



DMIF, University of Udine

Data Management and Analysis with Business Applications

A Brief Introduction to Data Mining

Andrea Brunello

andrea.brunello@uniud.it

24th May 2020



- 1 What is Data Mining
- 2 Types of Learning

What is Data Mining

Data \approx stored events/facts.

Information can be considered as the set of concepts, patterns, regularities that are hidden in the data.

Data Mining is the task by which useful, previously unknown information can be extracted from (possibly large) quantities of data.

- > It is a process of abstraction, that leads to the definition of a *model*.

Machine Learning represents the “technical basis” of Data Mining.



What are Patterns Good for?

The models that capture the patterns can be used to:

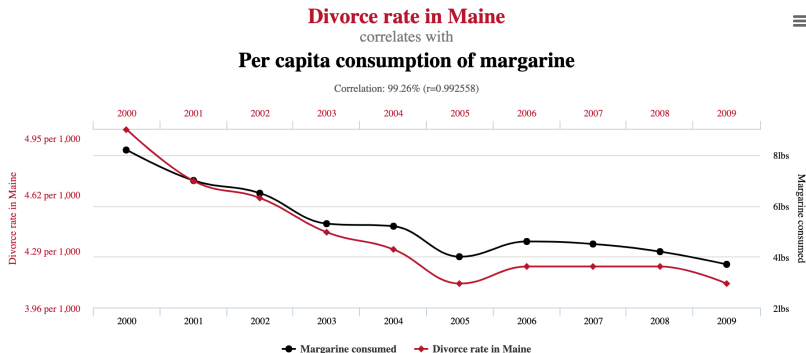
- **know:** that some population groups are more likely to buy a specific good
- **explain:** what are the reasons behind customer churn
- **predict:** whether an increase in advertising budget will bring to more sales

Sometimes, goals may overlap. For instance, think about a model that gives the value of a house based on a series of its characteristics.



Caveats

Sometimes, the discovered patterns may be trivial, produced by random correlation, or simply wrong.



Data sources: National Vital Statistics Reports and U.S. Department of Agriculture

tylervigen.com

<https://www.tylervigen.com/spurious-correlations>

To summarize:

- Data Mining is a task that relies on Machine Learning
- to (semi-)automatically extract
- information, useful patterns
- from (possibly large) quantities of data

Input of the process:

- instances, examples of the concepts that you want to learn

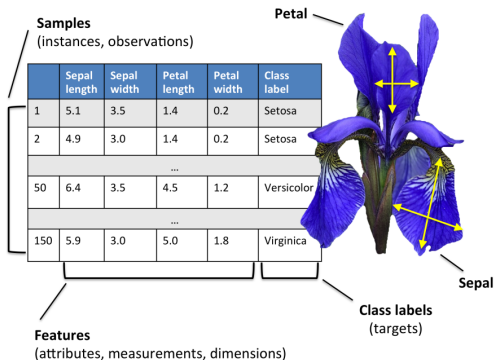
Output of the process:

- predictions
- models

Types of Learning

We will consider tabular datasets, i.e.,

- each row corresponds to an instance
- each column corresponds to a characteristic (feature)
- there may be a column with a special role (label)





A Short Taxonomy of Learning

We may identify the following, main, categories of learning:

- Supervised Learning:
 - Classification tasks
 - Regression tasks
- Unsupervised Learning:
 - Association Rule Discovery
 - Clustering
 - ...

Each instance in the dataset is characterized by a set of categorical or numerical features that are used as predictors to determine the value of a specific label.

Given a training dataset of instances, each with feature values $x_1, x_2, \dots, x_n \in X_1 \times X_2 \times \dots \times X_n$ and a label value $l \in L$, we want to learn a function $f : X_1 \times X_2 \times \dots \times X_n \rightarrow L$, such that:

$$f(x_1, \dots, x_n) = \hat{l} \approx l$$

Function f is encoded into a model, that can be used to predict the value of l for new instances.



In classification tasks, the label l is categorical, thus its domain of values is discrete and finite. For instance, a set of colors, topics, ...

Classical models:

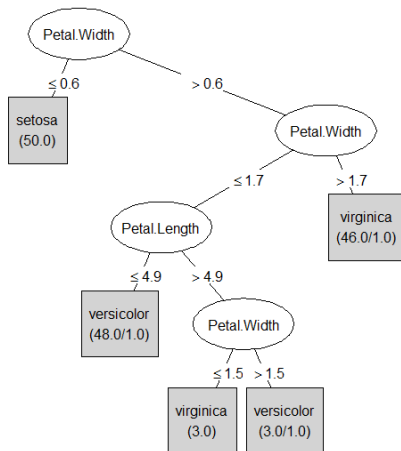
- decision trees and their ensembles
- logistic regression
- naive bayes classifier
- support vector machines

Exemplary tasks:

- text/image/video classification
- credit card fraud detection
- customer churn prediction

Decision Tree Example

J48 decision tree with 98% accuracy on the *Iris* dataset (using 10-fold cross-validation).





In regression tasks, the label l is numerical, thus its domain is continuous. For instance, real estate values, probability of a failure, ...

Classical models:

- linear regression
- decision tree ensembles
- support vector regression

Exemplary tasks:

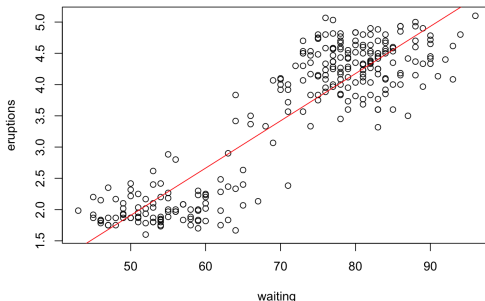
- predictive maintenance
- sentiment analysis
- revenue forecasting



Linear Regression Example

Dataset *faithful*, recordings about the Old Faithful geyser in Yellowstone National Park.

Eruption duration	Waiting time
2.883	55
1.883	54
1.600	52
1.750	47





We are given a dataset of instances, each one with feature values $x_1, x_2, \dots, x_n \in X_1 \times X_2 \times \dots \times X_n$.

There is no label, the goal here is to look for any kind of interesting pattern that can be found among the features.

Still, the output of the process can be considered a model, that encodes such relationships between the features.



The goal is that of discovering “interesting” relations between features in a large dataset.

For instance, the rule $\{onions, potatoes\} \Rightarrow \{burger\}$ found in the sales data of a supermarket would indicate that if a customer buys onions and potatoes together, they are likely to also buy hamburger meat.

Such information can be used as the basis for decisions about activities such as promotional pricing or product placements.

Many algorithms to mine association rules have been presented in the literature. Historically, the most important one is *Apriori* (Agrawal and Srikant, 1994).



Clustering is the task of grouping a set of instances in such a way that objects in the same group (cluster) are more similar to each other than to those in other groups.

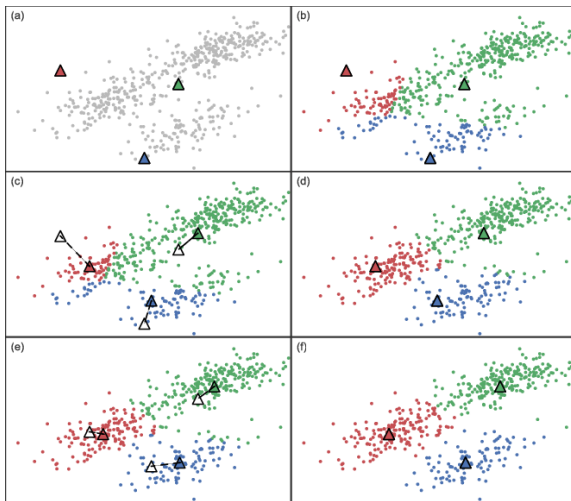
Similarity calculation relies on metrics (e.g., euclidean distance) that are applied on the instances' features.

Many kinds of clustering: soft vs hard, hierarchical vs partitional, ...

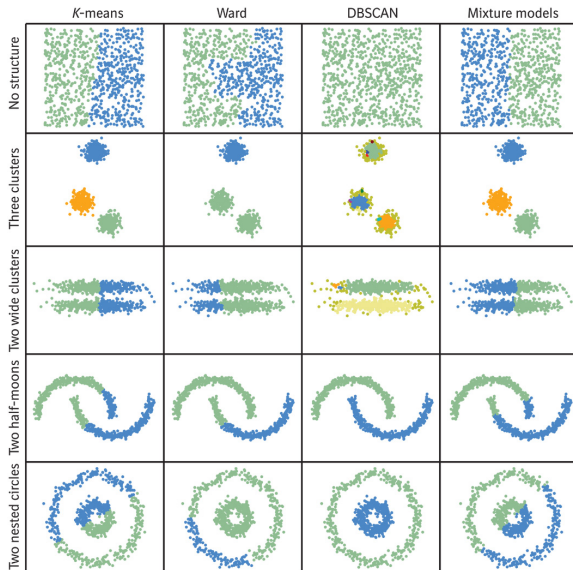
Useful, for instance, to perform customer segmentation.

A popular, partitional clustering algorithm is *K-Means*.

K-Means Example



Clustering is a Hard Task!





M. Hall, I. H. Witten, E. Frank, C. J. Pal, *Data Mining: Practical Machine Learning Tools and Techniques*, 4th Edition, 2016.

R. Tibshirani, T. Hastie, *An Introduction to Statistical Learning*, 2nd Edition, 2009.

F. Chollet, *Deep Learning with Python*, 2017.