COMP338: **ARTIFICIAL INTELLIGENCE**

Machine Learning - Introduction

Dr. Radi Jarrar
Department of Computer Science

What is Machine Learning?

- The task of building knowledge and storing it in some form in the computer (such as algorithm or mathematical model) that can help detect patterns
- It is concerned with predicting a particular outcome given some data
- Training set is used to train the algorithm
- The algorithm, typically, has a number of parameters that are learnt from the data

Machine Learning Applications

Image classification	Text translation	Medical Diagnosis	Stock market prediction
Online fraud detection	Email spam filter	Self driving cars	Product recommendation
Traffic prediction	Speech recognition	Face detection	Face recognition

What is Machine Learning?

- Machine Learning exists at the intersection of Mathematics & Statistics with Software Engineering & Computer

- **Machine Learning** is concerned with *teaching* computers something about the world, so that they can think more clearly about the world in order to make better decisions

Machine Learning

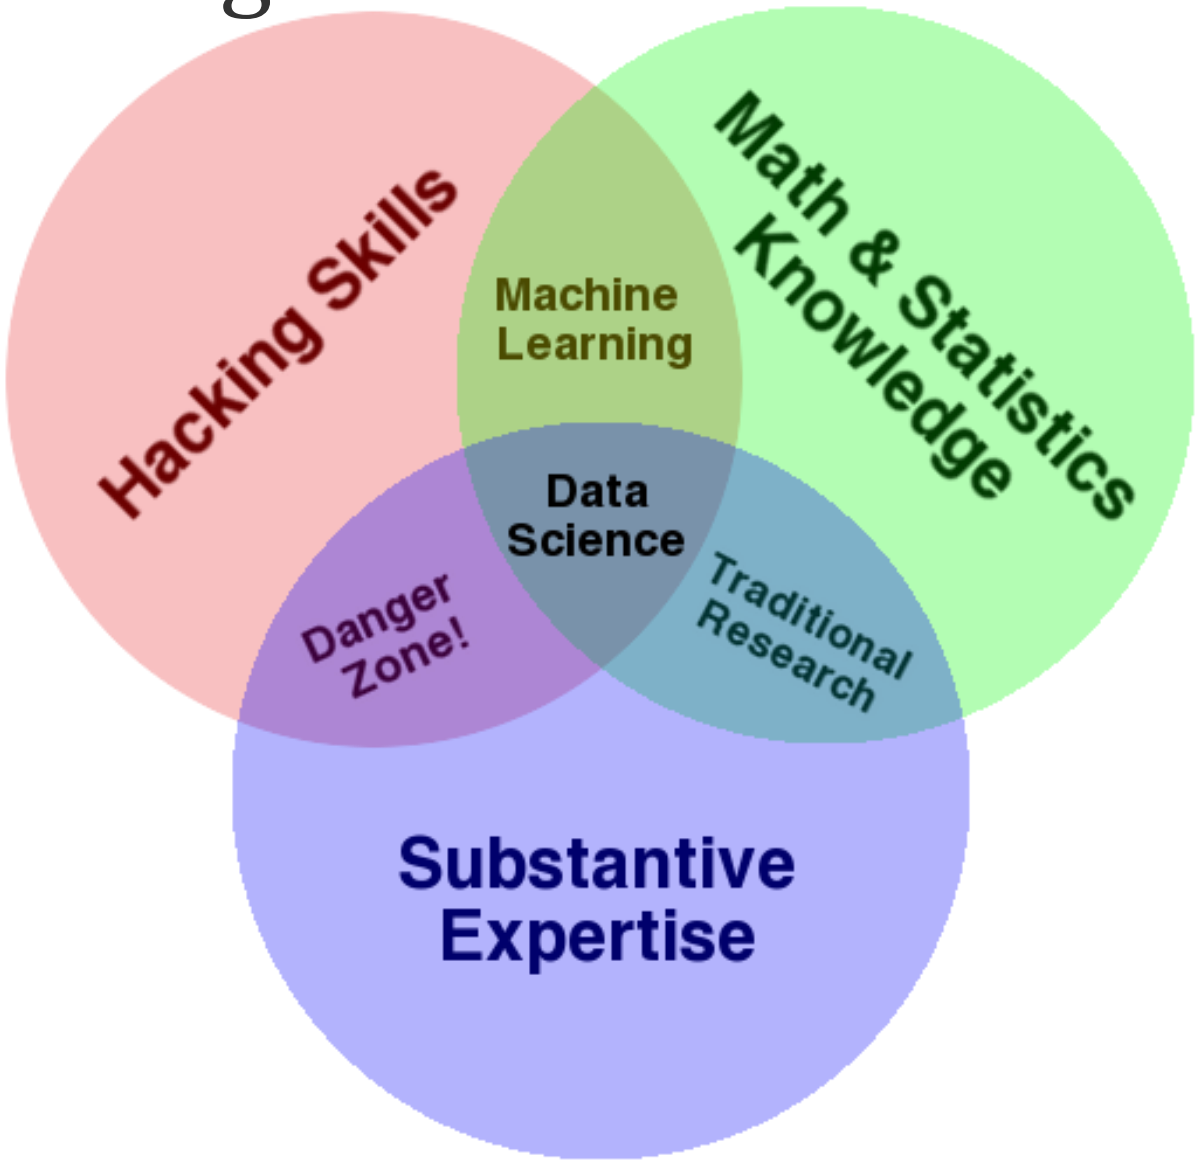

ML & Artificial Intelligence (AI)

- AI is the broad family of Machine Learning
- AI is the study of how to create intelligent agents
- How to program the computer to behave and perform as a intelligent agent (i.e., a human)
- This may not involve training or learning from data!

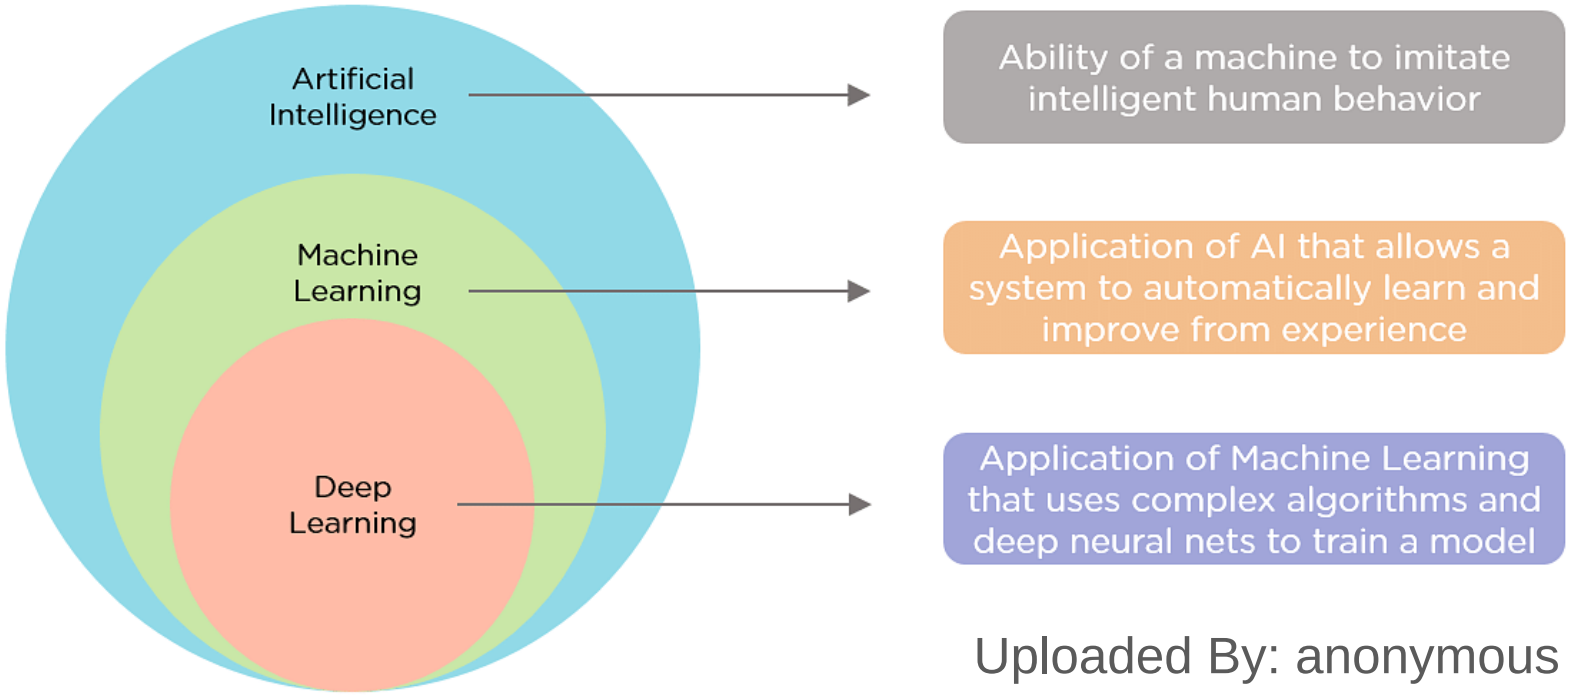

ML & Statistics

- Statistics is concerned with learning something interpretable from data; whereas Machine Learning is concerned with turning data into something practical and usable
- Machine learning is concerned with teaching *computers* something about the world, so that they can use that knowledge to perform other tasks. Statistics is more concerned with developing tools for teaching *humans* something about the world, so that they can think more clearly about the world in order to make better decisions.

ML & Data Mining

- Often confused!
- Machine learning can be seen as a pre-requisite for Data mining
- Machine Learning focuses on prediction based on ‘known’ properties learnt from the training data
- Data Mining is the identification of correlations and patterns within data. It focuses on discovery of ‘unknown’ properties and patterns of data

Deep Learning

- Deep learning is a machine learning technique that teaches computers by imitating the human thinking
- Meaning, do what comes naturally to humans (learn by example)
- Deep learning has come into widespread and is a key technology behind driverless cars, voice control in consumer devices like phones, tablets, etc.

Deep Learning

- Deep learning requires large size of training data and substantial computing power
- Examples of Deep learning applications
  - Self-driving cars
  - Automatic machine translation
  - Automatic handwriting generation
  - Colourisation of black/white images
  - Pixel restoration

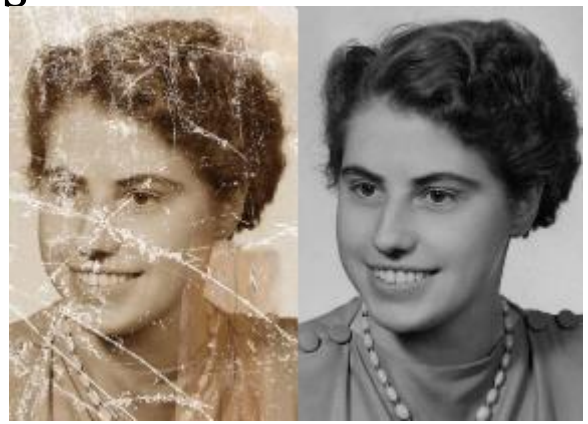

Deep Learning vs. Machine Learning

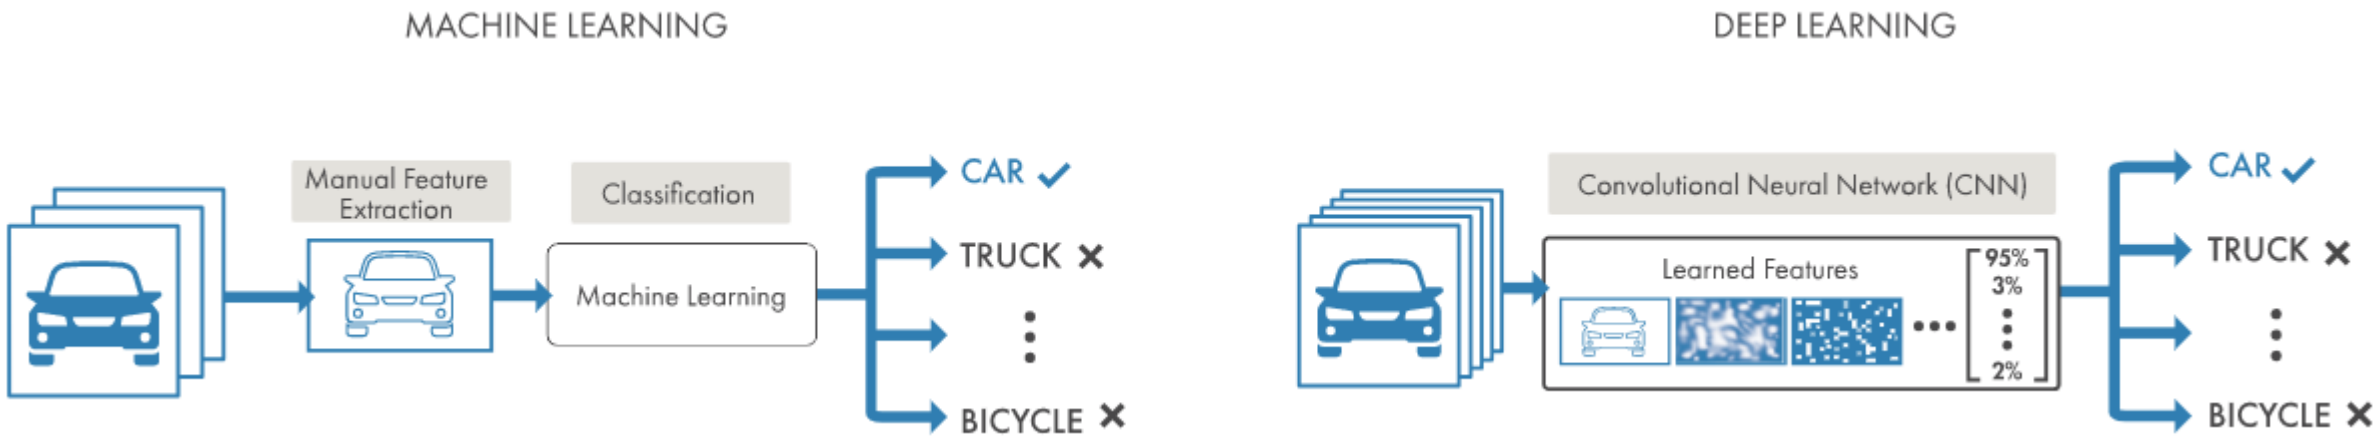

FORMULATING MACHINE LEARNING PROBLEMS

Example: Handwritten digit recognition

- Input: images of handwritten digits
- Output: the digit in images
- Setup:
  - A large collection of image examples ‘labeled’ with the correct numbers
- Task
  - Predict on new images with their numbers
- Features
  - Image pixels!

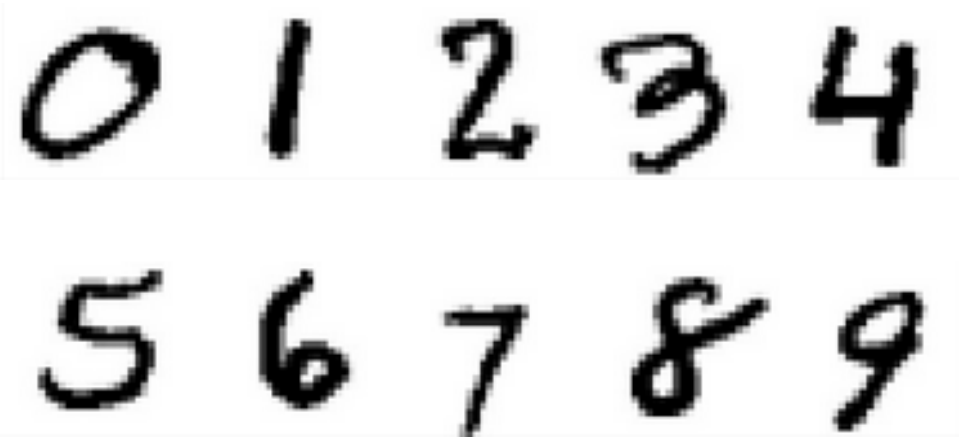

Example: Handwritten digit recognition

- Each image is
  15 x 15 pixels

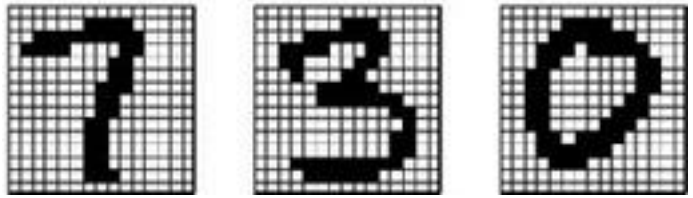

- The input ‘features’
  of each image can be
  represented as a vector $x \in \mathcal{R}^{225}$

- The classification problem is
  a function f(x) such that
  $f : x \rightarrow \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$

Example: Face detection

- Detect faces within images
- Detection might specify frontal or side-face

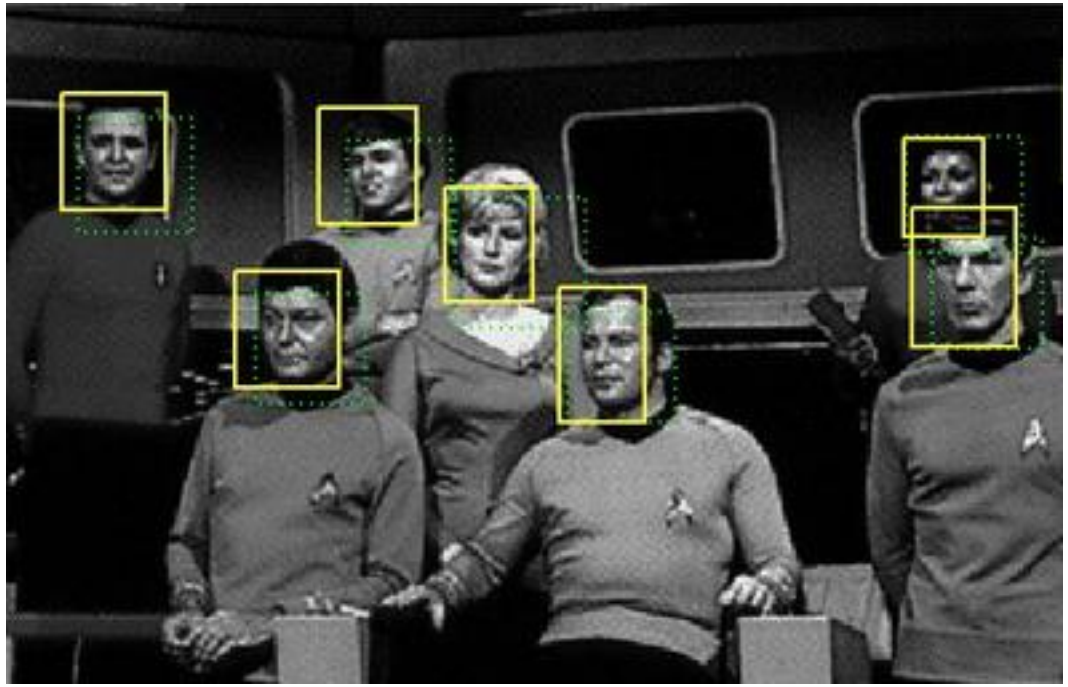

Example: Text classification

- Classify text documents into pre-defined categories
- Documents might be classified into a single or multiple categories
- Helps in
  - Archiving old texts
  - Classifying news
  - …

Example: Stock price prediction

- Predict stock price at future
- This task is a bit different as the output is continuous valued (rather than category or discrete value); i.e., Regression

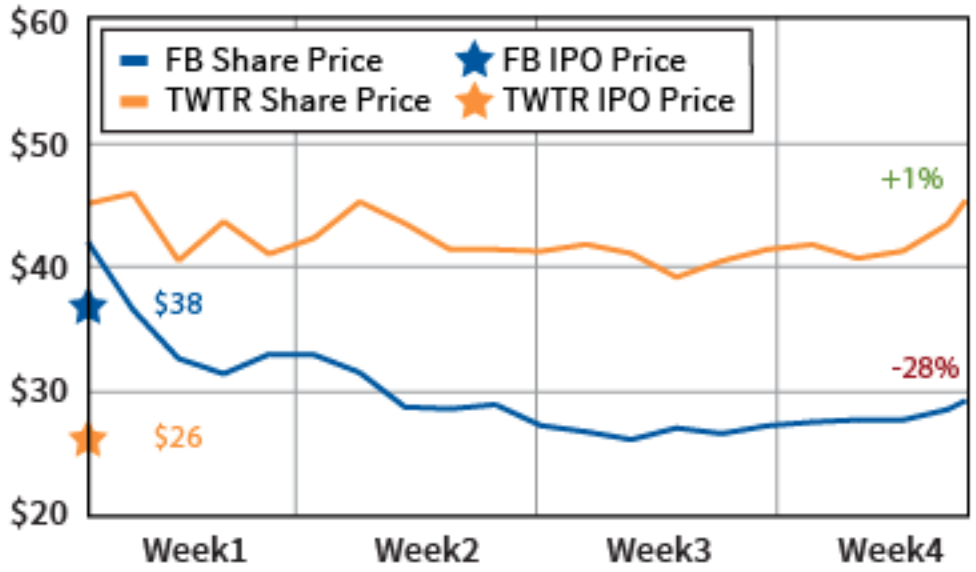

Example: Image classification

- Classify images based on their contents (e.g., indoor vs. outdoor)
- Training annotated images are used to generate the classifier

Example 7: Recommendation systems

- Give users suggestions on items based on previous data (history) from other users (or the same user!)

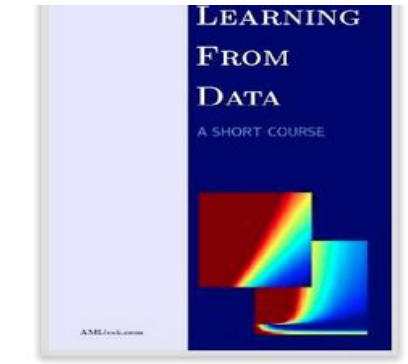

See all 2 images

See all formats and editions

Hardcover
from $28.00

11 Used from $40.00
6 New from $28.00

This book, together with specially prepared online material freely accessible to our readers, provides a complete introduction to Machine Learning, the technology that enables computational systems to adaptively improve their performance with experience accumulated from the observed data. Such techniques are widely applied in engineering, science, finance, and commerce. This book is designed for a short course on machine learning. It is a short course, not a hurried course. From over a decade of teaching this material, we have distilled what we believe to be the core topics that every student of the

Read more

Customers Who Bought This Item Also Bought

Page 2 of 25 | Start over

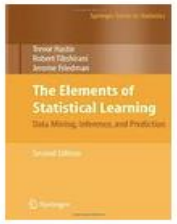

The Elements of Statistical

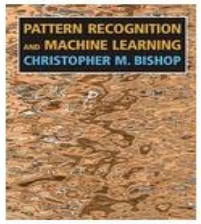

Pattern Recognition and

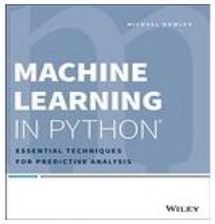

Machine Learning in

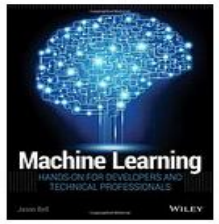

Machine Learning: Hands-

FEATURES AND TYPES OF LEARNINGS

Features

- Machine learning algorithms receives data and it manages to classify data
- There is a “garbage in; garbage out” aspect to machine learning
- Meaning, the data should be good, clean, and representative in order to make the algorithm learn how to apply hypothesis and learn data

Features

- Feature engineering is crucial to machine learning

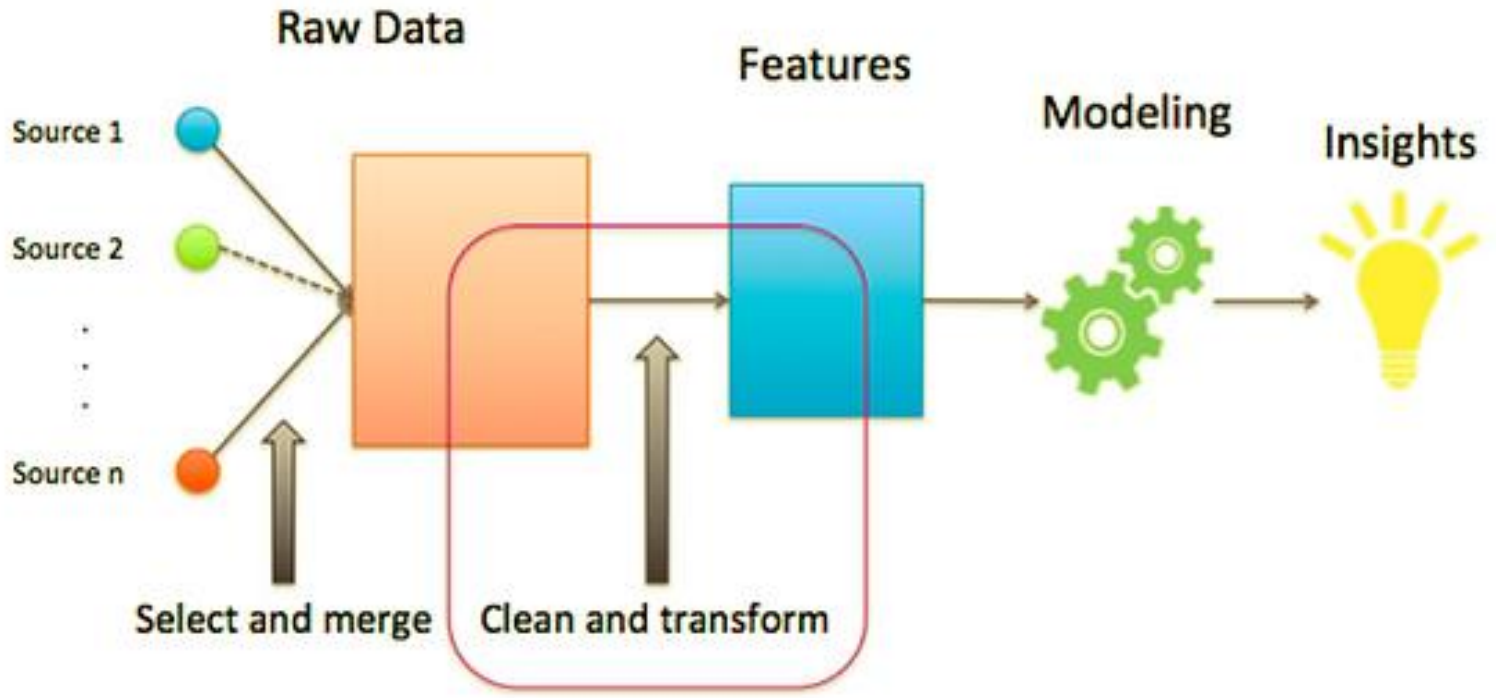

https://www.datainsightonline.com/post/feature-engineering-in-machine-learning

Features

- Can be seen as the language that we use to describe certain object
- E.g., email, image, historical stock information,…
- Through proper methods/algorithms, features are extracted from the objects from which the learner is being trained
- They determine the much of the success of the model: a model is only as good as its features

Features

- E.g., imagine you have a 100x100 pixel image, an easy feature representation of the image could be a 30,000 dimensional feature vector
- Each dimension corresponds to the red, blue, or green component of some pixel in the image
- The first element in the vecto could be the red element of the first vector, the second is the blue of the same pixel, and so on
- This is the pixel representation of the image

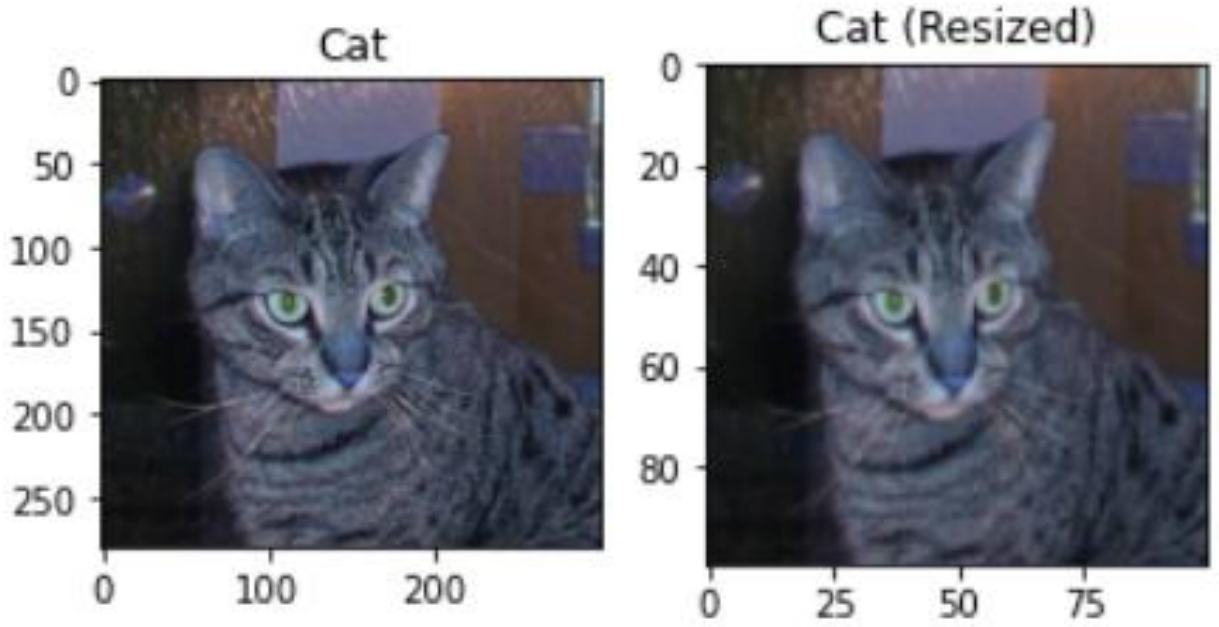

Features

- Text (as well images) can be represented as Bags-of-Words, which considers documents as collection words regardless of their location in the document
- The terms that occur the highest will be considered as features and their values is how many times those words occurred in a document

Features

- Missing values is considered a problem when dealing with ML
- Some learner do not accept missing values
- Easiest method is to remove rows with missing features (however, this might remove some other good values of other features!)
- Impute missing values:
  - Add a constant value (such as 0) to distinguish from other features
  - Add a value from another randomly selected record
  - Add the mean, median or the mode value for the column
  - Use algorithms that support Missing values!

Features

- Another factors to consider when working with features including Normalisation, Scaling, and Pruning

Types of Machine Learning

Supervised	Unsupervised	Reinforcement
• Task driven • Data with defined output are provided • Classification (image classification) and regression (forecasting)	• Data driven • Machine understands data (identifies patterns/structures) • Clustering (customer segmentation)	• Algorithms learns to react to an environment • Game AI (learn from mistakes) and Robot navigation

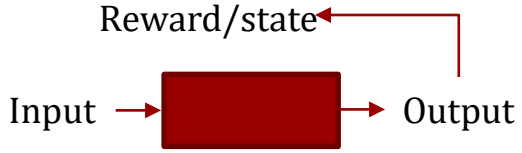

Supervised learning

- A model creation process where the model describes the relationship between a set of features (attributes) and a predefined variable called target class (also called label)

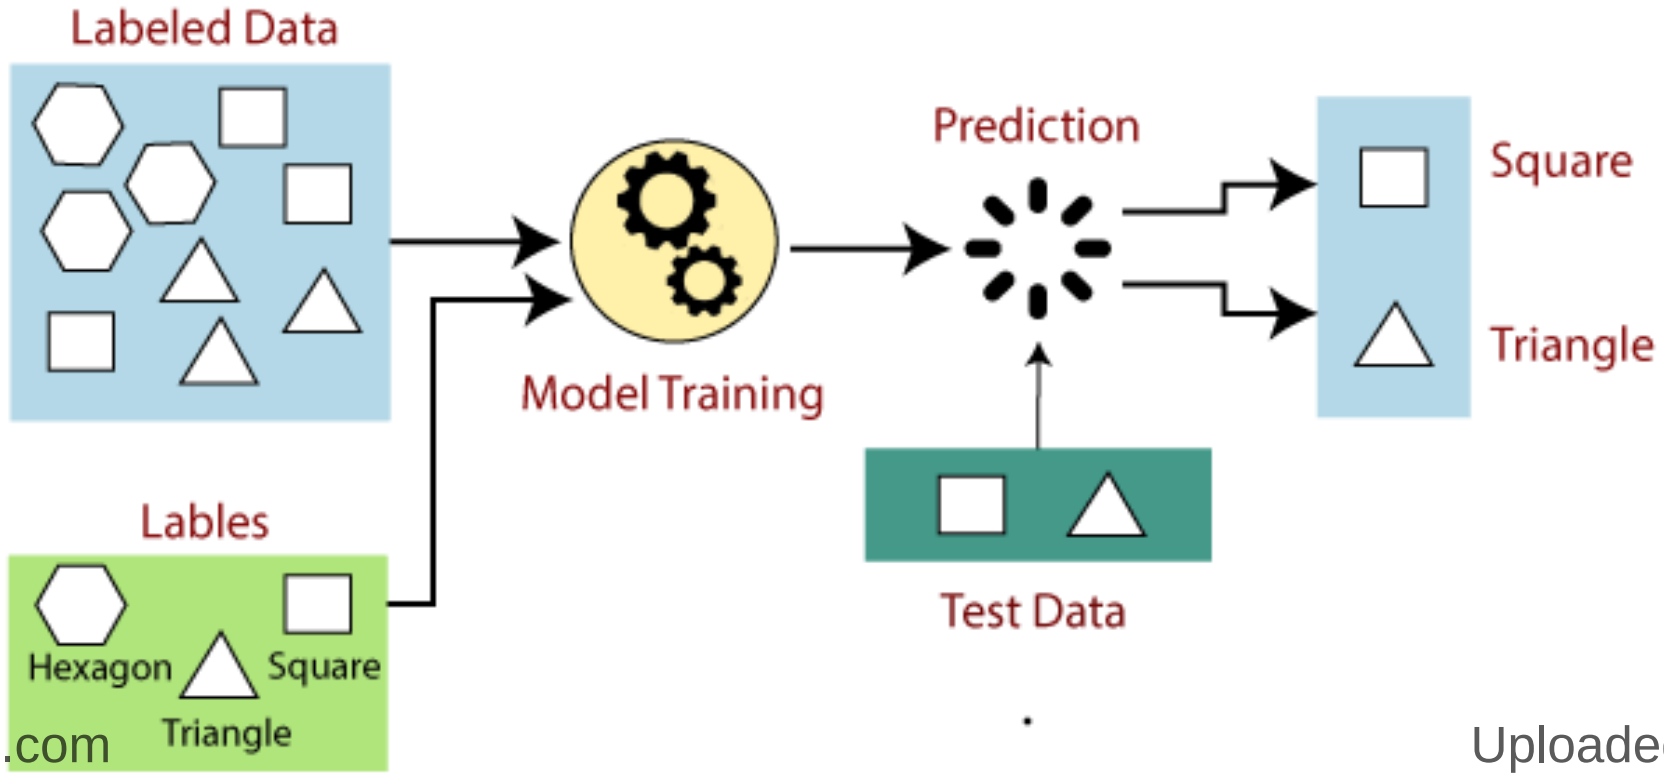

Classification and Regression

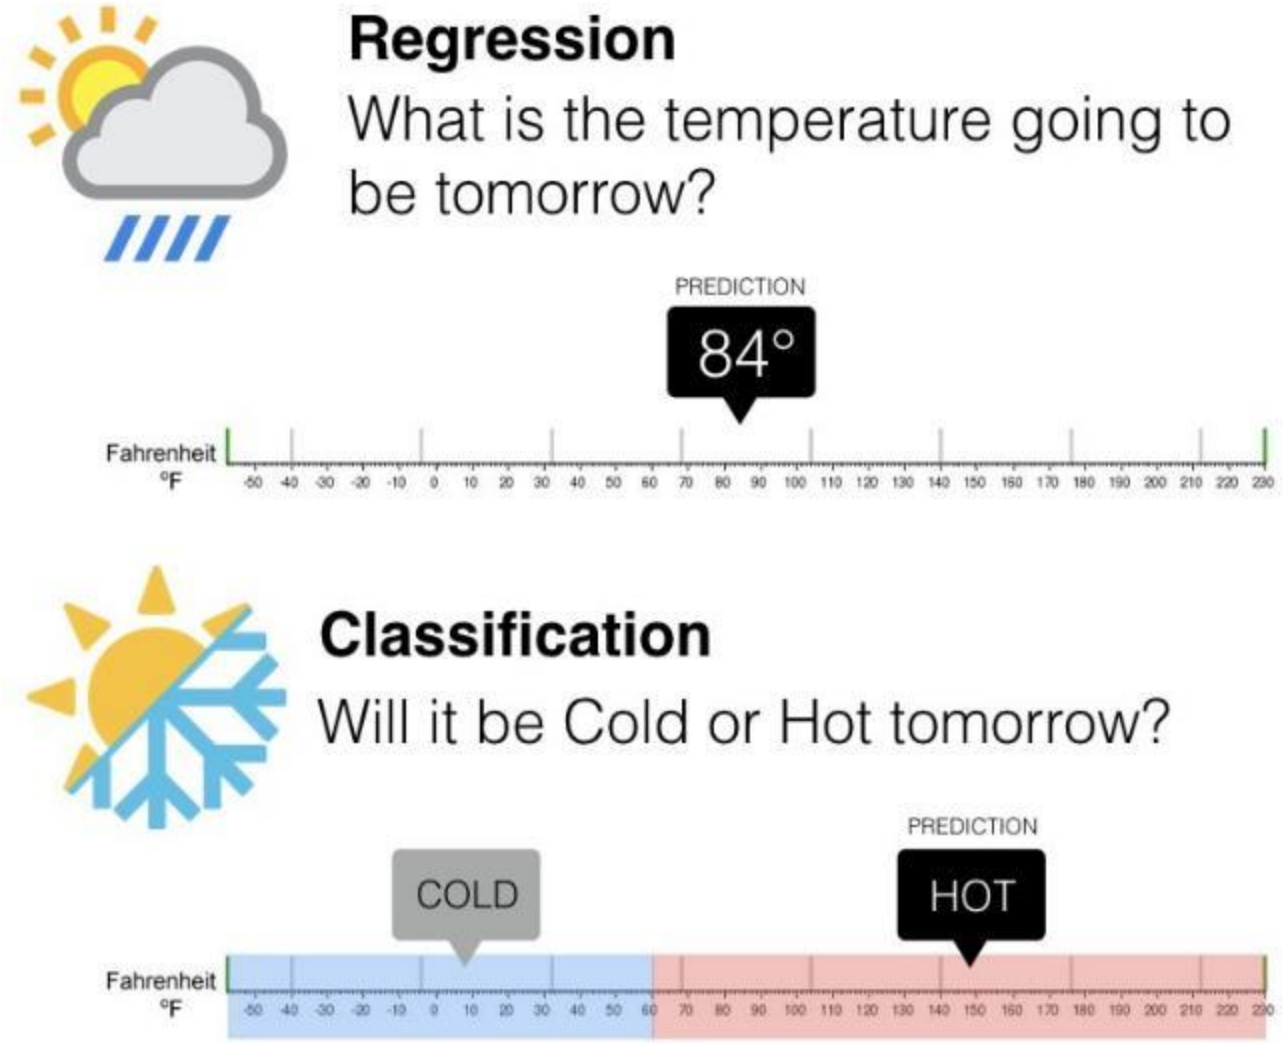

https://medium.com/@ali_88273/regression-vs-classification-87c224350d69

Classification and Regression

- Classification: the task of classifying some input features into classes (categories)
- As the name implies, the output of classification is categorical (‘class’ name)
- Image classification, Face detection and recognition, Handwritten recognition, Speech recognition, Document classification, …

Classification and Regression

- Regression: the output variables of a regression problem is a continuous value (i.e., predict the price of stocks in the future)

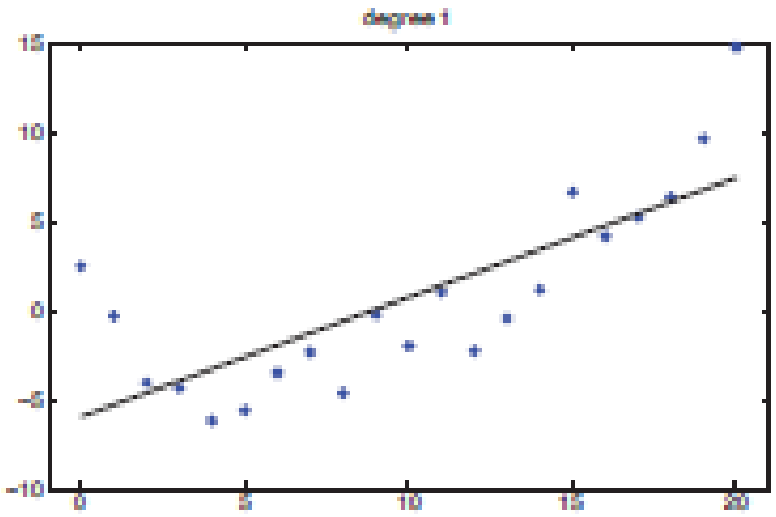

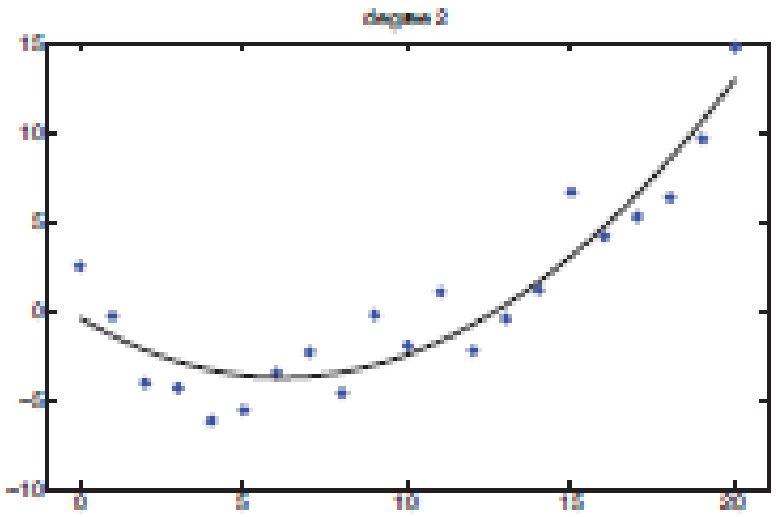

Classification and Regression

- Examples on regression:
  - Predict the stock market given current market conditions
  - Predict the prices of real-estate given some data about it
  - Predict the age of a viewer watching a video on YouTube
  - Predict the weather given some measures (temperature, humidity, wind, …)
  - Predict the temperature inside a building using weather data, time, door sensors, …

Model

- A Model is a simplified representation of reality created to serve a purpose
- Simplified based on assumptions of what is and what is not important for the specific purpose
- Sometimes based on constraints on information

Binary classification vs. Multi-class classification

- The outcome of a binary classification problem is two classes
- Usually the classes are encoded as {0, 1}
- A multiclass classification problem has many target classes
- For example, image scene classification problem where the outcome of the classification is a class label of the scene

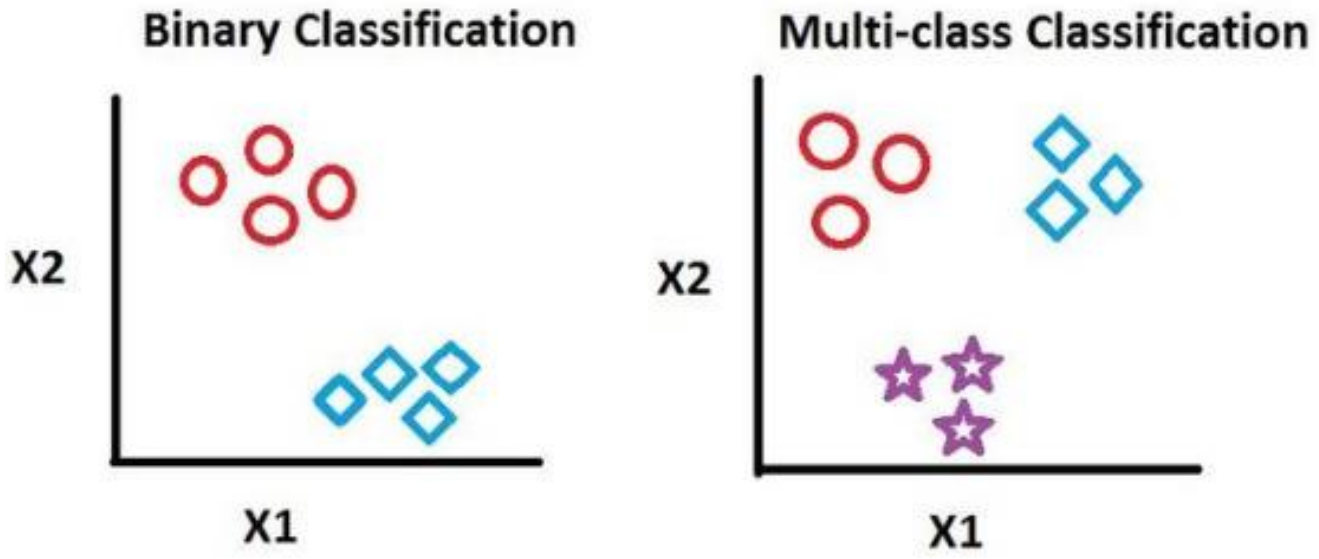

Unsupervised learning

- In unsupervised learning, the training data consists of input vectors without any corresponding target values
- The goal might be, for instance, to discover similar examples within the data (clustering such as K-means algorithm)
- Discover patterns/structures hidden inside the data
- This is referred to as Knowledge Discovery

Unsupervised learning

- The tasks here are generally formulated as density estimations in order to build the models
- More widely applicable than supervised learning as there is not required human annotation of the data as in supervised learning

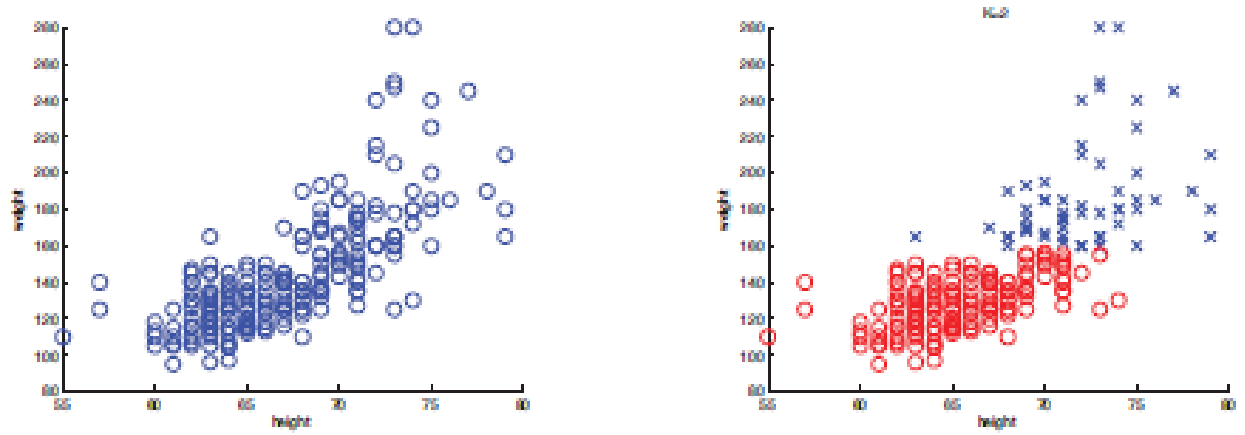

(a) The height and weight of some people. (b) A possible clustering using $K = 2$ clusters.

Clustering

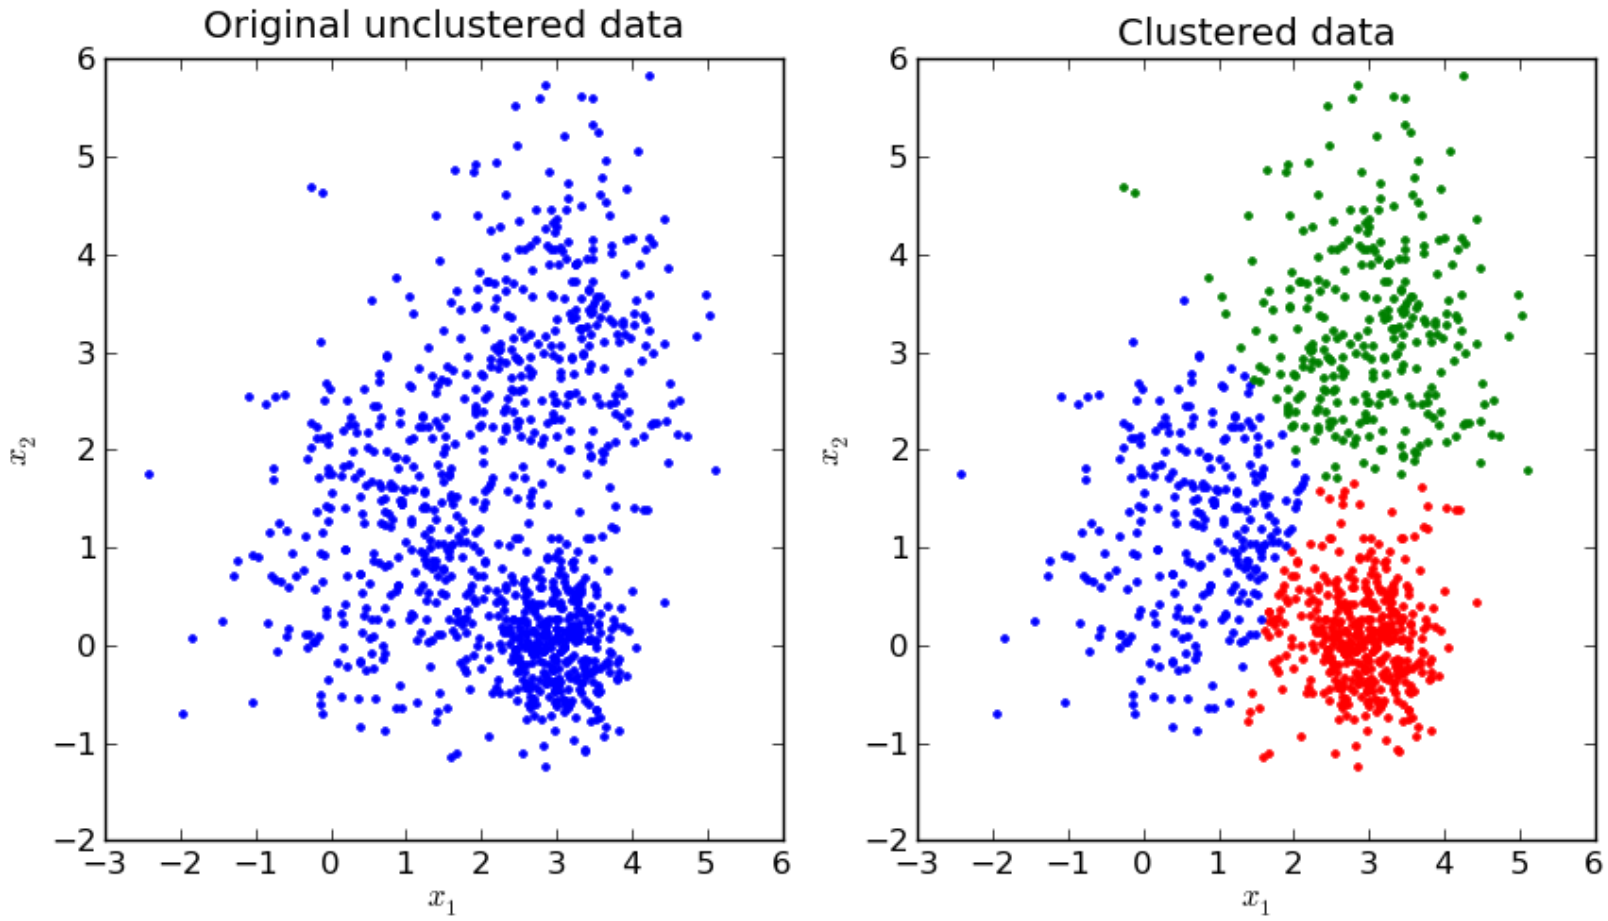

Clustering

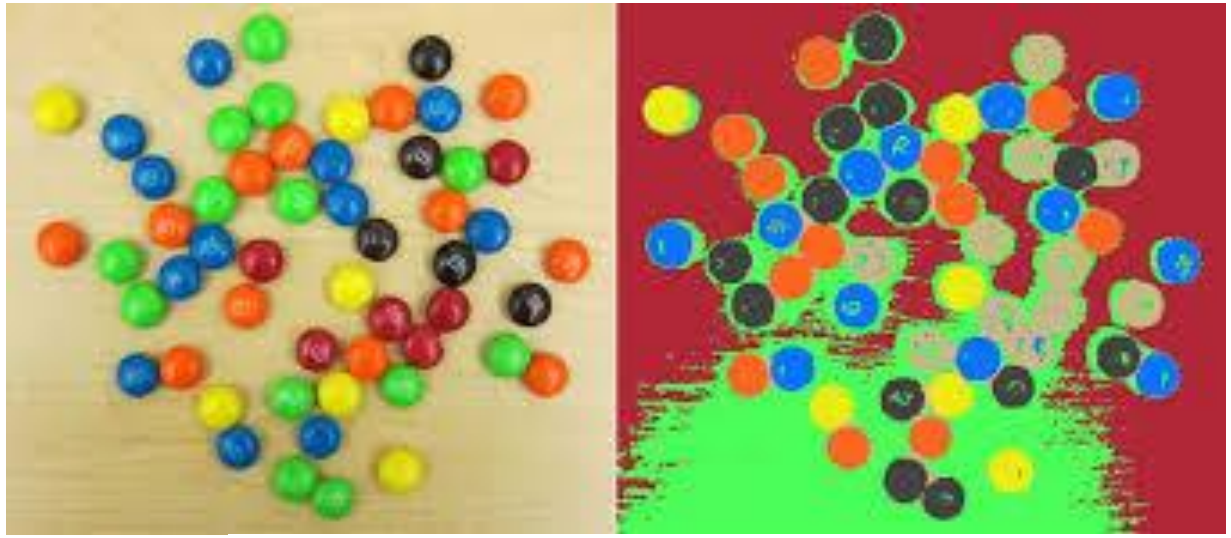

Datasets

- The dataset is a matrix of variables (features, attributes) that represent the observations of the real world
- Each row contains a set of attributes (features, variables)
- Each row can be seen as an instance and is referred to as a feature vector

Training/test data

- The training data is the data that is used as input for the learning algorithm for inducing the model
- It is the data (experience) that is used by the model to build its hypothesis (whether classification or regression)
- The test dataset is the subset of the data that is used to measure the accuracy of the model