Data Preparation

(Data pre-processing)

Data Preparation

- Introduction to Data Preparation
- Types of Data and Basic statistics
- Discretization of Continuous Variables
- Working in the R environment
- Outliers
- Data Transformation
- Missing Data
- Data Integration
- Data Reduction

INTRODUCTION TO DATA PREPARATION

Why Prepare Data?

- Some data preparation is needed for all mining tools
- The purpose of preparation is to transform data sets so that their information content is best exposed to the mining tool
- Error prediction rate should be lower (or the same) after the preparation as before it

Why Prepare Data?

- Preparing data also prepares the miner so that when using prepared data the miner produces better models, faster
- GIGO - good data is a prerequisite for producing effective models of any type

Why Prepare Data?

- Data need to be formatted for a given software tool
- Data need to be made adequate for a given method
- Data in the real world is dirty
  - incomplete: lacking attribute values, lacking certain attributes of interest, or containing only aggregate data
    - e.g., occupation=""
  - noisy: containing errors or outliers
    - e.g., Salary="-10", Age="222"
  - inconsistent: containing discrepancies in codes or names
    - e.g., Age="42" Birthday="03/07/1997"
    - e.g., Was rating "1,2,3", now rating "A, B, C"
    - e.g., discrepancy between duplicate records
    - e.g., ***Endereço:*** travessa da Igreja de Nevogilde ***Freguesia:*** Paranhos

Major Tasks in Data Preparation

- Data discretization
  - Part of data reduction but with particular importance, especially for numerical data
- Data cleaning
  - Fill in missing values, smooth noisy data, identify or remove outliers, and resolve inconsistencies
- Data integration
  - Integration of multiple databases, data cubes, or files
- Data transformation
  - Normalization and aggregation
- Data reduction
  - Obtains reduced representation in volume but produces the same or similar analytical results

Data Preparation as a step in the Knowledge Discovery Process

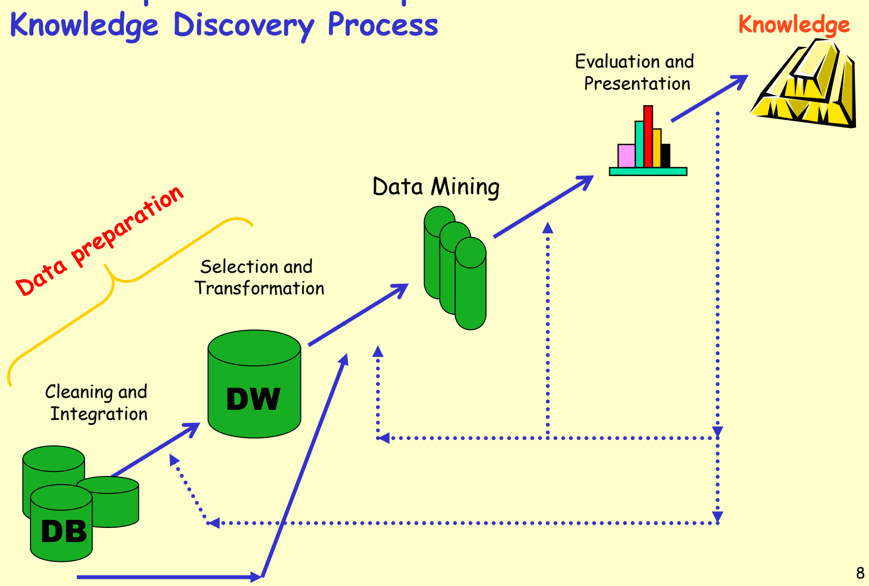

CRISP-DM

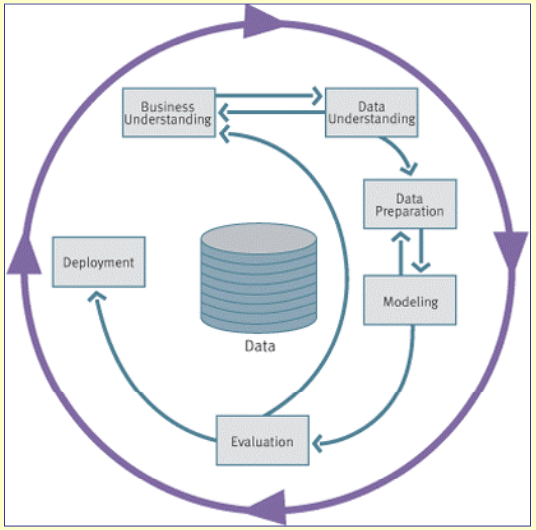

CRISP-DM is a comprehensive data mining methodology and process model that provides anyone—from novices to data mining experts—with a complete blueprint for conducting a data mining project.

CRISP-DM breaks down the life cycle of a data mining project into six phases.

CRISP-DM Phases and Tasks

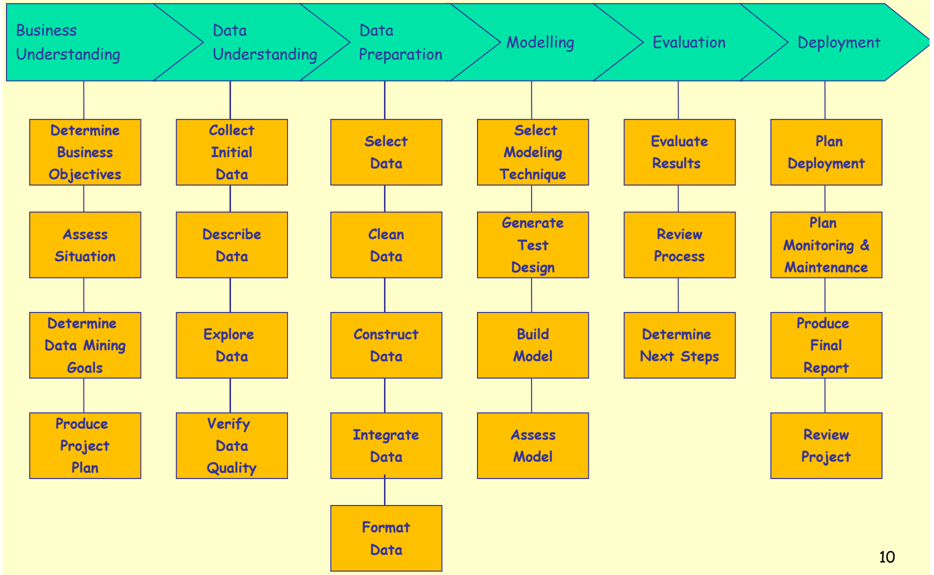

CRISP-DM Phases and Tasks

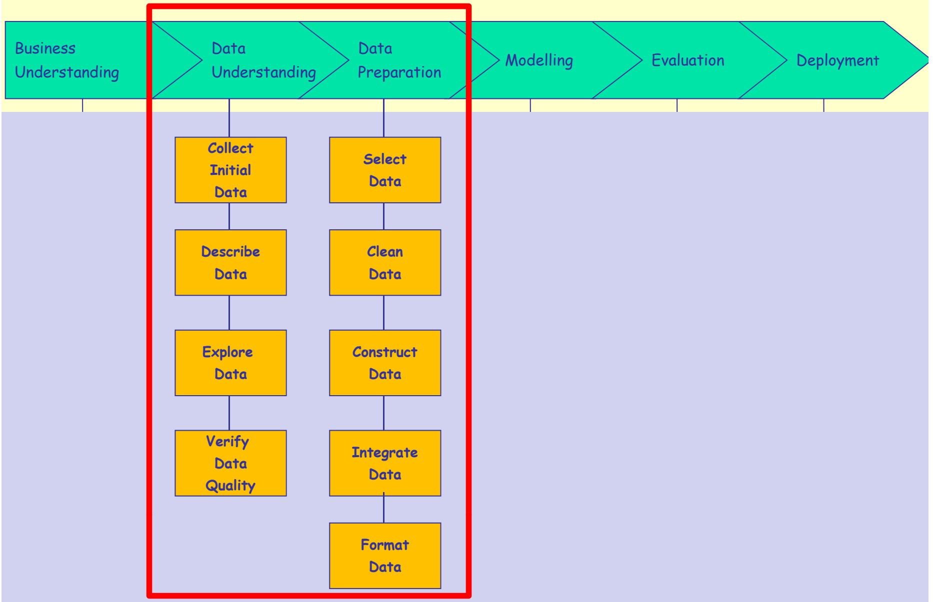

CRISP-DM: Data Understanding

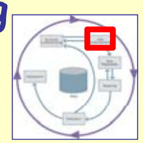

- **Collect data**
  - List the datasets acquired (locations, methods used to acquire, problems encountered and solutions achieved).
- **Describe data**
  - Check data volume and examine its gross properties.
  - Accessibility and availability of attributes. Attribute types, range, correlations, the identities.
  - Understand the meaning of each attribute and attribute value in business terms.
  - For each attribute, compute basic statistics (e.g., distribution, average, max, min, standard deviation, variance, mode, skewness).

CRISP-DM: Data Understanding

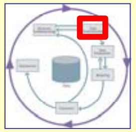

- **Explore data**
  - Analyze properties of interesting attributes in detail.
    - *Distribution, relations between pairs or small numbers of attributes, properties of significant sub-populations, simple statistical analyses.*
- **Verify data quality**
  - Identify special values and catalogue their meaning.
  - Does it cover all the cases required? Does it contain errors and how common are they?
  - Identify missing attributes and blank fields. Meaning of missing data.
  - Do the meanings of attributes and contained values fit together?
  - Check spelling of values *(e.g., same value but sometime beginning with a lower case letter, sometimes with an upper case letter).*
  - Check for plausibility of values, e.g. all fields have the same or nearly the same values.

CRISP-DM: Data Preparation

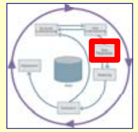

- **Select data**
  - Reconsider data selection criteria.
  - Decide which dataset will be used.
  - Collect appropriate additional data (internal or external).
  - Consider use of sampling techniques.
  - Explain why certain data was included or excluded.
- **Clean data**
  - Correct, remove or ignore noise.
  - Decide how to deal with special values and their meaning *(99 for marital status).*
  - Aggregation level, missing values, etc.
  - Outliers?

CRISP-DM: Data Preparation

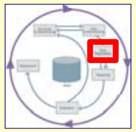

- **Construct data**
  - Derived attributes.
  - Background knowledge.
  - How can missing attributes be constructed or imputed?
- **Integrate data**
  - Integrate sources and store result (new tables and records).
- **Format Data**
  - Rearranging attributes *(Some tools have requirements on the order of the attributes, e.g. first field being a unique identifier for each record or last field being the outcome field the model is to predict).*
  - Reordering records *(Perhaps the modelling tool requires that the records be sorted according to the value of the outcome attribute).*
  - Reformatted within-value *(These are purely syntactic changes made to satisfy the requirements of the specific modelling tool, remove illegal characters, uppercase lowercase).*

TYPES OF DATA AND BASIC STATISTICS

Types of Data Measurements

- Measurements differ in their nature and the amount of information they give
- Qualitative vs. Quantitative

Types of Measurements

- Nominal scale
- Categorical scale
- Ordinal scale
- Interval scale
- Ratio scale

Qualitative

Quantitative

Discrete or Continuous

More information content

Types of Measurements: Examples

- Nominal:
  - ID numbers, Names of people
- Categorical:
  - eye color, zip codes
- Ordinal:
  - rankings (e.g., taste of potato chips on a scale from 1-10), grades, height in {tall, medium, short}
- Interval:
  - calendar dates, temperatures in Celsius or Fahrenheit, GRE score
- Ratio:
  - temperature in Kelvin, length, time, counts

Types of Measurements: Examples

Day	Outlook	Temperature	Humidity	Wind	*PlayTennis?*
1	Sunny	85	85	Light	No
2	Sunny	80	90	Strong	No
3	Overcast	83	86	Light	Yes
4	Rain	70	96	Light	Yes
5	Rain				
6	Rain				
7	Overcast				
8	Sunny				
9	Sunny				
10	Rain				
11	Sunny				
12	Overcast				
13	Overcast				
14	Rain				

Day	Outlook	Temperature	Humidity	Wind	*PlayTennis?*
1	Sunny	Hot	High	Light	No
2	Sunny	Hot	High	Strong	No
3	Overcast	Hot	High	Light	Yes
4	Rain	Mild	High	Light	Yes
5	Rain	Cool	Normal	Light	Yes
6	Rain	Cool	Normal	Strong	No
7	Overcast	Cool	Normal	Strong	Yes
8	Sunny	Mild	High	Light	No
9	Sunny	Cool	Normal	Light	Yes
10	Rain	Mild	Normal	Light	Yes
11	Sunny	Mild	Normal	Strong	Yes
12	Overcast	Mild	High	Strong	Yes
13	Overcast	Hot	Normal	Light	Yes
14	Rain	Mild	High	Strong	No

Basic Statistics

http://www.liaad.up.pt/~ltorgo/Regression/DataSets.html

A20645 | f_x -121.24

	A	B	C	D	E	F	G	H	I	J	K
1	http://www.liaad.up.pt/~ltorgo/Regression/DataSets.html										
2	California Housing										
3											
4											
5	longitude	latitude	housing Median Age	total Rooms	total Bedrooms	population	households	median Income	median House Value		
6	-122.23	37.88	41	880	129	322	126	8.3252	452600		
7	-122.22	37.86	21	7099	1106	2401	1138	8.3014	358500		
8	-122.24	37.85	52	1467	190	496	177	7.2574	352100		
9	-122.25	37.85	52	1274	235	558	219	5.6431	341300		
10	-122.25	37.85	52	1627	280	565	259	3.8462	342200		
11	-122.25	37.85	52	919	213	413	193	4.0368	269700		
12	-122.25	37.84	52	2535	489	1094	514	3.6591	299200		
20638	-121.45	39.26	15	2319	416	1047	385	3.125	115600		
20639	-121.53	39.19	27	2080	412	1082	382	2.5495	98300		
20640	-121.56	39.27	28	2332	395	1041	344	3.7125	116800		
20641	-121.09	39.48	25	1665	374	845	330	1.5603	78100		
20642	-121.21	39.49	18	697	150	356	114	2.5568	77100		
20643	-121.22	39.43	17	2254	485	1007	433	1.7	92300		
20644	-121.32	39.43	18	1860	409	741	349	1.8672	84700		
20645	-121.24	39.37	16	2785	616	1387	530	2.3886	89400		
20646											
20647											

data | summary | Sheet3

Ready

Summary Statistics

(Excel)

4								
5	N	20640	20640	20640	20640	20640	20640	20640
6	Average	28.64	2635.76	537.90	1425.48	499.54	3.87	206855.82
7	Min	1	2	1	3	1	0.4999	14999
8	Max	52	39320	6445	35682	6082	15.0001	500001
9	Range	51	39318	6444	35679	6081	14.5002	485002
10	Median	29	2127	435	1166	409	3.5348	179700
11	Standard Deviation	12.59	2181.62	421.25	1132.46	382.33	1.90	115395.62
12	Variance	158.40	4759445.11	177449.80	1282470.46	146176.04	3.61	13316148163.04
13	Mode	52	1527	280	891	306	3.125	500001
14	Skewness	0.06	4.15	3.45	4.94	3.41	1.65	0.98

16	longitude	latitude	housing Median Age	total Rooms	total Bedrooms	population	households	median Income	median House Value
17	-122.23	37.88	41	880	129	322	126	8.3252	452600
18	-122.22	37.86	21	7099	1106	2401	1138	8.3014	358500

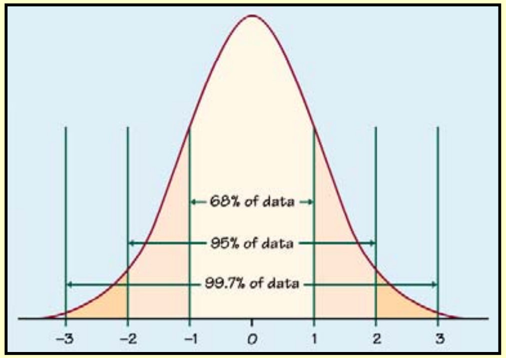

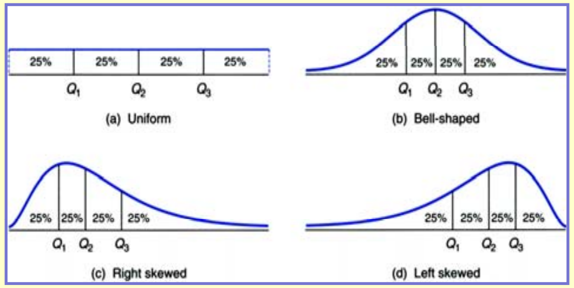

Histograms

(SPSS)

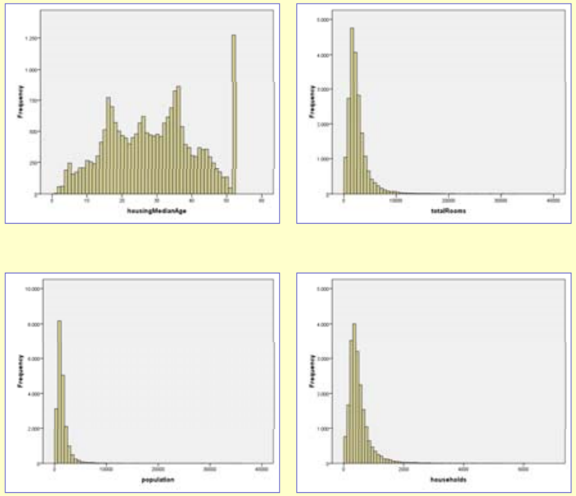

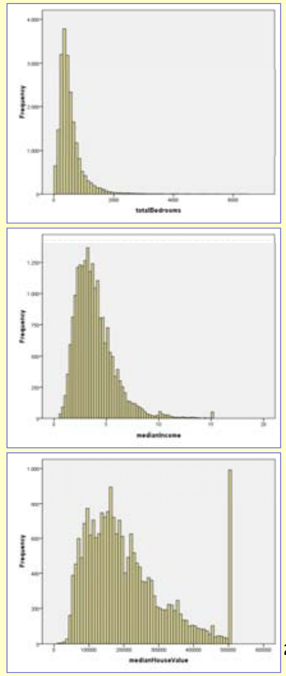

Box Plots

(SPSS)

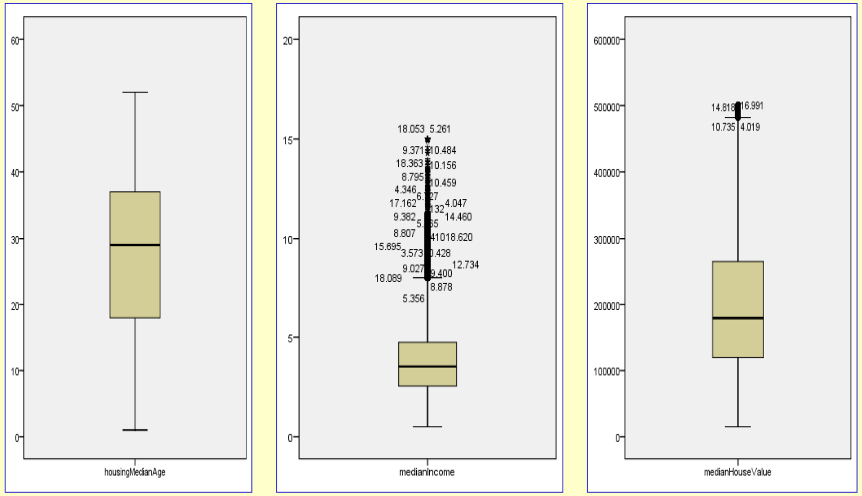

Data Conversion

- Some tools can deal with nominal values but other need fields to be numeric
- Convert ordinal fields to numeric to be able to use ">" and "<" comparisons on such fields.
  - A → 4.0
  - A- → 3.7
  - B+ → 3.3
  - B → 3.0
- Multi-valued, unordered attributes with small no. of values
  - e.g. Color=Red, Orange, Yellow, …, Violet
  - for each value *v* create a binary "flag" variable *C_v* , which is 1 if Color=*v*, 0 otherwise

Conversion: Nominal, Many Values

- Examples:
  - US State Code (50 values)
  - Profession Code (7,000 values, but only few frequent)
- Ignore ID-like fields whose values are unique for each record
- For other fields, group values "naturally":
  - e.g. 50 US States → 3 or 5 regions
  - Profession – select most frequent ones, group the rest
- Create binary flag-fields for selected values

DISCRETIZATION OF CONTINUOUS VARIABLES

Discretization

- Divide the range of a continuous attribute into intervals
  - Some methods require discrete values, e.g. most versions of Naïve Bayes, CHAID
  - Reduce data size by discretization
  - Prepare for further analysis
  - Discretization is very useful for generating a summary of data
- Also called "binning"

Equal-width Binning

- It divides the range into *N* intervals of equal size (range): uniform grid
- If *A* and *B* are the lowest and highest values of the attribute, the width of intervals will be: $W = (B - A)/N$.

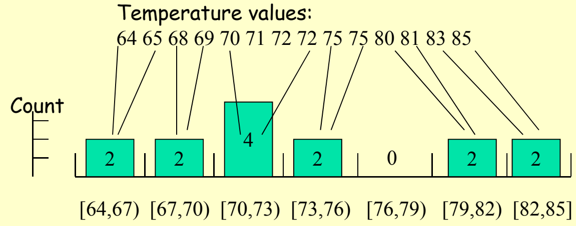

Equal Width, bins Low <= value < High

Equal-width Binning

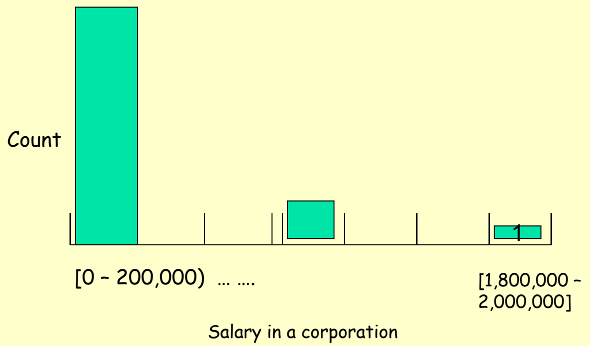

Disadvantage

(a) Unsupervised

(b) Where does N come from?

(c) Sensitive to outliers

Advantage

(a) simple and easy to implement

(b) produce a reasonable abstraction of data

Equal-depth (or height) Binning

- It divides the range into *N* intervals, each containing *approximately* the same number of samples
  - Generally preferred because avoids clumping
  - In practice, "almost-equal" height binning is used to give more intuitive breakpoints
- Additional considerations:
  - don't split frequent values across bins
  - create separate bins for special values (e.g. 0)
  - readable breakpoints (e.g. round breakpoints

Equal-depth (or height) Binning

Temperature values:
64 65 68 69 70 71 72 72 75 75 80 81 83 85

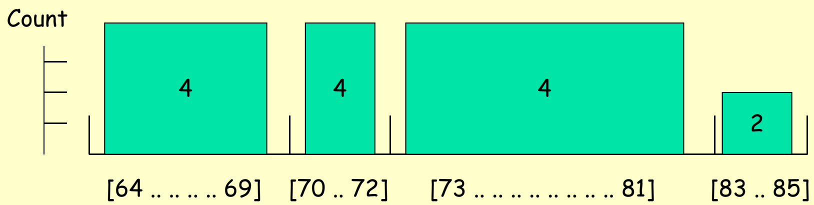

Equal Height = 4, except for the last bin

Discretization considerations

- Class-independent methods
  - Equal Width is simpler, good for many classes
    - can fail miserably for unequal distributions
  - Equal Height gives better results
- Class-dependent methods can be better for classification
  - Decision tree methods build discretization on the fly
  - Naïve Bayes requires initial discretization
- Many other methods exist …

Method 1R

- Developed by Holte (1993).
- It is a supervised discretization method using binning.
- After sorting the data, the range of continuous values is divided into a number of disjoint intervals and the boundaries of those intervals are adjusted based on the class labels associated with the values of the feature.
- Each interval should contain a given minimum of instances (6 by default) with the exception of the last one.
- The adjustment of the boundary continues until the next values belongs to a class different to the majority class in the adjacent interval.

1R Example

Interval contains at leas 6 elements

Adjustment of the boundary continues until the next values belongs to a class different to the majority class in the adjacent interval.

	1	2	3	4	5	6	1	2	3	4	5	6	7	1	2	3	4	5	6	7	1	2	3	4	5	6	7	1	2	3	4
Var	65	78	79	79	81	81	82	82	82	82	82	82	83	83	83	83	83	84	84	84	84	84	84	84	84	84	85	85	85	85	85
Class	2	1	2	2	2	1	1	2	1	2	2	2	2	1	2	2	2	1	2	2	1	1	2	2	1	1	1	2	2	2	2
majority						2						2							2							1					
new class	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	3	3	3	3

Exercise

- Discretize the following values using EW and ED binning
- 13, 15, 16, 16, 19, 20, 21, 22, 22, 25, 30, 33, 35, 35, 36, 40, 45

Entropy Based Discretization

Class dependent (classification)

1. Sort examples in increasing order
2. Each value forms an interval ('m' intervals)
3. Calculate the entropy measure of this discretization
4. Find the binary split boundary that minimizes the entropy function over all possible boundaries. The split is selected as a binary discretization.

$$E(S,T) = \frac{|S_1|}{|S|}Ent(S_1) + \frac{|S_2|}{|S|}Ent(S_2)$$

5. Apply the process recursively until some stopping criterion is met, e.g.,

$$Ent(S) - E(T,S) > \delta$$

Entropy

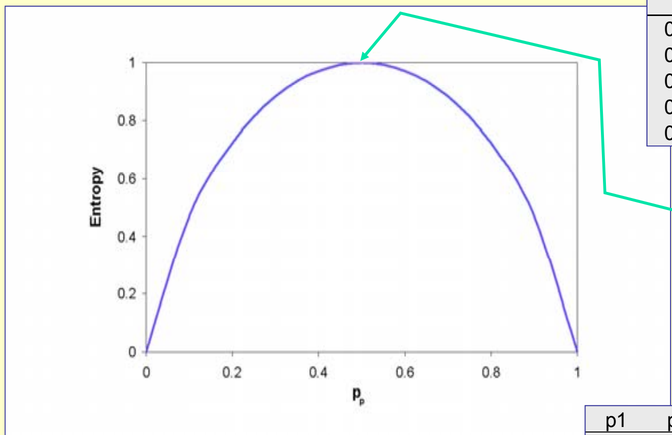

p	1-p	Ent
0.2	0.8	0.72
0.4	0.6	0.97
0.5	0.5	1
0.6	0.4	0.97
0.8	0.2	0.72

$\log_2(2)$

p1	p2	p3	Ent
0.1	0.1	0.8	0.92
0.2	0.2	0.6	1.37
0.1	0.45	0.45	1.37
0.2	0.4	0.4	1.52
0.3	0.3	0.4	1.57
0.33	0.33	0.33	1.58

$$Ent = -\sum_{c=1}^{N} p_c \cdot \log_2 p_c$$

$\log_2(3)$

Entropy/Impurity

- S - training set, $C_1,\ldots,C_N$ classes
- Entropy E(S) - measure of the impurity in a group of examples
  - p_c - proportion of C_c in S

$$\text{Impurity}(S) = -\sum_{c=1}^{N} p_c \cdot \log_2 p_c$$

Impurity

Very impure group

Less impure

Minimum impurity

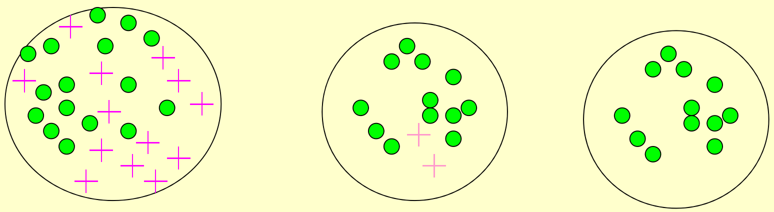

An example of entropy disc.

Test split temp < 71.5

Temp.	Play?
64	Yes
65	No
68	Yes
69	Yes
70	Yes
71	No
72	No
72	Yes
75	Yes
75	Yes
80	No
81	Yes
83	Yes
85	No

(4 yes, 2 no)

(5 yes, 3 no)

	yes	no
< 71.5	4	2
> 71.5	5	3

$$Ent(split\ 71.5) = \frac{6}{14} \cdot \left(\frac{4}{6} \log \frac{4}{6} + \frac{2}{6} \log \frac{2}{6} \right) + \frac{8}{14} \cdot \left(\frac{5}{8} \log \frac{5}{8} + \frac{3}{8} \log \frac{3}{8} \right) = 0.939$$

	yes	no
< 77	7	3
> 77	2	2

$$Ent(split\ 77) = \frac{10}{14} \cdot \left(\frac{7}{10} \log \frac{7}{10} + \frac{3}{10} \log \frac{3}{10} \right) + \frac{4}{14} \cdot \left(\frac{2}{4} \log \frac{2}{4} + \frac{2}{4} \log \frac{2}{4} \right) = \boxed{0.915}$$

An example (cont.)

Temp.	Play?
64	Yes
65	No
68	Yes
69	Yes
70	Yes
71	No
72	No
72	Yes
75	Yes
75	Yes
80	No
81	Yes
83	Yes
85	No

6th split

5th split

4th split

3rd split

2nd split

1st split

The method tests all split possibilities and chooses the split with smallest entropy.

In the first iteration a split at 84 is chosen.

The two resulting branches are processed recursively.

The fact that recursion only occurs in the first interval in this example is an artifact. In general both intervals have to be split.

The stopping criterion

Previous slide did not take into account the stopping criterion.

$$Ent(S) - E(T,S) > \delta$$

$$\partial > \frac{\log(N-1)}{N} + \frac{\Delta(T,S)}{N}$$

$$\Delta(S,T) = \log_2(3^c - 2) - [cEnt(S) - c_1 Ent(S_1) - c_2 Ent(S_2)]$$

c is the number of classes in S

c_1 is the number of classes in S_1

c_2 is the number of classes in S_2.

This is called the Minimum Description Length Principle (MDLP)

Exercise

- *Compute the gain of splitting this data in half*

Humidity	play
65	Yes
70	No
70	Yes
70	Yes
75	Yes
80	Yes
80	Yes
85	No
86	Yes
90	No
90	Yes
91	No
95	No
96	Yes

WORKING IN THE R ENVIRONMENT

Brief Introduction to R

- http://www.r-project.org/
- http://cran.r-project.org/doc/contrib/Short-refcard.pdf

- **Examples of Expressions:**
- 3+5*6
- a <- 2+2 (atribuir resultado de expressão a uma variável)
- 3^(3+2)
- b <- 1:10 (define sequência)
- b*3
- log(b)
- b+2
- seq(1,15,2) (define sequência)

more R examples

- ?log – help on a function
- help.search("clustering")
- objects() – lists existing objects
- rm(obj1, obj2,...) – removes existing objects
- str(obj) – displays the internal structure of an object
- Menu "File; Change dir..."
- dir()

- v <- c(1,2,3,4,5) - defines a vector
- m <- matrix(c(1,2,3,4),2,2) - defines 2x2 matrix de 2x2
- a <- array(1:8, c(2,2,2)) - defines 2x2x2 array
- m*2
- m[1,1]
- m[1,]

The california housing dataset in R

- File/change dir - to the directory with the dataset
- cal_housing <- read.table("aula_02_dataset_california.txt")
- cal_housing[1:10,] - first 10 rows
- cal_housing <- read.table("aula_02_dataset_california.txt", header = TRUE) - with headers
- summary(cal_housing) – summary statistics
- hist(cal_housing$totalRooms) – histogram
- hist(cal_housing[,4:4])
- pairs(cal_housing[,3:8]) – scatters for pairs of variables
- plot(cal_housing$population,cal_housing$households) – scatter 2 vars
- cor(cal_housing[,3:8]) – correlation matrix
- boxplot(cal_housing[,3:8]) - boxplots

Discritization with R

- **Load Dataset**
  - data <- read.table("aula_02_1R_exemplo.txt")
- **Load Data Preparation Package**
  - library(dprep)
- **Equal Width**
  - disc_data_ew <- disc.ew(data,1:1)
  - disc_data_ew
- **Equal Depth**
  - disc_data_ef <- disc.ef(data,1:1,3)
  - disc_data_ef
- **Holte 1R**
  - disc_data_1r <- disc.1r(data,1:1,6)
  - disc_data_1r
- **Entropy**
  - disc_data_ent <- disc.mentr(data,1:2)
  - disc_data_ent

OUTLIERS

Outliers

- Outliers are values thought to be out of range.
  - ***"An outlier is an observation that deviates so much from other observations as to arouse suspicion that it was generated by a different mechanism"***
  - Can be detected by standardizing observations and label the standardized values outside a predetermined bound as outliers
  - Outlier detection can be used for fraud detection or data cleaning
- Approaches:
  - do nothing
  - enforce upper and lower bounds
  - let binning handle the problem

Outlier detection

- **Univariate**
  - Compute mean and std. deviation. For k=2 or 3, x is an outlier if outside limits (normal distribution assumed)

$$(\bar{x} - ks, \bar{x} + ks)$$

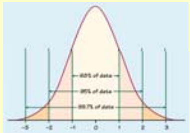

  - Boxplot: An observation is an **extreme** outlier if

    (Q1-3×IQR, Q3+3×IQR), where IQR=Q3-Q1

    (IQR = Inter Quartile Range)

    and declared a **mild** outlier if it lies outside of the interval

    (Q1-1.5×IQR, Q3+1.5×IQR).

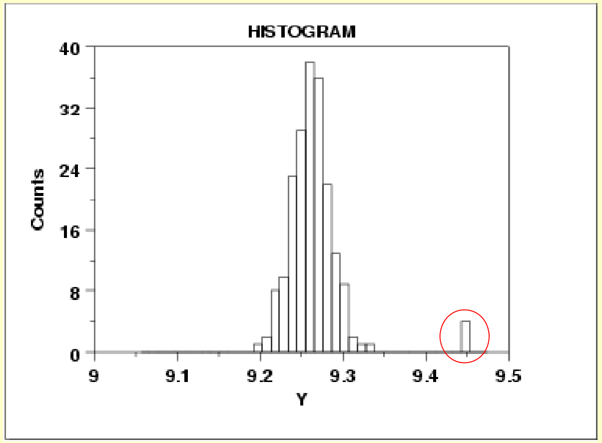

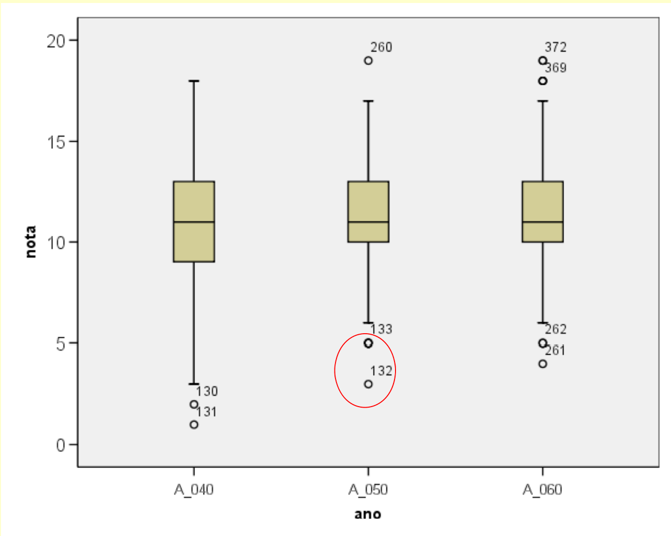

Outlier detection

- Multivariate
  - Clustering
    - Very small clusters are outliers
  - Distance based
    - An instance with very few neighbors within λ is regarded as an outlier

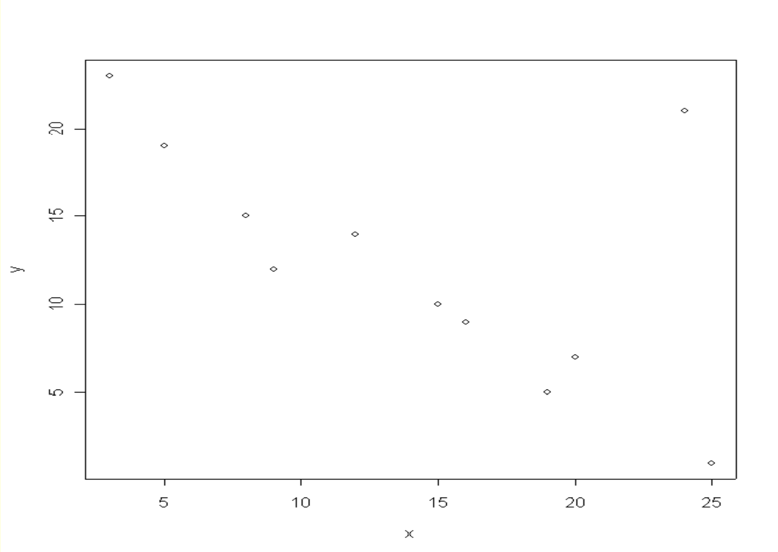

A bi-dimensional outlier that is not an outlier in either of its projections.

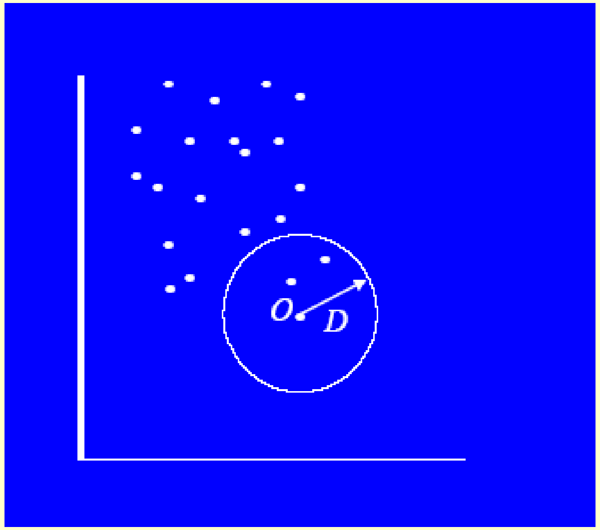

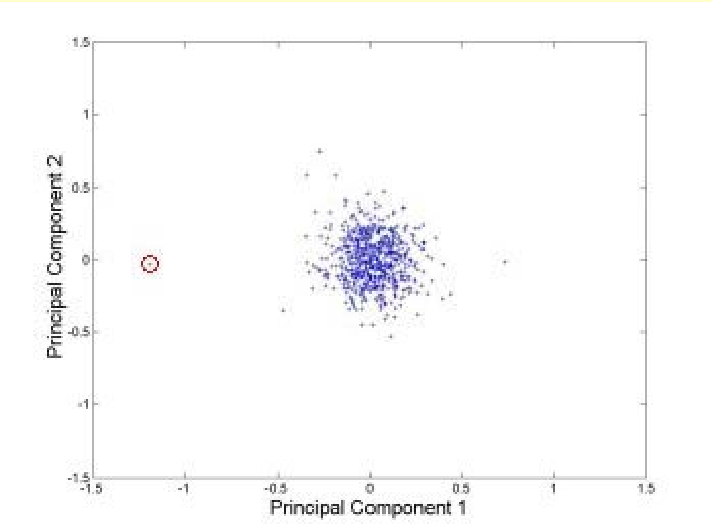

DATA TRANSFORMATION

Data Transformation

- Smoothing: remove noise from data (binning, regression, clustering)
- Aggregation: summarization, data cube construction
- Generalization: concept hierarchy climbing
- Attribute/feature construction
  - New attributes constructed from the given ones (add att. area which is based on height and width)
- Normalization
  - Scale values to fall within a smaller, specified range

Data Cube Aggregation

- Data can be aggregated so that the resulting data summarize, for example, sales per year instead of sales per quarter.

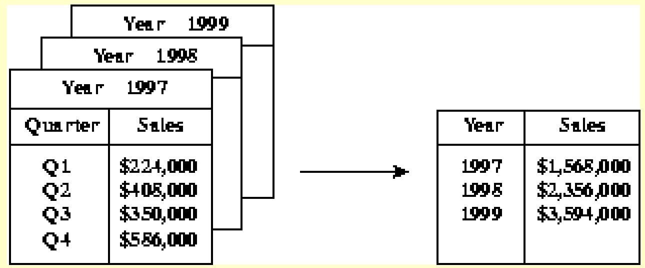

- Reduced representation which contains all the relevant information if we are concerned with the analysis of yearly sales

Concept Hierarchies

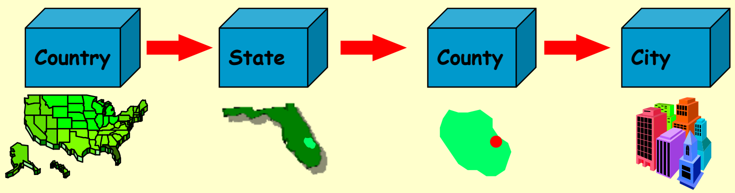

Jobs, food classification, time measures...

Normalization

- For distance-based methods, normalization helps to prevent that attributes with large ranges out-weight attributes with small ranges
  - min-max normalization
  - z-score normalization
  - normalization by decimal scaling

Normalization

- min-max normalization

In R
mmnorm(data,minval=0,maxval=1)

$$v' = \frac{v - min_v}{max_v - min_v}(new\_max_v - new\_min_v) + new\_min_v$$

- z-score normalization

$$v' = \frac{v - \bar{v}}{\sigma_v}$$

In R
boxplot(znorm(cal_housing[,3:8]))

- normalization by decimal scaling

$$v' = \frac{v}{10^j}$$

Where j is the smallest integer such that Max(| v' |)<1

range: -986 to 917 => j=3 -986 -> -0.986 917 -> 0.917

MISSING DATA

Missing Data

- Data is not always available
  - E.g., many tuples have no recorded value for several attributes, such as customer income in sales data
- Missing data may be due to
  - equipment malfunction
  - inconsistent with other recorded data and thus deleted
  - data not entered due to misunderstanding
  - certain data may not be considered important at the time of entry
  - not register history or changes of the data
- Missing data may need to be inferred.
- Missing values may carry some information content: e.g. a credit application may carry information by noting which field the applicant did not complete

Missing Values

- There are always MVs in a real dataset
- MVs may have an impact on modelling, in fact, they can destroy it!
- Some tools ignore missing values, others use some metric to fill in replacements
  - The modeller should avoid default automated replacement techniques
    - Difficult to know limitations, problems and introduced bias
- Replacing missing values without elsewhere capturing that information removes information from the dataset

How to Handle Missing Data?

- Ignore records (use only cases with all values)
  - Usually done when class label is missing as most prediction methods do not handle missing data well
  - Not effective when the percentage of missing values per attribute varies considerably as it can lead to insufficient and/or biased sample sizes
- Ignore attributes with missing values
  - Use only features (attributes) with all values (may leave out important features)
- Fill in the missing value manually
  - tedious + infeasible?

How to Handle Missing Data?

- Use a global constant to fill in the missing value
  - e.g., “unknown”. (May create a new class!)
- Use the attribute mean to fill in the missing value
  - It will do the least harm to the mean of existing data
  - If the mean is to be unbiased
  - What if the standard deviation is to be unbiased?
- Use the attribute mean for all samples belonging to the same class to fill in the missing value

How to Handle Missing Data?

- Use the most probable value to fill in the missing value
  - Inference-based such as Bayesian formula or decision tree
  - Identify relationships among variables
    - Linear regression, Multiple linear regression, Nonlinear regression
  - Nearest-Neighbour estimator
    - Finding the k neighbours nearest to the point and fill in the most frequent value or the average value
    - Finding neighbours in a large dataset may be slow

Nearest-Neighbour

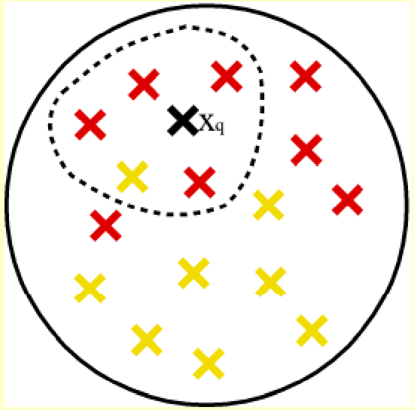

How to Handle Missing Data?

- Note that, it is as important to avoid adding bias and distortion to the data as it is to make the information available.
  - bias is added when a wrong value is filled-in
- No matter what techniques you use to conquer the problem, it comes at a price. The more guessing you have to do, the further away from the real data the database becomes. Thus, in turn, it can affect the accuracy and validation of the mining results.

Missing Data with R

- library(dprep)
- data(hepatitis) - loads dataset
- str(hepatitis) - gives dataset structure
- summary(hepatitis)
- short_hep <- hepatitis[1:15,]
- ?ce.impute - gives information about the fill missing values method
- res <- ce.impute(short_hep,"median",19)
- ?clean()ce.impute(hepatitis,"median",1:19)
- ce.impute(hepatitis,"knn",k1=10)
- clean() – eliminates rows and columns that have more than the set limit missings
- clean(res,0.3,0.2)
- imagmiss(hepatitis) – gives the percentage of missing values

DATA INTEGRATION

Data Integration

- Turn a collection of pieces of information into an integrated and consistent whole
- Detecting and resolving data value conflicts
  - For the same real world entity, attribute values from different sources may be different
    - Which source is more reliable ?
    - Is it possible to induce the correct value?
  - Possible reasons: different representations, different scales, e.g., metric vs. British units

Data integration requires knowledge of the "business"

Types of Inter-schema Conflicts

- Classification conflicts
  - Corresponding types describe different sets of real world elements.
    DB1: authors of journal and conference papers;
    DB2 authors of conference papers only.
    - Generalization / specialization hierarchy
- Descriptive conflicts
  - naming conflicts : synonyms , homonyms
  - cardinalities: first name - one , two , N values
  - domains: salary : $, Euro ... ; student grade : [0 : 20] , [1 : 5]
    - Solution depends upon the type of the descriptive conflict

Data type inconsistency example

- **1999 Sep 23**
  The $125 million Mars Climate Orbiter was presumed lost after it hit the Martian atmosphere. The crash was later blamed on navigation confusion due to 2 teams using conflicting English and metric units.
- http://en.wikipedia.org/wiki/Mars_Climate_Orbiter

Types of Inter-schema Conflicts

- Structural conflicts
  - DB1 : Book is a class; DB2 : books is an attribute of Author
    - Choose the less constrained structure (Book is a class)
- Fragmentation conflicts
  - DB1: Class Road_segment ; DB2: Classes Way_segment , Separator
    - Aggregation relationship

Handling Redundancy in Data Integration

- Redundant data occur often when integrating databases
  - The same attribute may have different names in different databases
  - False predictors are fields correlated to target behavior, which describe events that happen at the same time or *after* the target behavior
    - Example: Service cancellation date is a leaker when predicting attriters
  - One attribute may be a "derived" attribute in another table, e.g., annual revenue
  - For numerical attributes, redundancy may be detected by correlation analysis

$$r_{XY} = \frac{\frac{1}{N-1}\cdot\sum_{n=1}^{N}(x_n-\bar{x})\cdot(y_n-\bar{y})}{\sqrt{\frac{1}{N-1}\cdot\sum_{n=1}^{N}(x_n-\bar{x})^2}\cdot\sqrt{\frac{1}{N-1}\cdot\sum_{n=1}^{N}(y_n-\bar{y})^2}} \quad (-1 \le r_{XY} \le 1)$$

Scatter Matrix

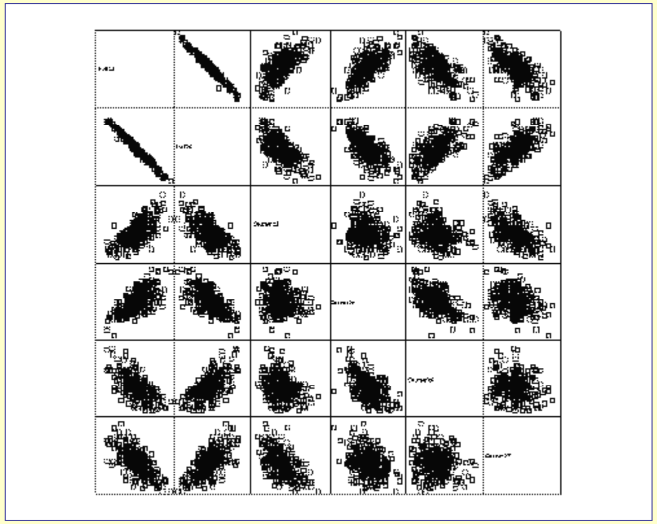

(Almost) Automated False Predictor Detection

- For each field
  - Build 1-field decision trees for each field
  - (or compute correlation with the target field)
- Rank all suspects by 1-field prediction accuracy (or correlation)
- Remove suspects whose accuracy is close to 100% (Note: the threshold is domain dependent)
- Verify top “suspects” with domain expert

DATA REDUCTION

Data Reduction

- Selecting Most Relevant Attributes
  - If there are too many attributes, select a subset that is most relevant (according to your knowledge of the business).
  - Select top N fields using 1-field predictive accuracy as computed for detecting false predictors.
- Attribute Numerosity Reduction
  - Parametric methods
    - Assume the data fits some model, estimate model parameters, store only the parameters, and discard the data (except possible outliers), Regression
  - Non-parametric methods
    - Do not assume models
    - Major families: histograms, clustering, sampling

Clustering

- Partition a data set into clusters makes it possible to store cluster representation only
- Can be very effective if data is clustered but not if data is "smeared"
- There are many choices of clustering definitions and clustering algorithms, further detailed in next lessons

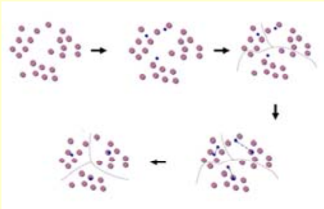

Histograms

- A popular data reduction technique
- Divide data into buckets and store average (sum) for each bucket
- Can be constructed optimally in one dimension using dynamic programming:
  - Optimal histogram has minimum variance. Hist. variance is a weighted sum of the variance of the source values in each bucket.

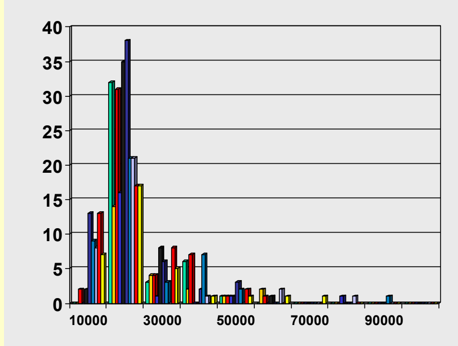

Increasing Dimensionality

- In some circumstances the dimensionality of a variable need to be increased:
  - Color from a category list to the RGB values
  - ZIP codes from category list to latitude and longitude

Sampling

- The cost of sampling is proportional to the sample size and not to the original dataset size, therefore, a mining algorithm's complexity is potentially sub-linear to the size of the data
- Choose a representative subset of the data
  - Simple random sampling (SRS) (with or without reposition)
  - Stratified sampling:
    - Approximate the percentage of each class (or subpopulation of interest) in the overall database
    - Used in conjunction with skewed data

Unbalanced Target Distribution

- Sometimes, classes have very unequal frequency
  - Attrition prediction: 97% stay, 3% attrite (in a month)
  - medical diagnosis: 90% healthy, 10% disease
  - eCommerce: 99% don't buy, 1% buy
  - Security: >99.99% of Americans are not terrorists
- Similar situation with multiple classes
- Majority class classifier can be 97% correct, but useless

Handling Unbalanced Data

- With two classes: let positive targets be a minority
- Separate raw held-aside set (e.g. 30% of data) and raw train
  - put aside raw held-aside and don't use it till the final model
- Select remaining positive targets (e.g. 70% of all targets) from raw train
- Join with equal number of negative targets from raw train, and randomly sort it.
- Separate randomized balanced set into balanced train and balanced test

Building Balanced Train Sets

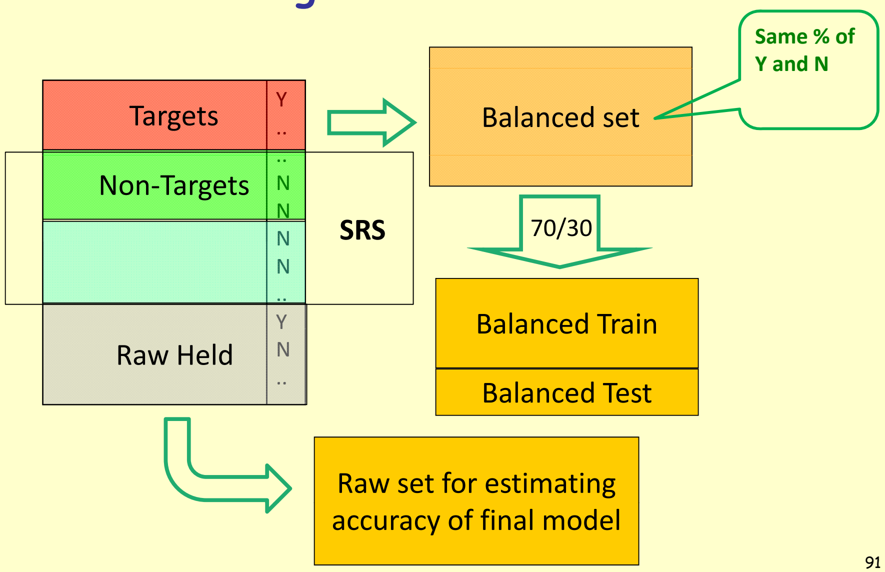

Summary

- Every real world data set needs some kind of data pre-processing
  - Deal with missing values
  - Correct erroneous values
  - Select relevant attributes
  - Adapt data set format to the software tool to be used
- In general, data pre-processing consumes more than 60% of a data mining project effort

References

- 'Data preparation for data mining', Dorian Pyle, 1999
- 'Data Mining: Concepts and Techniques', Jiawei Han and Micheline Kamber, 2000
- 'Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations', Ian H. Witten and Eibe Frank, 1999
- 'Data Mining: Practical Machine Learning Tools and Techniques second edition', Ian H. Witten and Eibe Frank, 2005
- DM: Introduction: Machine Learning and Data Mining, Gregory Piatetsky-Shapiro and Gary Parker (http://www.kdnuggets.com/data_mining_course/dm1-introduction-ml-data-mining.ppt)
- ESMA 6835 Mineria de Datos (http://math.uprm.edu/~edgar/dm8.ppt)

Thank you !!!