 1 | 0 |
|  | FP | 5 | 5 | 4 | 4 | 3 | 2 | 1 | 1 | 0 | 0 | 0 |
|  | TN | 0 | 0 | 1 | 1 | 2 | 3 | 4 | 4 | 5 | 5 | 5 |
|  | FN | 0 | 1 | 1 | 2 | 2 | 2 | 2 | 3 | 3 | 4 | 5 |
|  | TPR | 1 | 0.8 | 0.8 | 0.6 | 0.6 | 0.6 | 0.6 | 0.4 | 0.4 | 0.2 | 0 |
| $\longrightarrow$ | FPR | 1 | 1 | 0.8 | 0.8 | 0.6 | 0.4 | 0.2 | 0.2 | 0 | 0 | 0 |

$T P R=T P /(T P+F N)$
$F P R=F P /(T N+F P)$

ROC Curve:


| Inst. | $\mathrm{P}(+\mid \mathrm{A})$ | True <br> Class |
| :---: | :---: | :---: |
| 1 | 0.95 | + |
| 2 | 0.93 | + |
| 3 | 0.87 | - |
| 4 | 0.85 | - |
| 5 | 0.85 | - |
| 6 | 0.85 | + |
| 7 | 0.76 | - |
| 8 | 0.53 | + |
| 9 | 0.43 | - |
| 10 | 0.25 | + |

- The lift curve is a popular technique in direct marketing.
- The input is a dataset that has been "scored" by appending to each case the estimated probability that it will belong to a given class.
- The cumulative lift chart (also called gains chart) is constructed with the cumulative number of cases (descending order of probability) on the $x$-axis and the cumulative number of true positives on the $y$-axis.
- The dashed line is a reference line. For any given number of cases (the $x$-axis value), it represents the expected number of positives we would predict if we did not have a model but simply selected cases at random. It provides a benchmark against which we can see performance of the model.

Notice: "Lift chart" is a rather general term, often used to identify also other kinds of plots. Don't get confused!

## Lift Chart-Example

| Serial no. | Predicted prob of 1 | Actual Class | Cumulative Actual class |
| ---: | ---: | ---: | ---: |
| 1 | 0.995976726 | 1 | 1 |
| 2 | 0.987533139 | 1 | 2 |
| 3 | 0.984456382 | 1 | 3 |
| 4 | 0.980439587 | 1 | 4 |
| 5 | 0.948110638 | 1 | 5 |
| 6 | 0.889297203 | 1 | 6 |
| 7 | 0.847631864 | 1 | 7 |
| 8 | 0.762806287 | 0 | 7 |
| 9 | 0.706991915 | 1 | 8 |
| 10 | 0.680754087 | 1 | 9 |
| 11 | 0.656343749 | 1 | 10 |
| 12 | 0.622419543 | 0 | 10 |
| 13 | 0.505506928 | 1 | 11 |
| 14 | 0.47134045 | 0 | 11 |
| 15 | 0.337117362 | 0 | 11 |
| 16 | 0.21796781 | 1 | 12 |
| 17 | 0.199240432 | 0 | 12 |
| 18 | 0.149482655 | 0 | 12 |
| 19 | 0.047962588 | 0 | 12 |
| 20 | 0.038341401 | 0 | 12 |
| 21 | 0.024850999 | 0 | 12 |
| 22 | 0.021806029 | 0 | 12 |
| 23 | 0.016129906 | 0 | 12 |
| 24 | 0.003559986 | 0 | 12 |



## Lift Chart - Application Example

- From Lift chart we can easily derive an "economical value" plot, e.g. in target marketing.
- Given our predictive model, how many customers should we target to maximize income?
- Profit $=$ UnitB*MaxR*Lift(X) - UnitCost*N*X/100
- UnitB = unit benefit, UnitCost = unit postal cost
- $\mathrm{N}=$ total customers
- MaxR = expected potential respondents in all population ( N )
- $\operatorname{Lift}(X)=$ lift chart value for $X$, in $[0, . ., 1]$


## Lift Chart-Application Example



## References

- Chapter 3. Classification: Basic Concepts and Techniques.



[^0]:    Accuracy $=T P+T N /(T P+T N+F N+F P)=\#$ correct $/ N$

