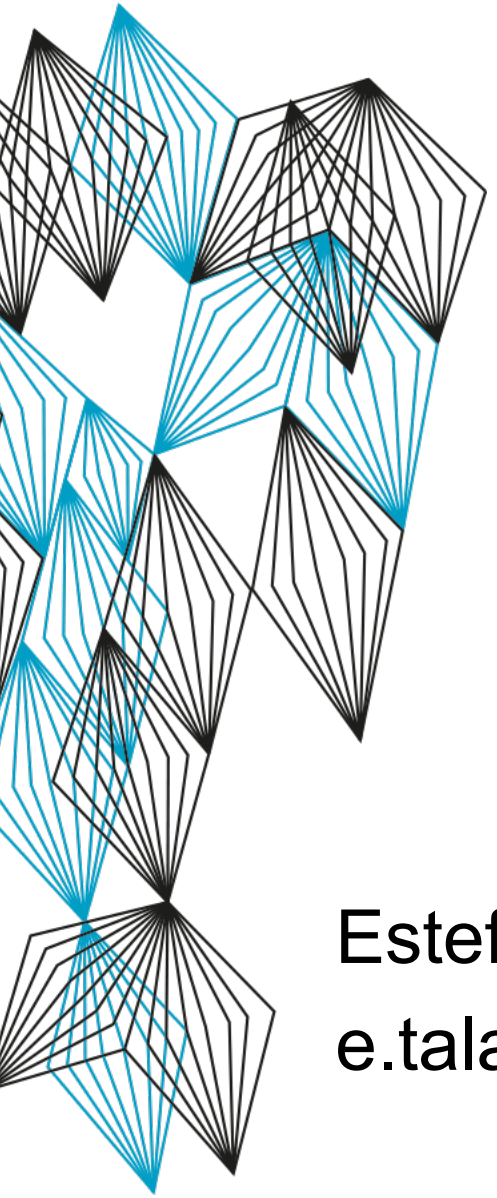

ARTIFICIAL INTELLIGENCE & CYBER SECURITY

INTRODUCTION TO MACHINE LEARNING

Estefanía Talavera
e.talaveramartinez@utwente.nl

Learning/Intelligence

Can *machines* (automata, computers, programs) be intelligent ?

… think ?

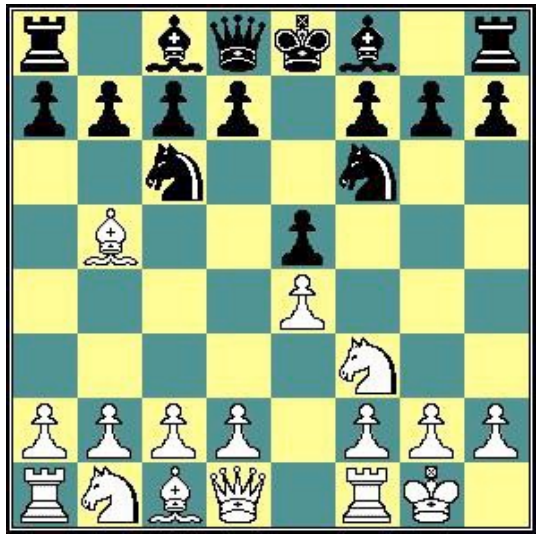

Computer chess

- Deep Blue beats Kasparov (May 1997)
- Matches expert level performance
- 'Thinks' differently from human expert …
  by examining ~ 200 million possible situations

real intelligence **or** ***"just computation"*** **???**

Intelligent systems

Some aspects of **Intelligent Systems** in Computer Science

Perception:
interaction with environment requires cognitive processes,
e.g. computer **vision**, speech **recognition**, motion **detection,**
scene **analysis,** object **classification**

Decision making:
processing of incoming information, **analysis** of a situation,
selection of possible **actions** in order to achieve a **goal**,
e.g. path finding, sorting of objects

Learning:
data driven **adaptation** of the system based on **observations** only (*unsupervised)* or together with **feed-back** from the environment (supervised), e.g. **classification, regression**

An actual IQ test problem

Which figure does not belong to the group ?

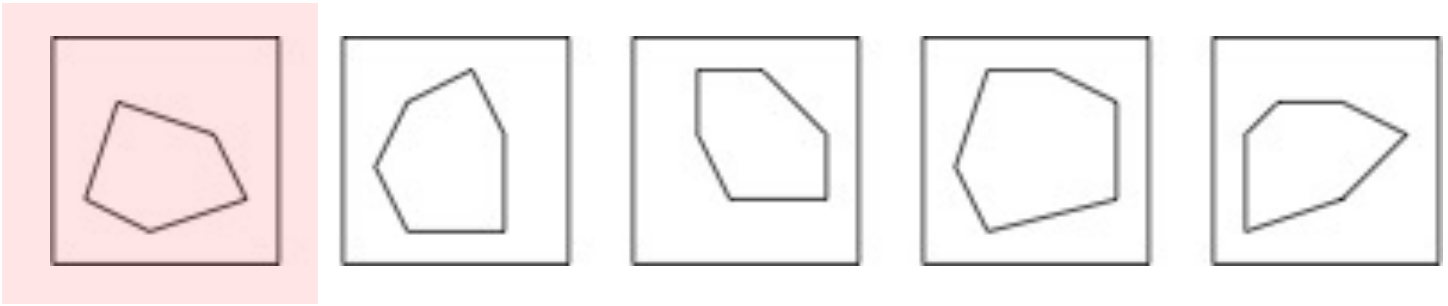

a) b) c) d) e)

Perception: “polygons!”

Feature selection: “count number of edges”

Sorting/clustering: “group objects according to number of edges”

Decision/Action: “(a) is different from the others”

more buzz-words related to intelligent systems

fuzzy logic

smart homes

brain inspired computing

expert systems

self-organization

optimization

artificial life

big data

deep learning

multi-agent systems

genetic algorithms

natural language understanding

evolutionary computation

organic computing

Types of machine learning problems

Unsupervised learning

input data, no labels

- Clustering: group similar data points.
- Density estimation
- Dimensionality reduction
- Outlier/novelty detection

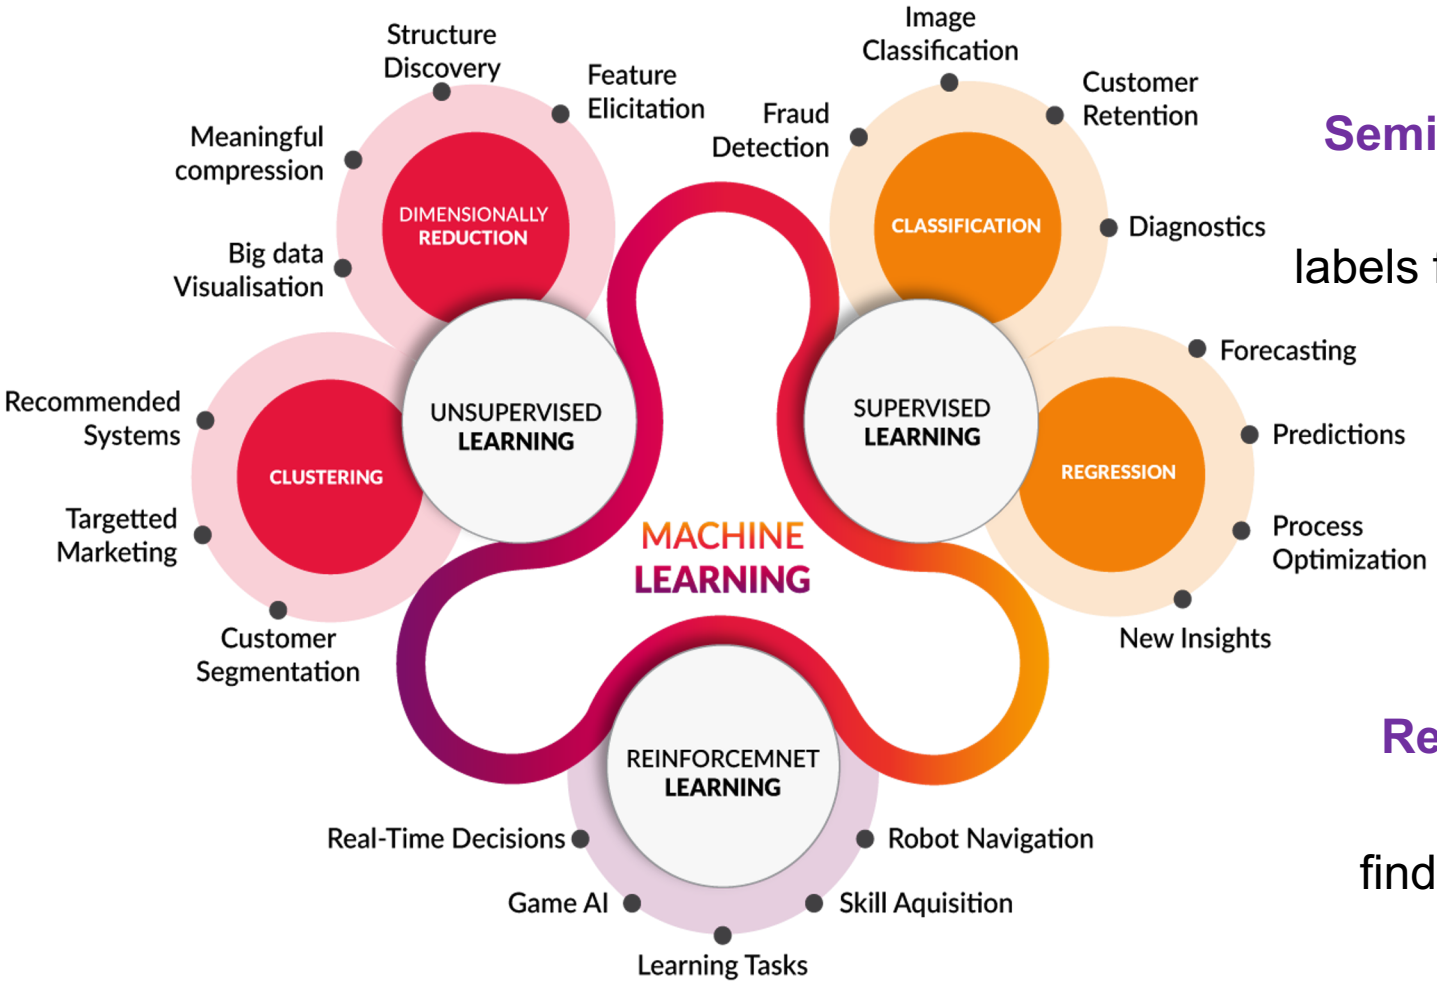

*Image from https://towardsdatascience.com/coding-deep-learning-for-beginners-types-of-machine-learning-b9e651e1ed9d

Supervised learning

labels are provided

- Classification
- Regression

Semi-Supervised learning

labels for just part of the data

Reinforcement learning

find a sequence of actions (policy) that reaches a target.

Supervised Learning

labels are provided

- Classification
- Regression

Basic framework

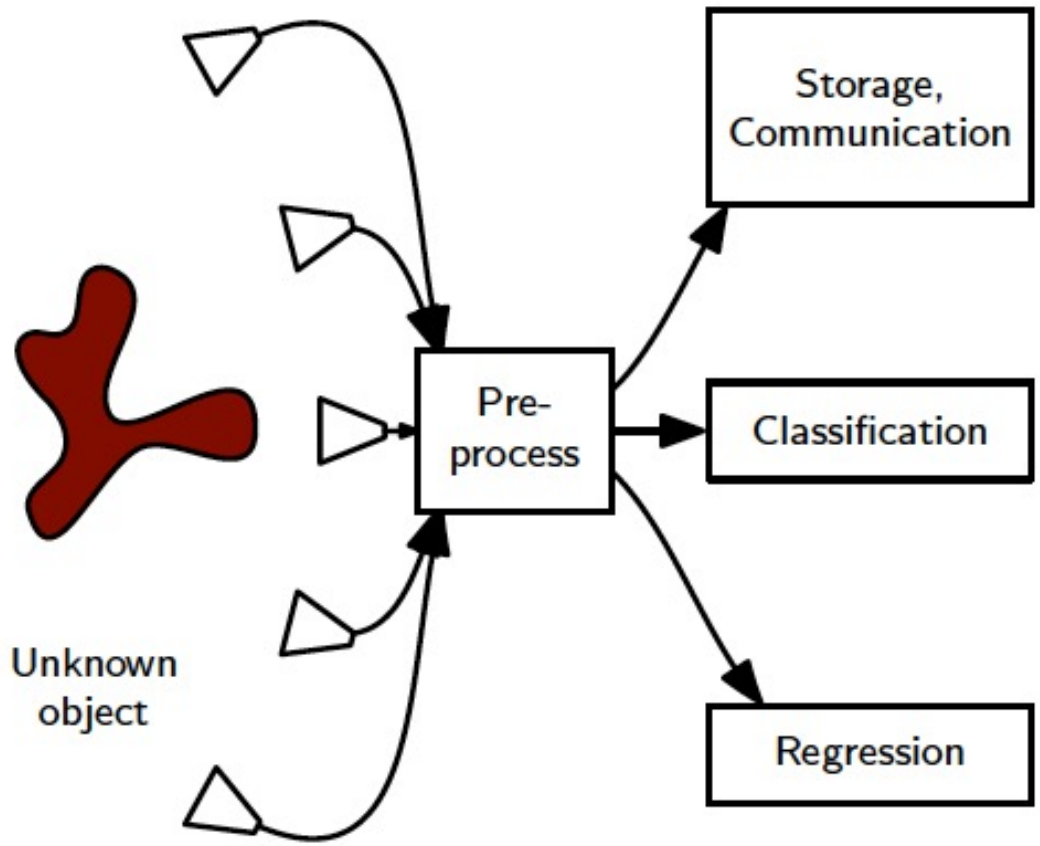

Basic framework

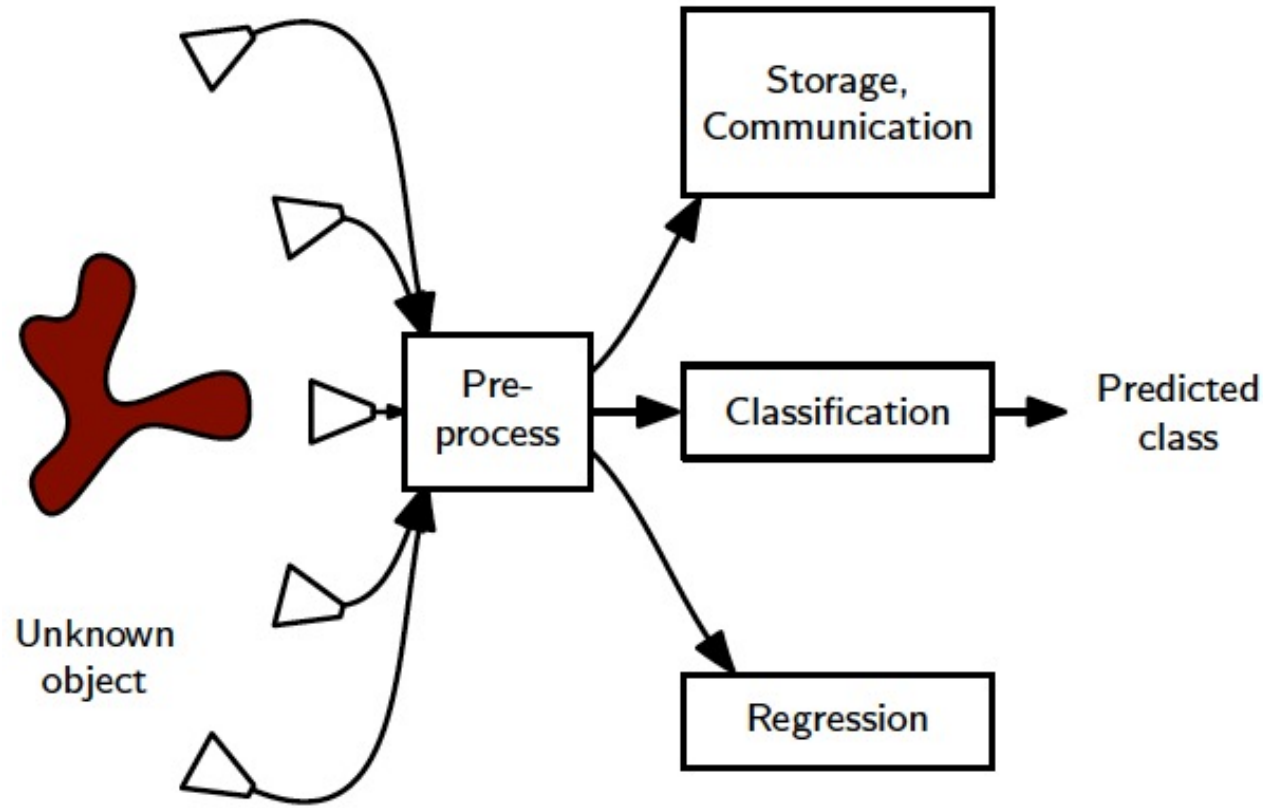

An example of classification

Example: Iris classification

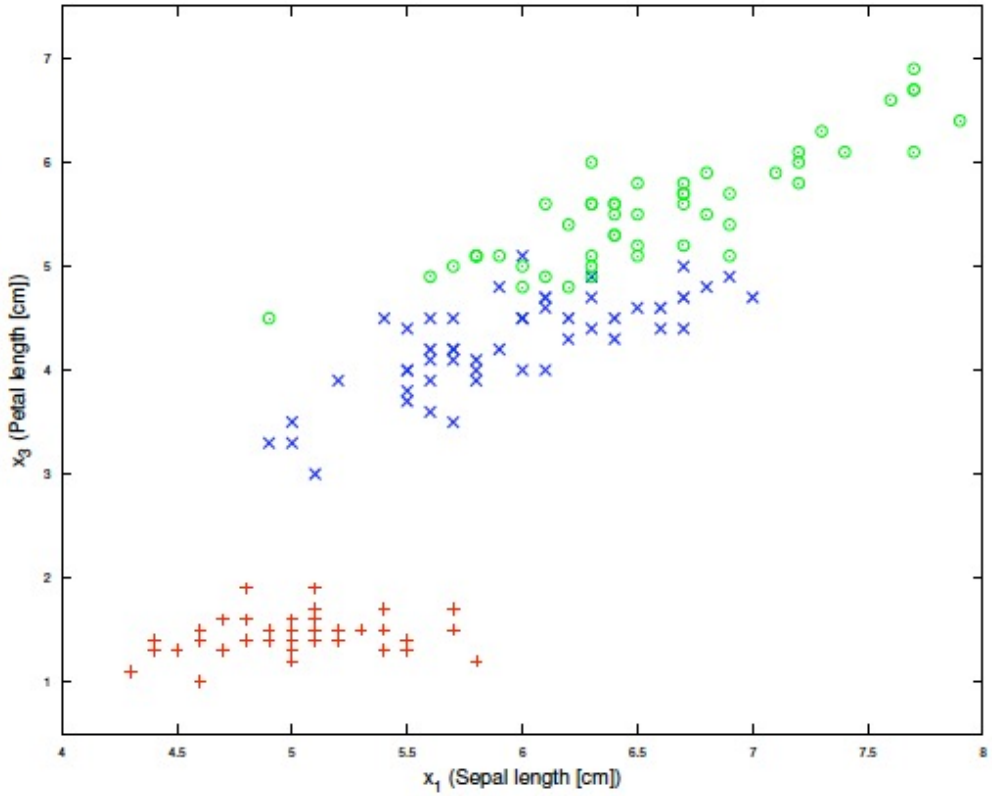

Basic framework

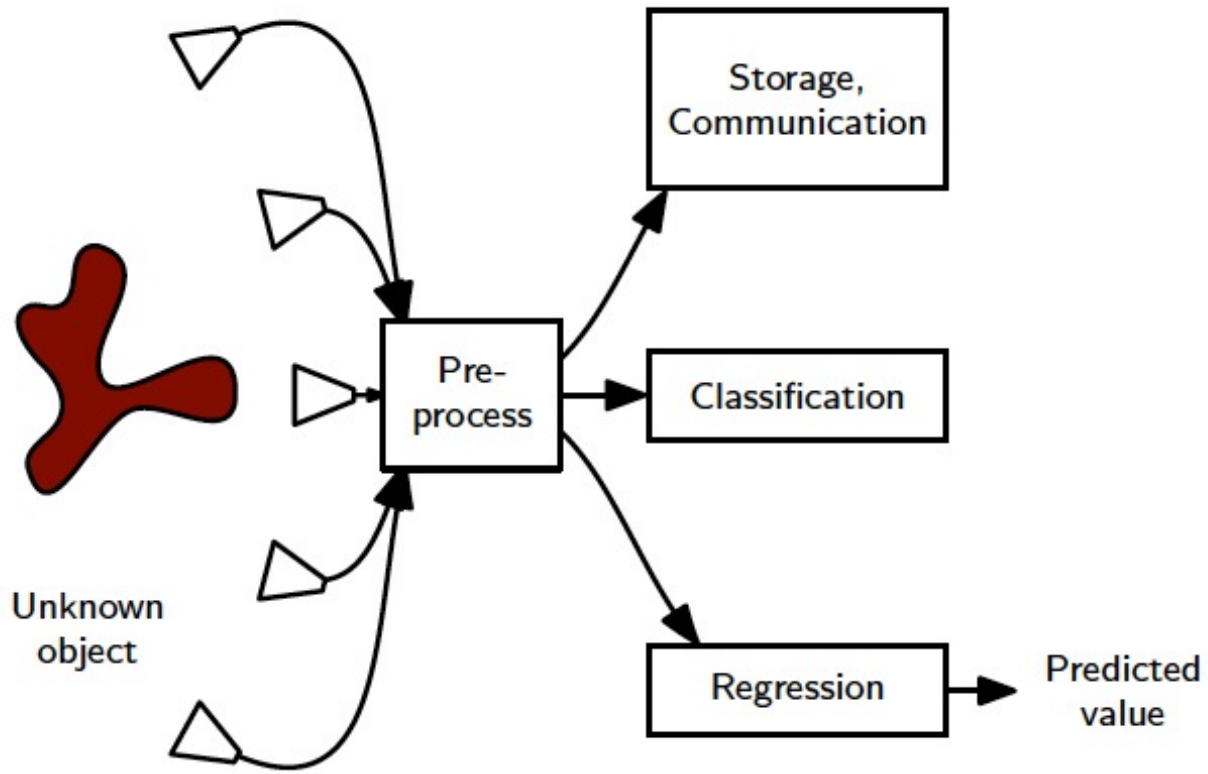

Regression example

Example: Evaluating temperature from cricket activity

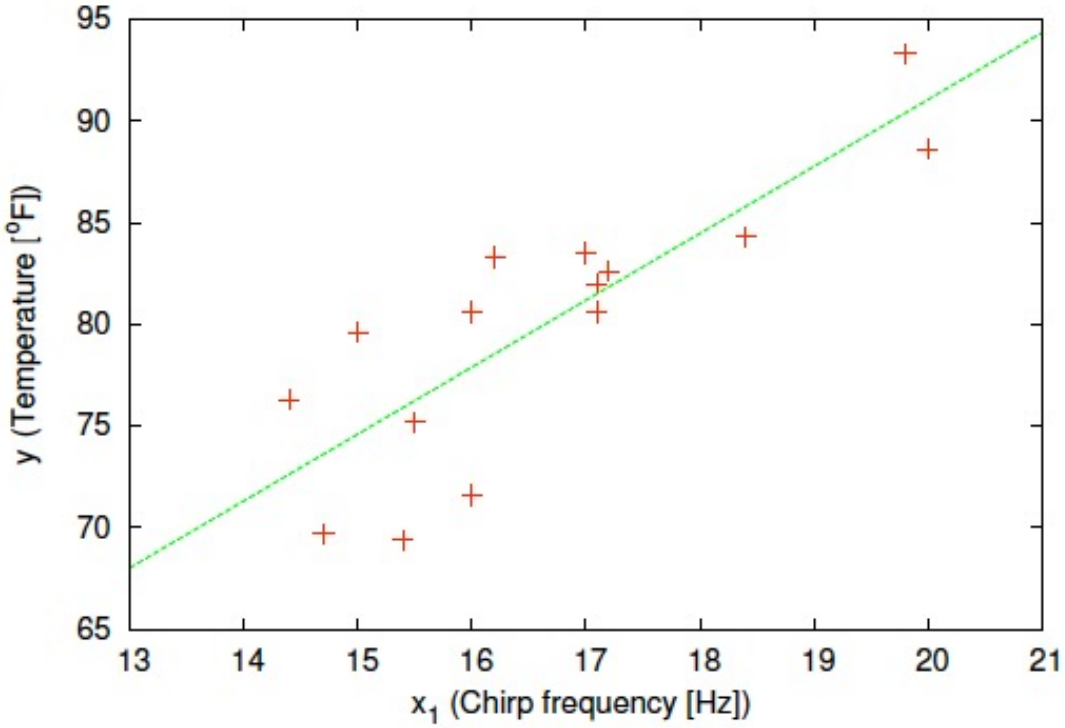

Classification vs Regression

- **Classification:** Predict a discrete label from features

  Example

  - Medicine: classify X-rays as \cancer" or \healthy"
  - SPAM detection: classify emails as spam or not
  - Face recognition, speech recognition, . . .

- **Regression:** Predict a continuous value

  Example

  - Weather forecasting (wind speed, mm rainfall, . . .)
  - In financial markets: predict tomorrow's stock price from past evolution and external factors
  - A robot learning its location in an environment

Supervised Training - Classification

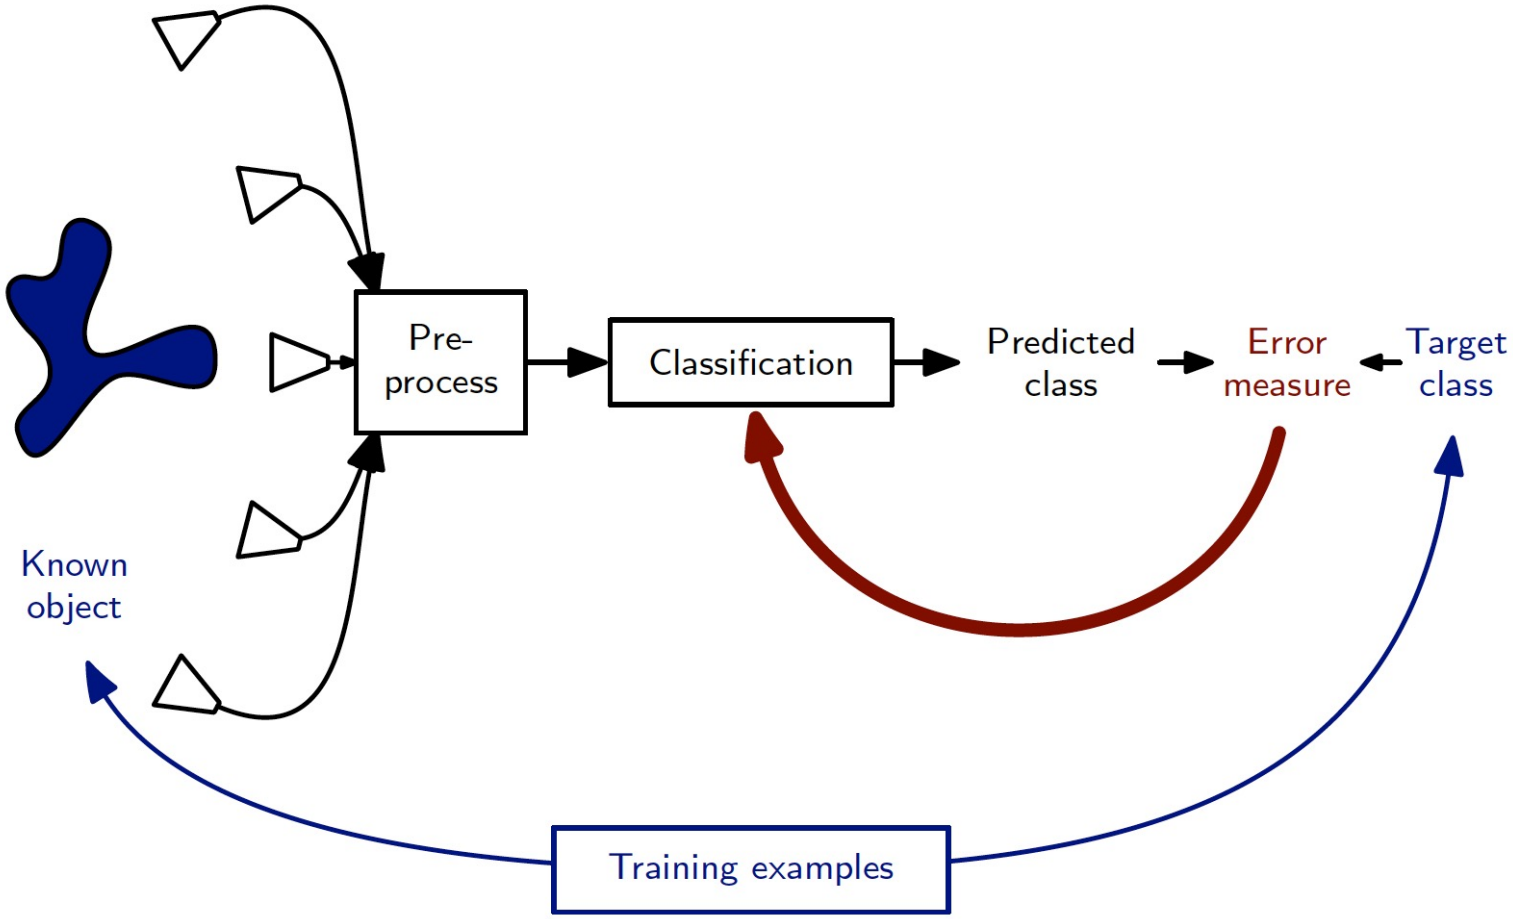

Supervised Training - Regression

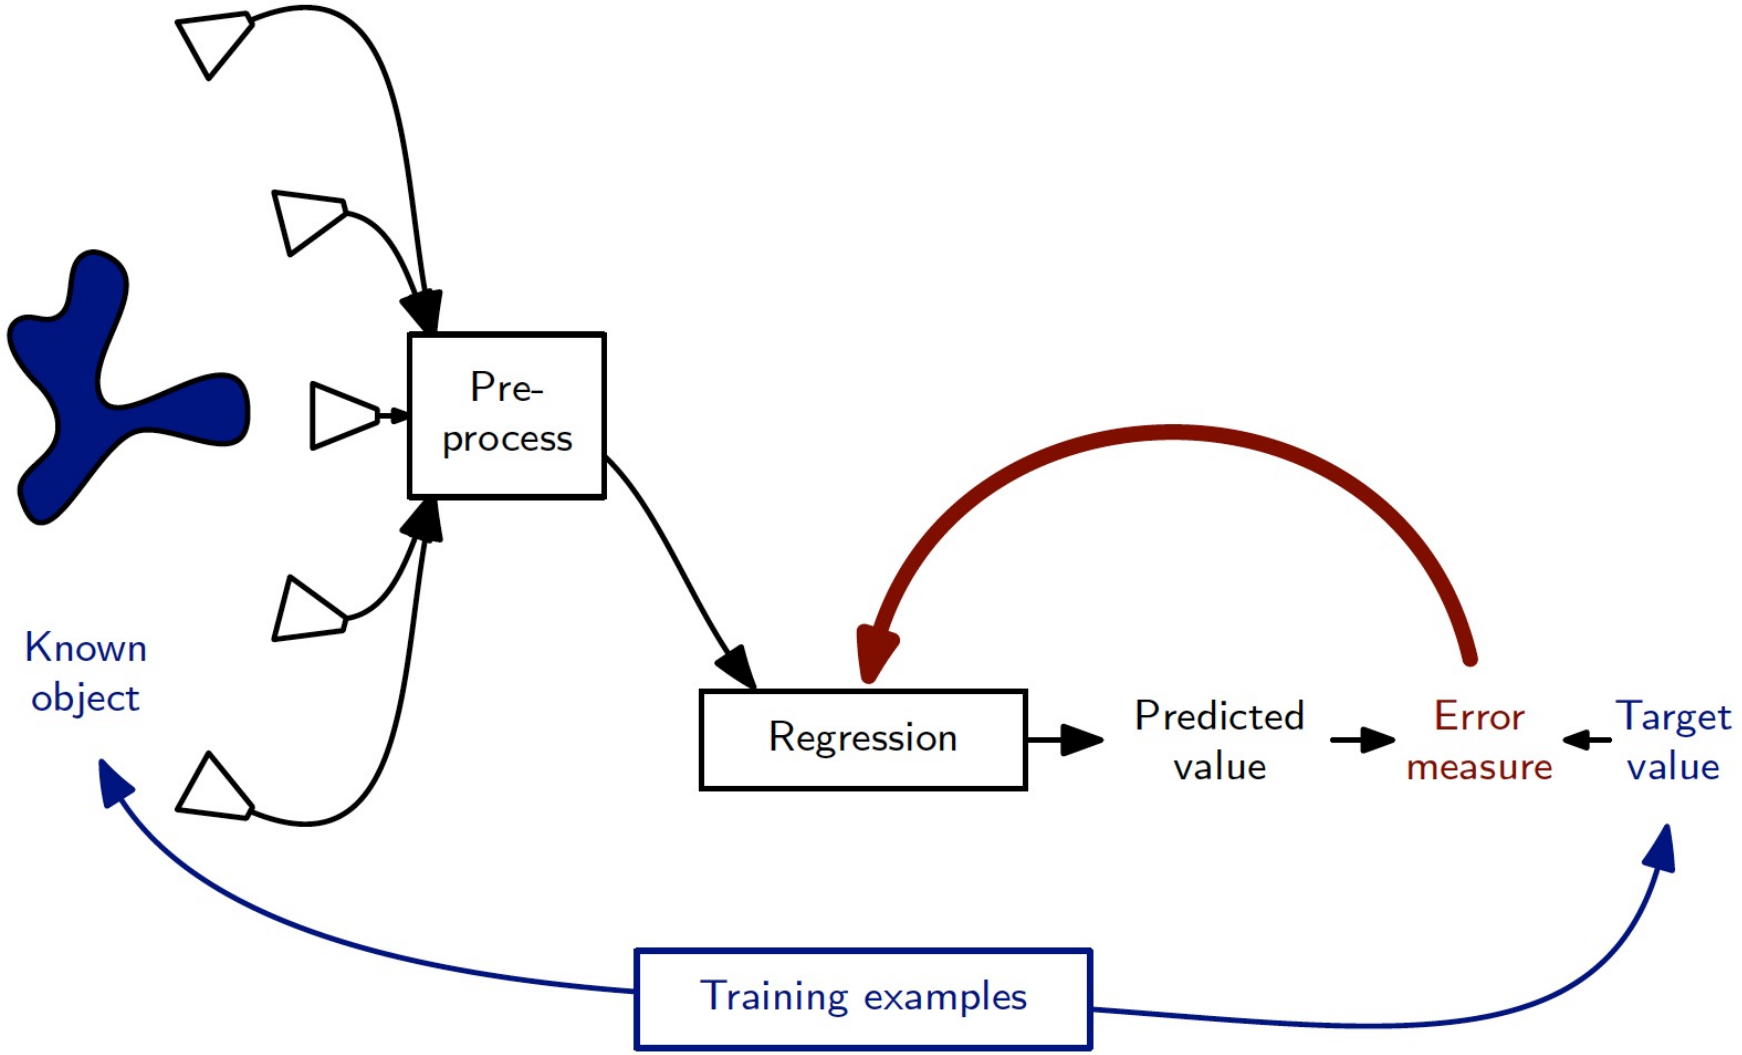

Semi-supervised Training

Semi-supervised learning, self-supervised learning

partially labelled data sets, first training based on labelled subset, extend by making predictions for unlabeled data, accept examples with high confidence for next training … iterative procedure

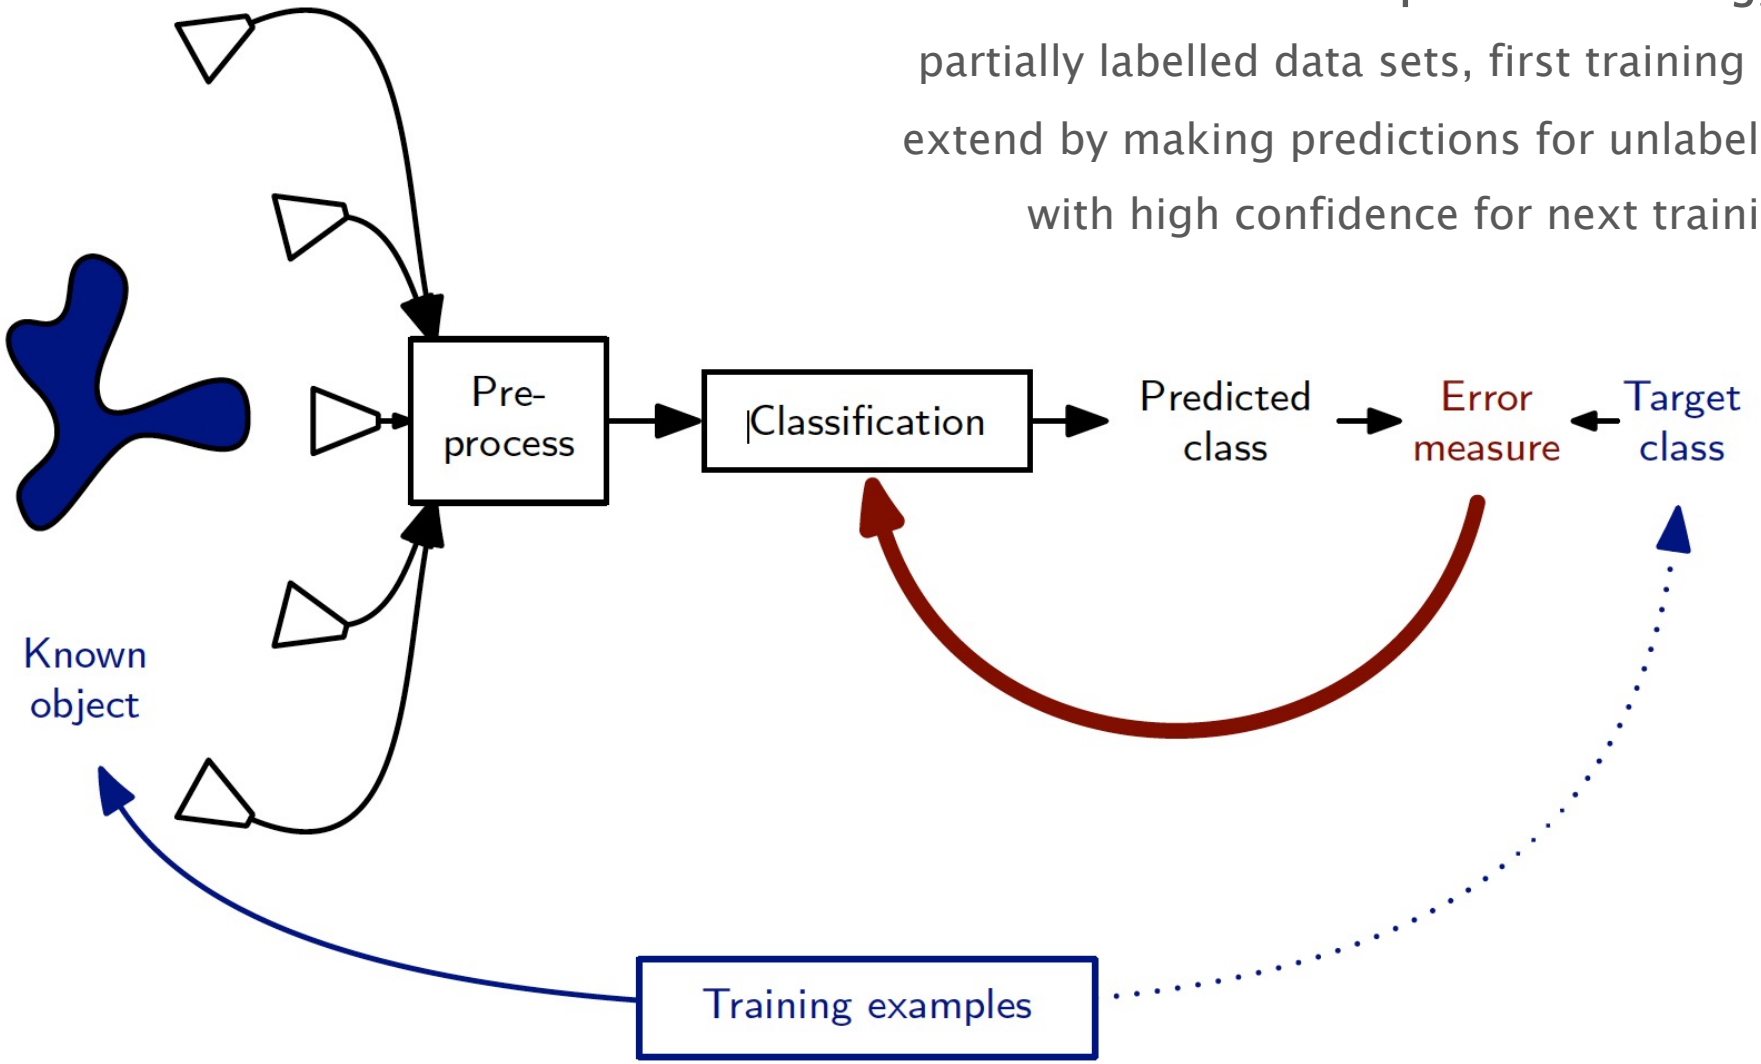

Unupervised Learning

input data, no labels

- Clustering: group similar data points.
- Density estimation
- Dimensionality reduction
- Outlier/novelty detection

Unsupervised Training

Clustering

Given a set of data points:

partition it into groups such that points within each group are similar (low inter-group variability) and groups are dissimilar (high intra-group variability)

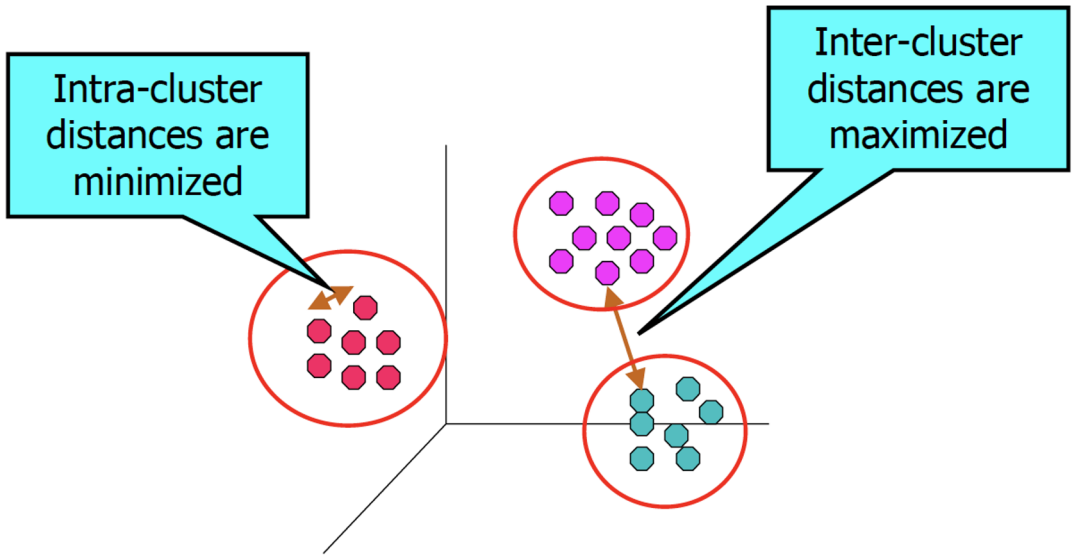

– Unsupervised learning

– Requires data, but no labels

– Detect patterns

e.g. : Group emails or search results, Customer shopping patterns, Regions of images

– Useful when don't know what you're looking for

– But: can be difficult to understand

© Image : Tan, Steinbach, & Kumar, 2005

Other forms of learning (examples)

Reinforcement learning

delayed reward (a form of supervision) after a number of decisions has been taken, e.g. in path finding problems or *games*

Transfer learning, few shot / single shot learning

adaptation of a pre-trained system (on a particular data set) to a (related) new data set or task (different statistics, classes, quality etc.) by means of efficient re-training

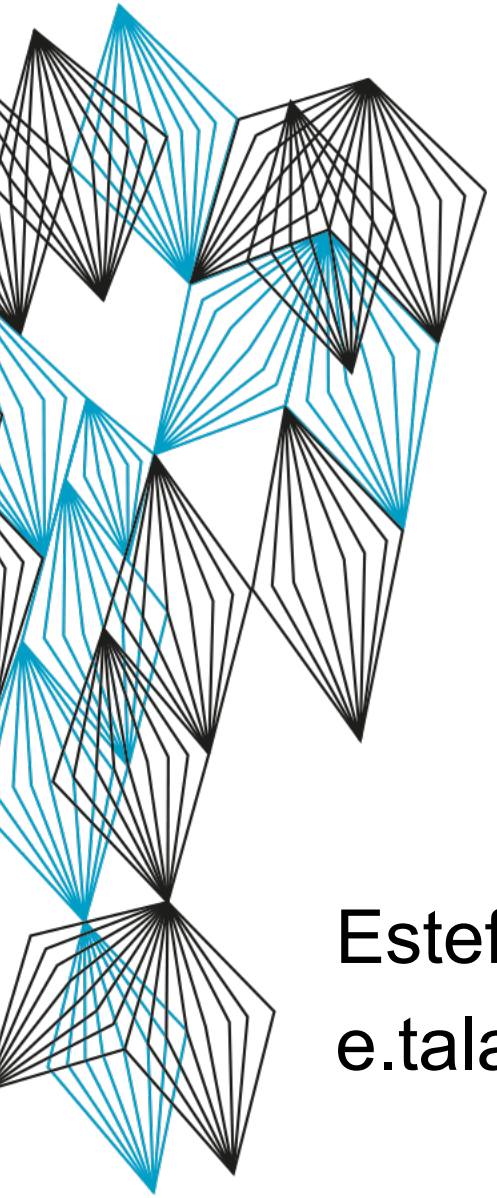

ARTIFICIAL INTELLIGENCE & CYBER SECURITY

INTRODUCTION TO MACHINE LEARNING

Estefanía Talavera
e.talaveramartinez@utwente.nl