Classification and Prediction

Introduction
Evaluation of Classifiers
Decision Trees
Bayesian Classification
Nearest-Neighbor Classification
Support Vector Machines
Multi-relational Classification
Regression Analysis

## Introduction

## The Classification Problem

- Let $O$ be a set of objects of the form $\left(o_{1}, \ldots, o_{d}\right)$
with attributes $A_{i}, 1 \leq i \leq d$, and class membership $c_{i}, c_{i} \in C=\left\{c_{1}, \ldots, c_{k}\right\}$
- Wanted:
class membership for objects from $D \backslash O$
a classifier $K: D \rightarrow C$
- Difference to clustering
classification: set of classes $C$ known apriori clustering: classes are output
- Related problem: prediction
predict the value of a numerical attribute

| NAME | RANK | YEARS | TENURED |
| :--- | :--- | :---: | :---: |
| Mike | Assistant Prof | 3 | no |
| Mary | Assistant Prof | 7 | yes |
| Bill | Professor | 2 | yes |
| Jim | Associate Prof | 7 | yes |
| Dave | Assistant Prof | 6 | no |
| Anne | Associate Prof | 3 | no |

## Evaluation of Classifiers

## Introduction

- Given a sample of labeled data $(O)$
- Want to build a classifier that labels the entire population in particular, $D \backslash O$
- Can only estimate the performance of the classifier on unseen data
- Need separate, disjoint training and test data (all labeled)
- Training data
for training the classifier (model construction)
- Test data
to evaluate the trained classifier


## Evaluation of Classifiers

## Approaches

- Train-and-Test
- partition set $O$ into two (disjoint) subsets: Training data and Test data
- not recommended for small $O$
- $m$-fold cross validation
- partition set $O$ into $m$ same size subsets
- train $m$ different classifiers using a different one of these $m$ subsets as test data and the other subsets for training
- average the evaluation results of the $m$ classifiers
- appropriate also for small $O$


## Evaluation Criteria

- Classification accuracy
- Interpretability
e.g. size of a decision tree insight gained by the user
- Efficiency
of model construction
of model application
- Scalability for large datasets
for secondary storage data
- Robustness
w.r.t. noise and unknown attribute values

Classification as optimization problem: score of a classifier

## Evaluation of Classifiers

## Classification Accuracy

- Let $K$ be a classifier, $T R \subseteq O$ the training data, $T E \subseteq O$ the test data. $C(o)$ : actual class of object $o$.
- classification accuracy of $K$ on $T E$ :
- classification error $^{\text {Accuracy }_{T E}(K)=\frac{|\{o \in T E \mid K(o)=C(o)\}|}{|T E|}}$

$$
\operatorname{Error}_{T E}(K)=\frac{|\{o \in T E \mid K(o) \neq C(o)\}|}{|T E|}
$$

aggregates over all classes $c_{i} \in C$
not appropriate if minority class is most important

## Evaluation of Classifiers

## Confusion Matrix

- Let $c_{1} \in C$ be the target (positive) class, the union of all other classes the contrasting (negative) class.
- Comparing the predicted and the actual class labels, we can distinguish four different cases:

|  | Predicted as positive | Predicted as negative |
| :--- | :--- | :--- |
| Actually positive | True Positive (TP) | False Negative (FN) |
| Actually negative | False Positive (FP) | True Negative (TN) |
|  |  |  |

Confusion matrix

## Evaluation of Classifiers

## Precision and Recall

- We define the following two measures of $K$ w.r.t. the given target class:

$$
\begin{aligned}
& \operatorname{Precision}(K)=\frac{|T P|}{|T P|+|F P|} \\
& \operatorname{Recall}(K)=\frac{|T P|}{|T P|++F N}
\end{aligned}
$$

- There is a trade-off between precision and recall.
- Therefore, we also define a measure combining precision and recall:

$$
\mathrm{F}-\operatorname{Measure}(K)=\frac{2 \cdot \operatorname{Precision}(K) \cdot \operatorname{Recall}(K)}{\operatorname{Precision}(K)+\operatorname{Recall}(K)}
$$

## Evaluation of Classifiers

## ROC Curves

- F-Measure captures only one of the possible trade-offs between precision and recall (or between TP and FP)
- True positive rate: percentage of positive data correctly predicted
- False positive rate: percentage of negative data falsely predicted as positive


8

## Evaluation of Classifiers

## Model Selection

- Given two classifiers and their (estimated!) classification accuracies
e.g., obtained from $m$-fold cross-validation
- Which of the classifiers is really better?
- Naive approach: just take the one with higher mean classification accuracy
- But: classification accuracy may vary greatly among the $m$ folds
- Differences in classification accuracies may be insignificant due only to chance


## Evaluation of Classifiers

## Model Selection

- We measure the classification error on a (small) test dataset $O \subseteq X$.
- Questions:

How to estimate the true classification error on the whole instance space $X$ ? How does the deviation from the observed classification error depend on the size of the test set?

- Random experiment to determine the classification error on test set (of size $n$ ): repeat $n$ times
(1) draw random object from $X$
(2) compare predicted vs. actual class label for this object
- Classification error is percentage of misclassified objects
$\rightarrow$ observed classification error follows a Binomial distribution with mean $=$ true classification error (unknown)


## Evaluation of Classifiers

## Binomial distribution

- $n$ repeated tosses of a coin with unknown probability $p$ of head head $=$ misclassified object
- Record the number $r$ of heads (misclassifications)
- Binomial distribution defines probability for all possible values of $r$ :

$$
P(r)=\frac{n!}{r!(n-r)!} p^{r}(1-p)^{n-r}
$$

- Random variable $Y$ counting the number of heads in $n$ coin tosses:

$$
\begin{aligned}
& E[Y]=n \cdot p \text { expected value } \\
& \operatorname{Var}[Y]=n p(1-p) \\
& \sigma_{Y}=\sqrt{n p(1-p)}
\end{aligned}
$$

## Evaluation of Classifiers

## Estimating the True Classification Error

- We want to estimate the unknown true classification error $(p)$.
- Estimator for $p: \quad E[Y]=n \cdot p=r \Rightarrow p=\frac{r}{n}$
- We want also confidence intervals for our estimate. $n$
- Standard deviation for the true classification error (Y/n):

$$
\begin{aligned}
\sigma_{\frac{r}{n}}=\frac{\sigma_{Y}}{n} & =\frac{\sqrt{n p(1-p)}}{n} \\
\sigma_{\frac{r}{n}} & \approx \sqrt{\frac{\frac{r}{n}\left(1-\frac{r}{n}\right)}{n}} \quad \text { use } \frac{r}{n} \text { as estimator for } p
\end{aligned}
$$

## Evaluation of Classifiers

## Estimating the True Classification Error

- For sufficiently large values of $n$, the Binomial distribution can be approximated by a Normal distribution with the same mean and standard deviation.
- Random variable $Y$ Normal distributed with mean $m$ and standard deviation $s$ and $y$ be the observed value of $Y$ :
the mean of $Y$ falls into the following interval with a probability of $N \%$

$$
y \pm z_{N} \sigma
$$

- In our context, $N \%$ confidence interval for the true classification error:

$$
\frac{r}{n} \pm z_{N} \sqrt{\frac{\frac{r}{n}\left(1-\frac{r}{n}\right)}{n}} \text { interval size decreases with increasing } n
$$

## Decision Trees

## Introduction

| ID | Age | Autotype | Risk |
| :---: | :---: | :---: | :---: |
| 1 | 23 | Family | high |
| 2 | 17 | Sports | high |
| 3 | 43 | Sports | high |
| 4 | 68 | Family | low |
| 5 | 32 | Truck | low |


disjunction of conjunction of attribute constraints and hierarchical structure

## Introduction

- A decision tree is a tree with the following properties:
- An inner node represents an attribute.
- An edge represents a test on the attribute of the father node.
- A leaf represents one of the classes of $C$.
- Construction of a decision tree

Based on the training data
Top-Down strategy

- Application of a decision tree

Traversal of the decision tree from the root to one of the leaves
Unique path
Assignment of the object to class of the resulting leaf

## Construction of Decision Trees

## Base algorithm

- Initially, all training data records belong to the root.
- Next attribute is selected and split (split strategy).
- Training data records are partitioned according to the chosen split.
- Method is applied recursively to each partition.
local optimization method (greedy)


## Termination conditions

- No more split attributes.
- All (most) training data records of the node belong to the same class.


## Decision Trees

## Example

| Day | Outlook | Temperature | Humidity | Wind | PlayTennis? |
| :---: | :--- | :---: | :---: | :---: | :---: |
| 1 | sunny | hot | high | weak | no |
| 2 sunny | hot | high | strong | no |  |
| 3 | overcast | hot | high | weak | yes |
| 4 | rainy | mild | high | weak | yes |
| 5 | rainy | cool | normal | weak | yes |
| 6 | rainy | cool | normal | strong | no |
| 7 | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |
|  |  |  |  |  |  |

## Decision Trees

## Example



## Types of Splits

## Categorical attributes

- Conditions of the form „attribute =a" or ,,attribute $\in$ set"
- Many possible subsets


Numerical attributes

- Conditions of the form ,,attribute $<a$ "
- Many possible split points



## Quality Measures for Splits

## Given

- a set $T$ of training data
- a disjoint, exhaustive partitioning $T_{1}, T_{2}, \ldots, T_{m}$ of $T$
- $p_{i}$ the relative frequency of class $c_{i}$ in $T$

Wanted

- A measure of the impurity of set $S$ (of training data) w.r.t. class labels
- A split of $T$ in $T_{1}, T_{2}, \ldots, T_{\mathrm{m}}$ minimizing this impurity measure
information gain, gini-index


## Decision Trees

## Information Gain

- Entropy: minimal number of bits to encode a message to transmit the class of a random training data record
- Entropy for a set $T$ of training data:

$$
\begin{aligned}
& \quad \begin{array}{r}
\operatorname{entropy}(T)
\end{array}(T)=-\sum_{i=1}^{k} p_{i} \cdot \log _{2} p_{i} \\
& \text { entropy }(T)=0, \text { if } p_{i}=1 \text { for some } i
\end{aligned}
$$

- Let attribute $A$ produce the partitioning $T_{1}, T_{2}, \ldots, T_{m}$ of $T$.
- The information gain of attribute $A$ w.r.t $T$ is defined as

$$
\operatorname{InformationGain}(T, A)=\operatorname{entropy}(T)-\sum_{i=1}^{m} \frac{\left|T_{i}\right|}{|T|} \cdot \operatorname{entropy}\left(T_{i}\right)
$$

## Decision Trees

## Gini-Index

- Gini index for a set $T$ of training data records

$$
\operatorname{gini}(T)=1-\sum_{j=1}^{k} p_{j}^{2}
$$

low gini index $\Leftrightarrow$ low impurity,
high gini index $\Leftrightarrow$ high impurity

- Let attribute $A$ produce the partitioning $T_{1}, T_{2}, \ldots, T_{m}$ of $T$.
- Gini index of attribute $A$ w.r.t. $T$ is defined as

$$
\operatorname{gini}_{A}(T)=\sum_{i=1}^{m} \frac{\left|T_{i}\right|}{|T|} \cdot \operatorname{gini}\left(T_{i}\right)
$$

## Decision Trees

## Example


$\operatorname{InformationGain}(T$, Humidity $)=0.94-\frac{7}{14} \cdot 0.985-\frac{7}{14} \cdot 0.592=0.151$


## Overfitting

Overfitting: there are two decision trees $T$ and $T^{\prime}$ with

- $T$ has a lower error rate than $T^{\prime}$ on the training data, but
- $T$ ' has a lower test error rate than $T$.



## Approaches for Avoiding Overfitting

- Removal of erroneous training data
in particular, inconsistent training data
- Choice of appopriate size of training data set
not too small, not too large
- Choice of appropriate minimum support
minimum support:
minimum number of training data records belonging to a leaf node

```
minimum support >> 1
```


## Approaches for Avoiding Overfitting

- Choice of appropriate minimum confidence
minimum confidence: minimum percentage of the majority class of a leaf node
minimum confidence $\ll 100 \%$
leaves can also absorb noisy / erroneous training data records
- Subsequent pruning of the decision tree


## remove overfitting branches

## see next section



## Error Reduction-Pruning [Mitchell 1997]

- Train-and-Test paradigm
- Construction of decision tree $T$ for training data set $T R$.
- Pruning of $T$ using test data set $T E$
- Determine subtree of $T$ such that its removal leads to the maximum reduction of the classification error on $T E$.
- Remove this subtree.
- Stop, if no more such subtree.
only applicable if enough labled data available



## Minimal Cost Complexity Pruning

[Breiman, Friedman, Olshen \& Stone 1984]

- Cross-Validation paradigm

Applicable even if only small number of labled data available

- Pruning of decision tree using training data set

Cannot use classification error as quality measure

- Novel quality measure for decision trees

Trade-off between (observed) classification error and tree size
Weighted sum of classification error and tree size

Small decision trees tend to generalize better to unseen data


## Decision Trees

## Notions

- Size $|T|$ of decision tree $T$ : number of leaves
- Cost complexity of $T$ w.r.t. training data set $T R$ and complexity parameter $\alpha \geq 0$ :

$$
C C_{T R}(T, \alpha)=\operatorname{error}_{T R}(T)+\alpha \cdot|T|
$$

- The smallest minimizing subtree $T(\alpha)$ of $T$ w.r.t. $\alpha$ has the following properties :
(1) There is no subtree of $T$ with smaller cost complexity.
(2) If $T(\alpha)$ and $T^{\prime}$ satisfy condition (1), then $T(\alpha)$ is a subtree of $T^{\prime}$.
- $\alpha=0: T(\alpha)=T$
- $\alpha=\infty: T(\alpha)=$ root of $T$
- $0<\alpha<\infty: T(\alpha)=$ true subtree of $T$ (more than the root)


## Decision Trees

## Notions

- $T_{\mathrm{e}}$ : subtree of $T$ with root $e,\{e\}$ : tree consisting only of node $e$ $T>T^{6}$ : subtree relationship
- For small values of $\alpha: C C_{\mathrm{TR}}\left(T_{\mathrm{e}}, \alpha\right)<C C_{\mathrm{TR}}(\{e\}, \alpha)$, for large values of $\alpha: C C_{\mathrm{TR}}\left(T_{\mathrm{e}}, \alpha\right)>C C_{\mathrm{TR}}(\{e\}, \alpha)$.
- critical value of $\alpha$ w.r.t. $e$

$$
\alpha_{c r i t}: C C_{\mathrm{TR}}\left(T_{\mathrm{e}}, \alpha_{\text {crit }}\right)=C C_{\mathrm{TR}}\left(\{e\}, \alpha_{\text {crit }}\right)
$$

for $\alpha \geq \alpha_{\text {crit }}$ the subtree of node $e$ should be pruned

- weakest link: node with minimal $\alpha_{\text {crit }}$ value


## Method

- Start with complete decision tree $T$.
- Iteratively, each time remove the weakest link from the current tree.
- If several weakest links: remove all of them in the same step. sequence of pruned trees $T\left(\alpha_{1}\right)>T\left(\alpha_{2}\right)>\ldots>T\left(\alpha_{m}\right)$
$\Longrightarrow$ with $\alpha_{1}<\alpha_{2}<\ldots<\alpha_{m}$
- Selection of the best $T\left(\alpha_{i}\right)$
estimate the true classification error of all $T\left(\alpha_{1}\right), T\left(\alpha_{2}\right), \ldots, T\left(\alpha_{m}\right)$
performing $l$-fold cross-validation on the training data set


## Decision Trees

## Example

| i | $\mathbf{T i} \mid$ | training error | estimated error | true error |
| ---: | ---: | :---: | :---: | :---: |
| 1 | 71 | 0,0 | 0,46 | 0,42 |
| 2 | 63 | 0,0 | 0,45 | 0,40 |
| 3 | 58 | 0,04 | 0,43 | 0,39 |
| 4 | 40 | 0,10 | 0,38 | 0,32 |
| 5 | 34 | 0,12 | 0,38 | 0,32 |
| 6 | 19 | 0,2 | 0,32 | 0,31 |
| 7 | 10 | 0,29 | 0,31 | 0,30 |
| 8 | 9 | 0,32 | 0,39 | 0,34 |
| 9 | 7 | 0,41 | 0,47 | 0,47 |
| 10 | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |

$T_{7}$ has the lowest estimated error and the lowest true error


## Bayesian Classification

## Introduction

- When building a probabilistic classifier, we would like to find the classifier (hypothesis) $h$ that has the maximum conditional probability given the observed data, i.e.

$$
\max _{h \in H} P(h \mid D)
$$

- But how to compute these conditional probabilities for all possible classifiers $h$ ?
- Bayes theorem


$$
\begin{aligned}
& P(h \mid D)=\frac{P(D \mid h) \cdot P(h)}{P(D)} \text { and } \\
& \max _{h \in H} P(h \mid D)=\max _{h \in H} P(D \mid h) \cdot P(h)
\end{aligned}
$$

## Bayesian Classification

## Introduction

$$
\max _{h \in H} P(h \mid D)=\max _{h \in H} P(D \mid h) \cdot P(h)
$$

$P(h \mid D)$ : posterior probability of $h$ given the data $D$ $P(D \mid h)$ : likelihood of the data $D$ given hypothesis $h$ $P(h)$ : prior probability of $h$

- The more training data $D$ we have, the higher becomes the influence of $P(D \backslash h)$.
- $P(h)$ is subjective.
- $P(h)$ can, e.g., favor simpler over more complex hypotheses.
- If there is no prior knowledge, i.e. $P(h)$ uniformly distributed, then we obtain the Maximum Likelihood Classifier as a special case.


## Bayesian Classification

## Introduction

－When applying a learned hypothesis $h$ to classify an object $o$ ， we could use the following decision rule：

$$
\operatorname{argmax} P\left(c_{j} \mid h\right)
$$

－$h$ depends on the attribute values of $o$ ，i．e．$o_{1}, \ldots c_{j} \mathscr{F}_{d}$ ．
－Therefore we determine

$$
\underset{\operatorname{argmax}}{\operatorname{argmax}} P\left(c_{j} \mid o_{1}, o_{d}\right)
$$

－Applying Bayes theorem，we ogbtain

$$
\begin{aligned}
& \underset{c_{j} \in C}{\operatorname{argmax}} P\left(c_{j} \mid o_{1}, \text { 団, }, o_{d}\right)=\underset{c_{j} \in C}{\operatorname{argmax}} \frac{P\left(o_{1}, \text {, }, o_{d} \mid c_{j}\right) \cdot P\left(c_{j}\right)}{P\left(o_{1}, \text { 國 }, o_{d}\right)} \\
& =\underset{c_{j} \in C}{\operatorname{argmax}} P\left(o_{1}, \mathrm{~K}_{\mathrm{k}}, o_{d} \mid c_{j}\right) \cdot P\left(c_{j}\right) \quad \text { Bayesian Classifier }
\end{aligned}
$$

## Bayesian Classification

## Naive Bayes Classifier

- Estimate the $P\left(c_{j}\right)$ using the observed frequencies of the individual classes.
- How to estimate the $P\left(o_{1}, \ldots, o_{d} \mid c_{j}\right)$ ?
- Assumption:
- Attribute values $o_{i}$ are conditionally independent, given class $c_{j}$
- $P\left(o_{i} \mid c_{j}\right)$ are easier to estimate from the training data than

$$
P\left(o_{1}, \ldots, o_{d} \mid c_{j}\right)
$$

§ $\sum_{i=1}^{d}\left|A_{i}\right|$ instead of $\prod_{i=1}^{d}\left|A_{i}\right|$ parameters to estimate

- Decision rule of the Naive Bayes-Classifier

$$
\underset{c_{j} \in C}{\operatorname{argmax}} P\left(c_{j}\right) \cdot \prod_{i=1}^{d} P\left(o_{i} \mid c_{j}\right)
$$

## Bayesian Classification

## Bayesian Networks

- Naive Bayes-Classifier is very efficient, but assumptions may be unrealistic $\Longrightarrow$ suboptimal classification accuracy
- Often, only some attributes are dependent, most are independent (given some class)
- Bayesian networks (Bayesian belief networks / Bayes nets)
allow you to specify all variable dependencies,
all other variables are assumed to be conditionally independent
- Network respresents subjective, a-priori beliefs


## Bayesian Classification

## Bayesian Networks

- Graph with nodes $=$ random variable (attribute) and

$$
\text { edge }=\text { conditional dependency }
$$

- Each random variable is (for given values of the predecessor variables) conditionally independent from all variables that are no successors.
- For each node (random variable): Table of conditional probabilities given values of the predecessor variables

Bayesian network can represent causal knowledge

## Example


$(\mathrm{FH}, \sim \mathrm{S}) \quad(\sim \mathrm{FH}, \sim \mathrm{S})$

|  | (FH,S) | $(\sim \mathbf{F H}, \mathbf{S})$ |  |  |
| :---: | :---: | :---: | :---: | :---: |
| LC | 0.8 | 0.5 | 0.7 | 0.1 |
| $\sim \mathrm{LC}$ | 0.2 | 0.5 | 0.3 | 0.9 |

Conditional probabilities for LungCancer

For given values of FamilyHistory and Smoker, the value of Emhysema does not provide any additional information about LungCancer

## Bayesian Classification

## Training Bayesian Networks

- With given network structure and fully observable random variables all attribute values of the training examples known estimate conditional probability tables by calculating the relative frequencies
- With given network structure and partially known random variables some attribute values of the training examples unknown expectation maximization (EM) algorithm random initialization of the unknown attribute values
- With apriori unknown network structure (very difficult!) assume fully observable random variables heuristic scoring functions for alternative network structures


## Bayesian Classification

## Interpretation of Raster Images

- Automatical interpretation of $d$ raster images of a given region for each pixel: a $d$-dimensional vector of grey values $\left(o_{1}, \ldots, o_{d}\right)$
- Assumption: different kinds of landuse exhibit characteristic behaviors of reflection / emission
(12),(17.5)



## Bayesian Classification

## Interpretation of Raster Images

- Application of the (optimal) Bayes classifier
- Estimate the $P\left(o_{1}, \ldots, o_{d} \mid c_{j}\right)$ without assuming conditional indepency
- Assume a $d$-dimensional Normal distribution of the grey value vectors of a given class



## Bayesian Classification

## Method

- Estimate from the training data
$\mu_{i}$ : d-dimensional mean vector of all feature vektors of class $c_{i}$
$\Sigma_{i}: d \cdot d$ covariance matrix of class $c_{i}$
- Problems of the decision rule
- Likelihood for the chosen cla very small
- Likelihood for several classes similar

unclassified regions


## Bayesian Classification

## Discussion

+ Optimality property
Standard for comparison with other classifiers
+ High classification accuracy in many applications
+ Incrementality
classifier can easily be adapted to new training objects
+ Integration of domain knowledge
- Applicability
the conditional probabilities may not be available
- Maybe inefficient

For high numbers of features
in particular, Bayesian networks

## Nearest-Neighbor Classification

## Motivation

- Optimal Bayes classifier assuming a $d$-dimensional Normal distribution Requires estimates for $\mu_{i}$ and $\Sigma_{i}$

Estimate for $\mu_{i}$ needs much less training data

- Goal
classifier using only the mean vectors per class
$\Longrightarrow$ Nearest-neighbor classifier


## Nearest-Neighbor Classification

## Example



Classifier:
Cat q is a dog!Instance-Based Learning
Related to Case-Based Reasoning

## Nearest-Neighbor Classification

## Overview

## Base method

- Training objects $o$ as feature (attribute) vectors $o=\left(o_{1}, \ldots, o_{d}\right)$
- Calculate the mean vector $\mu_{i}$ for each class $c_{i}$
- Assign unseen object to class $c_{i}$ with nearest mean vector $\mu_{i}$


## Generalisations

- Use more than one representative per class
- Consider $k>1$ neighbors
- Weight the classes of the $k$-nearest neighbors


## Nearest-Neighbor Classification

## Notions

- Distance function defines similarity (dissimilarity) for pairs of objects
- $k$ : number of neighbors considered
- Decision Set
set of $k$-nearest neighbors considered for classification
- Decision rule
how to determine the class of the unseen object from the classes of the decision set?


## Example



Uniform weight for the decision set
$k=1$ : classification as „+", $k=5$ classification as „,"
Inverse squared distance as weight for the decision set
$k=1$ and $k=5$ : classification as ,,+"

## Nearest-Neighbor Classification

## Choice of Parameter $k$

- „too small" $k$ : very sensitive to outliers
-,too large" $k$ : many objects from other clusters (classes) in the decision set
- medium $k$ : highest classification accuracy, often $1 \ll k<10$



## Nearest-Neighbor Classification

## Decision Rule

## Standard rule

Choose the majority class within the decision set

## Weighted decision rule

Weight the classes of the decision set

- By distance
- By class distribution (often skewed!)
class A: $95 \%$, class B $5 \%$
Decision set $=\{A, A, A, A, B, B, B\}$
Standard rule $\Rightarrow A$, Weighted rule $\Rightarrow B$


## Index Support for $k$-Nearest-Neighbor Queries

- Balanced index tree (such as X-tree or M-tree)
- Query point p
- PartitionList

BBs of subtrees that need to be processed, sorted in ascending order w.r.t. MinDist to $p$

- NN

Nearest neighbor of $p$ in the data pages read so far


## Index Support for $k$-Nearest-Neighbor Queries

- Remove all BBs from PartitionList that have a larger distance to $p$ than the currently best NN of p
- PartitionList is sorted in ascending order w.r.t. MinDist to p
- Always pick the first element of PartitionList as the next subtree to be explored

Does not read any unnecessary disk pages!

- Query processing limited to a few paths of the index structure

Average runtime $\mathrm{O}(\log n)$ for „not too many" attributesFor very large numbers of attributes: $\mathrm{O}(n)$

## Nearest-Neighbor Classification

## Discussion

+ Local method
Does not have to find a global decision function (decision surface)
+ High classification accuracy
In many applications
+ Incremental
Classifier can easily be adapted to new training objects
+ Can be used also for prediction
- Application of classifier expensive Requires $k$-nearest neighbor query
- Does not generate explicit knowledge about the classes


## Introduction [Burges 1998]

Input

$$
S=\left\{\left(x_{1}, y_{1}\right), \ldots,\left(x_{n}, y_{n}\right)\right\} \quad x_{i} \in X
$$

a training set
of objects and their known classes

$$
y_{i} \in\{-1,+1\}
$$

Output
a classifier

$$
f: X \rightarrow\{-1,+1\}
$$

Goal
Find the best separating hyperplane (e.g., lowest classification error)
Two-class problem


## Support Vector Machines

## Introduction

Half-space:
w. $\mathrm{x}+\mathrm{b}>0$
(Class +1 )

w. $\mathrm{x}+\mathrm{b}<0$
(Class -1)

- Classification based on the sign of the decision function

$$
f_{w, b}(x)=w \cdot x+b
$$

- "." denotes the inner product of two vectors

Hyperplane: w.x $+\mathrm{b}=0$

## Support Vector Machines

## Introduction



Choose hyperplane with largest margin (maximum distance to closest training object)

## Support Vector Machines



$$
\text { w. } x_{1}+b=0
$$

$$
\text { w. } x_{2}+b=1
$$

$$
\text { à w. }\left(\mathrm{x}_{2}-\mathrm{x}_{1}\right)=1
$$

$$
\text { à }\|w\|\left\|x_{2}-x_{1}\right\| \cos 0=1
$$

$$
\gamma=\left\|x_{2}-x_{1}\right\|=\frac{1}{\|w\|}
$$

$\gamma$ : margin

## Method

## Problem

- Minimize $\|w\|^{2}$
- Under the constraints $\forall i=1, \ldots, n: y_{i}\left(w \cdot x_{i}+b\right)-1 \geq 0$

Dual problem

- Introduce dual variables $\alpha_{i}$ for each training object $i$
- Find $\alpha_{i}$ maximizing

$$
L(\alpha)=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i, j=1}^{n} \alpha_{i} \cdot \alpha_{j} \cdot y_{i} \cdot y_{j} \cdot x_{i} \cdot x_{j}
$$

under the constraints $\alpha_{i} \geq 0$ and $\sum_{i=1}^{n} \alpha_{i} \cdot y_{i}=0$



- Only training objects with $\alpha_{i}>0$ contribute to w
- These training objects are the support vectors
$\left\{\begin{array}{l}\text { Typically, number of } \\ \text { support vectors } \ll n\end{array}\right.$


## Support Vector Machines

Non-Linear Classifiers


## Support Vector Machines

## Non-Linear Classifiers

- Decision function $\quad f_{w, b}(x)=w \cdot \Psi(x)+b$
- Kernel of two objects $\forall x, x^{\prime} \in X: \quad K\left(x, x^{\prime}\right)=\Psi(x) . \Psi\left(x^{\prime}\right)$
- Explicit computation of $\Psi(x) \quad$ is not necessary
- Example: $\quad \Psi(x)=\left(x_{1}^{2}, x_{2}^{2}, \sqrt{2} x_{1} x_{2}, \sqrt{2} x_{1}, \sqrt{2} x_{2}, 1\right)$

$$
K\left(x, x^{\prime}\right)=\Psi(x) \cdot \Psi\left(x^{\prime}\right)=\left(x \cdot x^{\prime}+1\right)^{2}
$$

## Support Vector Machines

## Kernels

－Kernel is a similarity measure
－ $\mathrm{K}(\mathrm{x}, \mathrm{x}$＇）is a kernel iff

$$
\forall x_{i} \in X:\left(\begin{array}{c}
K\left(x_{1}, x_{1}\right) K\left(x_{1}, x_{2}\right) \text { 匈 } \\
K\left(x_{2}, x_{1}\right) \\
K\left(x_{2}, x_{2}\right) \text { 娄 } \\
\text { 圆 }
\end{array}\right)
$$

is a symmetric and positive definite matrix

## Support Vector Machines

## SVM for Protein Classification [Leslie et al 2002]

- Two sequences are similar when they share many common substrings (subsequences)
..
and $|\mathrm{s}|$ denotes the length of string s
- Very high classification accuracy for protein sequences
- Variation of the kernel (when allowing gaps in the matching subsequences)

$$
K\left(x, x^{\prime}\right)=\sum_{s \text { common substring }} \lambda^{\text {length }(s, x)+\text { length }\left(s, x^{\prime}\right)}
$$

length(s,x): length of the subsequence of x matching s

## SVM for Prediction of Translation Initiation Sites [Zien et al 2000]

- Translation initiation site (TIS): starting position of a protein coding region in DNA All TIS start with the triplet "ATG"
- Problem: given an "ATG" triplet, does it belong to a TIS?
- Representation of DNA

Window of 200 nucleotides around candidate "ATG"
Encode each nucleotide with a 5 bit word $(00001,00010, \ldots, 10000)$ for A, C, G, T and unknown
$\rightarrow$ Vectors of 1000 bits

## Support Vector Machines

## SVM for Prediction of Translation Initiation Sites

- Kernels

$$
\begin{array}{ll}
K\left(x, x^{\prime}\right)=\left(\mathrm{x} . \mathrm{x}^{\prime}\right)^{\mathrm{d}} & \begin{array}{l}
\mathrm{d}=1: \text { number of common bits } \\
\mathrm{d}=2: \text { number of common pairs of bits }
\end{array}
\end{array}
$$

Locally improved kernel: compare only small window around "ATG"

- Experimental results

Long range correlations do not improve performanceLocally improved kernel performs best
Outperforms state-of-the-art methods

## Discussion

+ Strong mathematical foundation
+ Find global optimum
+ Scale well to very high-dimensional datasets
+ Very high classification accuracy
In many challenging applications
- Inefficient model construction

Long training times ( $\sim \mathrm{O}\left(n^{2}\right)$ )

- Model is hard to interpret

Learn only weights of features
Weights tend to be almost uniformly distributed

## Multi-relational Classification

## The Single Table Assumption

- Existing data mining algorithms expect data in a single table
- But in reality, DBs consist of multiple tables
- Naive solution: join all tables into a single one (universal relation) and apply (single-relational) data mining algorithm

Purchases

| Client\# | Date | Item | Quantity |
| :--- | :--- | :--- | :--- |
| 2765 | $02 / 25 / 2005$ | A | 5 |
| 3417 | $02 / 26 / 2005$ | B | 1 |
| 1005 | $02 / 26 / 2005$ | C | 12 |
| $\ldots$ |  |  |  |

Clients

| Client\# | Name | Age |
| :--- | :--- | :--- |
| 1005 | Jones | 35 |
| 1010 | Smith | 52 |
| 1054 | King | 27 |
| $\ldots$ |  |  |

## The Single Table Assumption

- Universal relation

| Client\# | Date | Item | Quantity | Name | Age |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 1005 | $02 / 26 / 2005$ | C | 12 | Jones | 35 |
| 1005 | $02 / 28 / 2005$ | B | 2 | Jones | 35 |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |
| 2765 | $02 / 25 / 2005$ | A | 5 | Bornman | 23 |
| $\ldots$ |  |  |  |  |  |

There are no more client entities!
What if rule depends on how many different items were purchased by a client?

## Multi-relational Classification

## Aggregating Related Tables

- Enhancing „target table" by aggregates of the related tuples in other tables
- Aggregation operators: COUNT, SUM, MIN, AVG, . . .

| Client\# | Name | Age | Overall <br> Quantity of <br> Item A | Overall <br> Quantity of <br> Item B | $\cdots$ |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 1005 | Jones | 35 | 0 | 10 | $\cdots$ |
| 1010 | Smith | 52 | 35 | 0 | $\cdots$ |
| $\ldots$ |  |  |  |  | $\cdots$ |

More meaningful! But what aggregates to consider? And what if attributes of the other clients that have purchased the same item are relevant?

## Multi-Relational Data Mining

- Data mining methods for multi-table databases
- Pattern search space much larger than for single tables
- Testing the validity of a pattern more expensive
- Similar data mining tasks
classification, clustering, association rules, . . .
. . . plus some tasks specific to the multi-relational case
- Single table (propositional) algorithms can be upgraded to multiple tables (first order predicate logic)


## Multi-relational Classification

## Inductive Logic Programming (ILP)

- Goal: learn logic programs from example data
- Knowledge representation is expressive and understandable
- Examples: tuples from multiple tables
- Hypotheses: sets of rules
- Use of background knowledge also set of rules


## Logic Programs and Databases

- Logic program: set of clauses
- Clause: rule of the form „Head $\leftarrow$ Body" where Head / Body consist of atoms connected using the logical operators

$$
\wedge, \vee \text { and } \checkmark
$$

- Atom: predicate applied to some terms
- Predicate: boolean function with arguments (terms)
- Term: constant (e.g., mary), variable (e.g., X), function symbol applied to some term


## Logic Programs and Databases

- Example rule

$$
\text { father }(X, Y) \vee \text { mother }(X, Y) \leftarrow \operatorname{parent}(X, Y)
$$

- Definite clauses: exactly one atom in the head

$$
\operatorname{parent}(X, Y) \leftarrow \operatorname{father}(X, Y) \vee \text { mother }(X, Y)
$$

- Horn clauses

One (positive) atom in the head, conjunction of body atoms

$$
\text { mother }(X, Y) \leftarrow \operatorname{parent}(X, Y) \wedge \text { female }(Y)
$$

## Classical Rule Induction Task

- Given:
set P of examples from target relation (positive examples) set N of examples not from target relation (negative examples)
background predicates B
hypothesis (rule) language
- Find a set of rules that explains all positive and none of the negative examples


## Multi-relational Classification

## Example

```
Training examples
Background knowledge
    daughter(mary,ann) + parent(ann,mary)female(ann) daughter(eve,tom)
    + parent(ann,tom) female(mary)
    daughter(tom,ann) - parent(tom,eve) female(eve)
    daughter(eve,ann) - parent(tom,ian)
```

Hypothesis language
definite clauses
Resulting rule

$$
\operatorname{daughter}(X, Y) \leftarrow \operatorname{parent}(Y, X) \wedge \text { female }(X)
$$

## The Sequential Covering Algorithm

Hypothesis (H) := \{\}

## Repeat

find a clause c that covers some positive and no negative examples; add c to H ;
delete all positive examples implied by c
Until no more (uncovered) positive examples

$$
B\left[\begin{array}{c}
W
\end{array} \cup\{c\}\right.
$$

Construction of new clauses: search of the space of clauses applying some refinement operator

## Structuring the Space of Clauses

- Substitution

$$
\theta=\left\{V_{1} / t_{1}, \text { 罒, } V_{n} / t_{n}\right\}
$$ assignment of terms $t_{i}$ to variables $V_{i}$

- Clauses as sets of atoms (literals)

$$
\begin{aligned}
& \text { Head } \leftarrow \operatorname{Body} \Leftrightarrow H \text { Head } \vee \neg \operatorname{Body} \\
& \text { e.g., daughter }(X, Y) \leftarrow \operatorname{parent}(Y, X): \\
& \{\operatorname{daughter}(X, Y), \neg \operatorname{parent}(Y, X)\}
\end{aligned}
$$

Clause $\boldsymbol{c} \theta$ - subsumes clause $c^{\prime}$
if there exists a substitution $\theta$ such that $c \theta \subseteq c^{\prime}$

## Multi-relational Classification

## Structuring the Space of Clauses

- Examples

$$
\begin{aligned}
& c=\operatorname{daughter}(X, Y) \leftarrow \operatorname{parent}(Y, X) \\
& \theta=\{X / \text { mary }, Y / \text { ann }\} \\
& c \theta=\operatorname{daughter}(\text { mary }, \text { ann }) \leftarrow \operatorname{parent}(\text { ann }, \text { mary })
\end{aligned}
$$

$$
\begin{aligned}
& c=\operatorname{daughter}(X, Y) \leftarrow \operatorname{parent}(Y, X) \\
& c^{\prime}=\operatorname{daughter}(X, Y) \leftarrow \operatorname{female}(X) \wedge \operatorname{parent}(Y, X) \\
& \theta=\{ \} \\
& c \theta=c \subseteq c^{\prime}, \text { i.e. } c \theta-\operatorname{subsumes} c^{\prime}
\end{aligned}
$$

## Structuring the Space of Clauses

- Syntactic notion of generality
clause c is at least as general as clause $\mathrm{c}^{\text {‘ }}\left(\mathrm{c} \leq\right.$ d ' $^{\prime} \mathrm{iff}$

$$
c \theta-\text { subsumes } \quad c^{\prime}
$$

c is more general than clause $\mathrm{c}^{\text {‘ }}$ iff

$$
c \leq c^{\prime} \wedge \neg\left(c^{\prime} \leq c\right)
$$

c is a generalization of $\mathrm{c}^{\text {c }}, \mathrm{c}^{\text {d }}$ a specialization of c
R If c does not cover an example, none of its specializations do If c covers an example, all of its generalizations do

## Searching the Space of Clauses

- Top-down approach:
start from most general clauses
recursively apply refinement operators
- Refinement operator
$\theta$ - subsumption-based
returns all most general specializations of a given clause
- Types of refinements
apply a substitution to a clause or
add a literal to the body of the clause


## Example



Refinement graph (lattice)

## Top－Down Search of Refinement Graphs

Hypothesis $(H):=\{ \}$
repeat
clause
$c:=p\left(X_{1}\right.$, 㕼，$\left.X_{n}\right) \leftarrow$

## repeat

build the set $S$ of all refinements of $c$ ；
$c:=$ the best element of $S$（according to some heuristic）
until stopping criterion satisfied（ c is consistent with $B$ 聯 $H$
$\operatorname{add} c$ to $H$ ；
delete all positive examples implied by $c($ using $B)[$ 网 $H$
until no more（uncovered）positive examples（i．e．，$H$ complete）

FOIL [Quinlan 1990]

- Top-down search of refinement graph
- Weighted information gain as heuristic to choose best clause
- Heuristic can be modified to allow clauses covering (some) negative examples
$\rightarrow$ handling of noisy data
- Declarative bias to reduce search space syntactic restrictions on clauses to be considered to be provided by the user


## Declarative Bias

- Argument types / domains (relational DBS)
- Input / output modes of arguments
argument must / must not be instantiated when predicate added
- Parametrized language bias
e.g., maximum number of variables, literals, . . . per clause
- Clause templates

Ex.: $\quad P(X, Y) \leftarrow Q(X, Z) \wedge R(Z, Y)$ where $\mathrm{P}, \mathrm{Q}, \mathrm{R}$ denote predicate variables

Declarative bias difficult to specify for user (syntactic!)

## Multi-relational Classification

## CrossMine [Yin, Han, Yang \& Yu 2004]

- Several improvements of FOIL and similar ILP classification methods
- Evaluation of alternative refinement operator requires joins, which are very expensive DB operations
$\rightarrow$ TupleID propagation (virtual joins)
propagate tupleIDs and their class labels from the target table to related tables
- Relationship tables have no attributes and may not yield a high information gain
would never been chosen by FOIL
$\rightarrow$ Increased look ahead (two instead of one literal)


## Multi-relational Classification

## TupleID Propagation



Related table

## Regression Analysis

## Prediction

## Commonality with classification

- First, construct a model
- Second, use model to predict unknown value

Major method for prediction is regression

- Simple and multiple regression
- Linear and non-linear regression

Difference from classification

- Classification refers to predict categorical class label
- Prediction models continuous-valued functions


## Regression Analysis

## Linear Regression

- Predict the values of the response variable $y$ based on a linear combination of the given values of the predictor variable(s) $x_{i}$

$$
\hat{y}=a_{0}+\sum_{j=1}^{d} a_{j} x_{j}
$$

- Simple regression: one predictor variable $\rightarrow$ regression line
- Multiple regression: several predictor variables $\rightarrow$ regression plane
- Residuals: differences between observed and predicted values

Use the residuals to measure the model fit

## Linear Regression

$$
y(i)=\hat{y}(i)+e(i)=a_{0}+\sum_{i=1}^{d} a_{j} x_{j}(i)+e(i), \quad 1 \leq i \leq n
$$

- $y$ : vector of the $y$ values for the $n$ training objects
- $\quad X$ : matrix of the values of the $d$ predictor variables for the $n$ training objects (and an additional column of 1s)
- $\quad e$ : vector of the residuals for the $n$ training objects
- Matrix notation:

$$
y=X a+e
$$

## Linear Regression

- Optimization goal: minimize $\sum_{i=1}^{n} e(i)^{2}=\sum_{i=1}^{n}\left[y(i)-\sum_{j=0}^{d} a_{j} x_{j}(i)\right]^{2}$
- Solution:
- Computational issues $\left(X^{T} X\right)^{-1} X^{T} y$
- $\mathrm{X}^{\mathrm{T}} \mathrm{X}$ must be invertible

Problems if linear dependencies between predictor variables

- Solution may be unstable

If predictor variables almost linear dependent
Equation solving e.g. using LU decomposition or SVD
Runtime complexity $\mathrm{O}\left(d^{2} n+d^{3}\right)$

## Locally Weighted Regression

Limitations of linear regression

- Only linear models
- One global model

Locally weighted regression
Construct an explicit approximation to $f$ over a local neighborhood of query instance $x q$
Weight the neighboring objects based on their distance to $x_{q}$
Distance-decreasing weight $K$
Related to nearest neighbor classification
$\rightarrow$ Minimize the squared local weighted error

## Locally Weighted Regression

## Local weighted error

- W.r.t. query instance $x q$
- Arbitrary approximating function
- Pairwise distance function $d$
- Three major alternatives:

$$
\begin{aligned}
& E\left(x_{q}\right)=\frac{1}{2} x_{x \in k_{-} \text {nearest_neighbors_of_- } x_{q}} \sum(f(x)-\hat{f}(x))^{2} \\
& E\left(x_{q}\right)=\frac{1}{2} \sum_{x \in \hat{D}}^{[f(x)-\hat{f}(x)]^{2} \cdot K\left(d\left(x_{q}, x\right)\right)} \\
& E\left(x_{q}\right)=\frac{1}{2}_{x \in k_{-} \text {nearest_neighbors_of_} x_{q}} \sum_{x^{\prime}}(f(x)-\hat{f}(x))^{2} K\left(d\left(x_{q}, x\right)\right)
\end{aligned}
$$

## Discussion

+ Strong mathematical foundation
+ Simple to calculate and to understand
For moderate number of dimensions
+ High classification accuracy
In many applications
- Many dependencies are non-linear

Can be generalized

- Model is global

Cannot adapt well to locally different data distributions
But: Locally weighted regression

