


Artificial Intelligence (AI) and Machine Learning (ML)

AI: Computer programs that mimic human intelligence to perform complex tasks

ML: adaptable algorithms that train their parameters from data



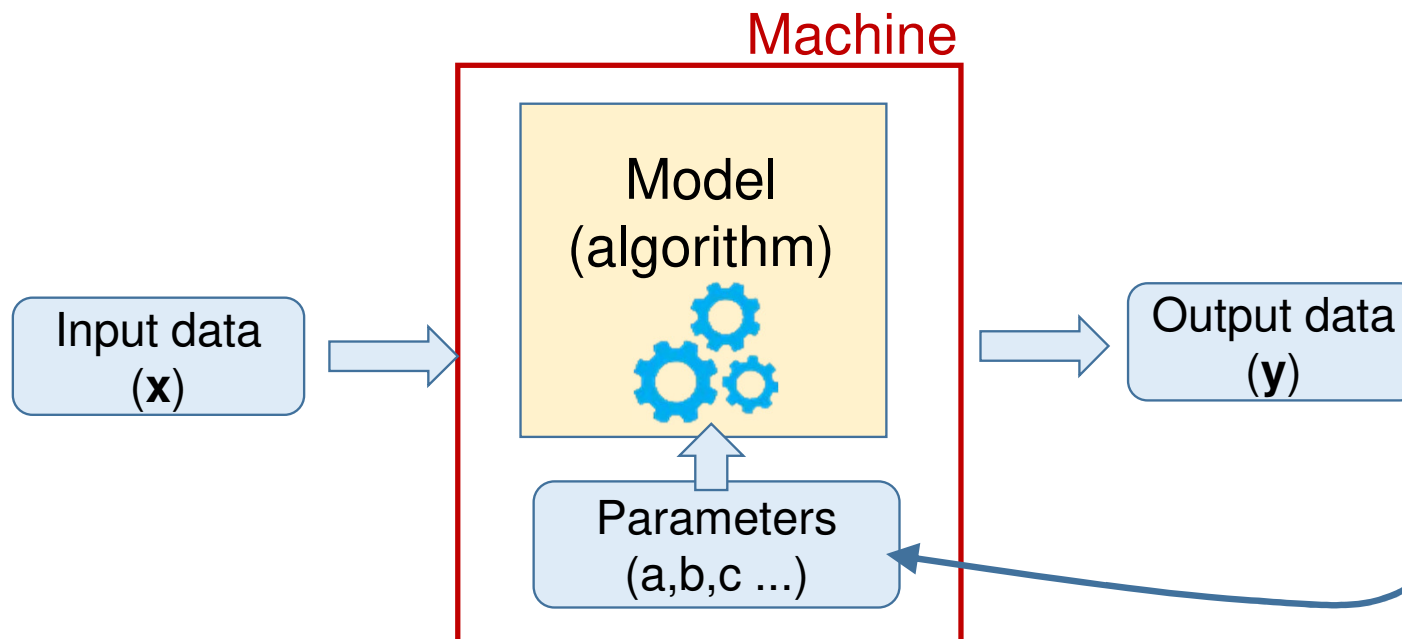
ANN (Artificial Neural Networks): ML algorithm based on a structure that resembles the neuron connections in the brain

DL (Deep Learning): ANNs formed by several layers of artificial neurons

Very short introduction to Machine Learning

Definition

- ❑ Machine Learning (ML) indicates all kinds of strategies and methods that allow tuning an artificial system ("the machine") to make predictions about future data the basis of past data.
- ❑ ML is all about "self-learning" algorithms.
- ❑ ML is subfield of Artificial Intelligence (AI).



In a ML system, the correct values of the parameters to obtain the desired response is autonomously determined by the machine after a "training" phase


Example of Applications of Machine Learning

- ❑ Every kind of feature recognition in images (faces, objects, hazards, etc.)
- ❑ Voice recognition (form sounds to words)
- ❑ Text interpretation (spam filters, web-search tasks, document selection ...)
- ❑ Pattern recognition from homogeneous or heterogeneous sensor data (chemical component recognition, medical diagnosis, food quality ...)

Advanced application in sensor systems




The three different cases of Machine Learning

- 1) Supervised learning  *Frequently used in sensor applications*
- 2) Unsupervised learning
- 3) Reinforcement learning

Supervised Learning

Single input data
"example"



Label y

Set of "features", or
"predictors" that can
be arranged into a row
or column vector x

Data in this format
are called "labeled data"

Training:

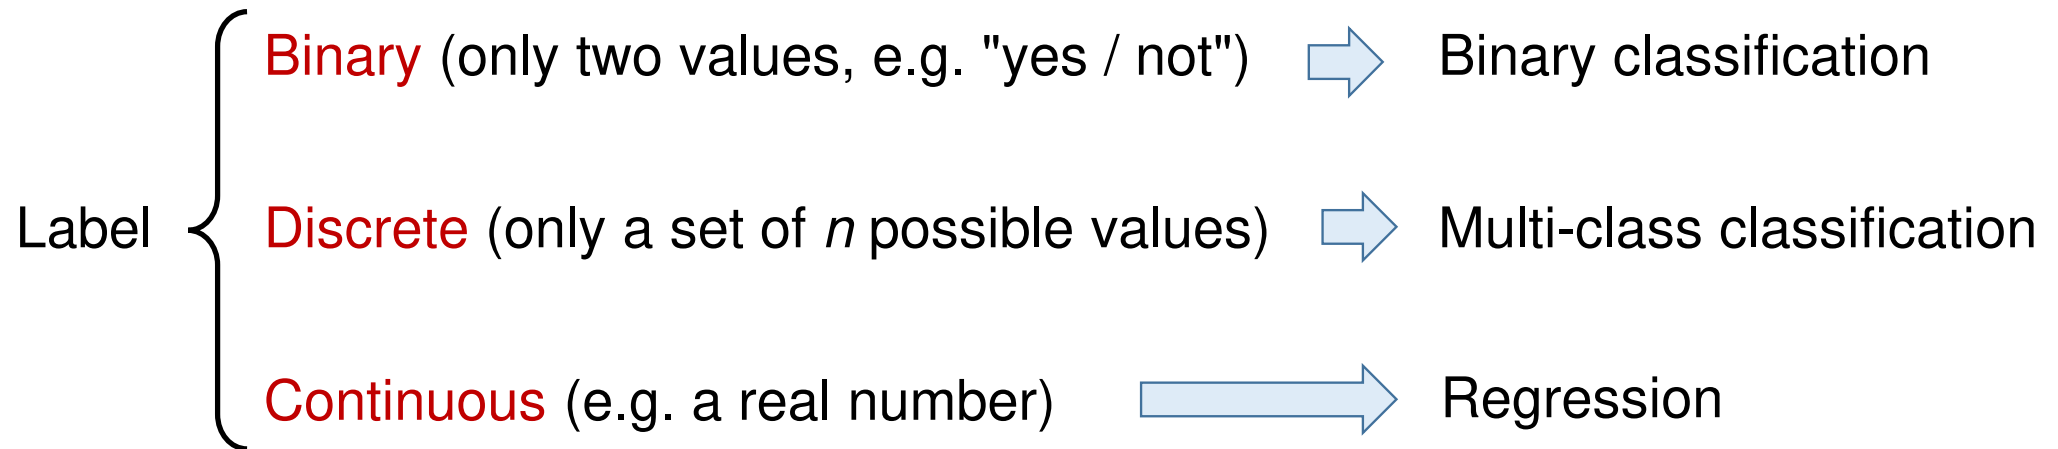
A large number of examples (labelled data) is proposed to the machine that learns to associate different labels (y) to different properties of the feature vector (x).

Goal:

Tune the algorithm (i.e. its parameters) is such a way that it gets able to predict the label of all new x feature vector that are proposed to it after training.

Supervised Learning

There are three main cases of supervised learning, depending on the type of label

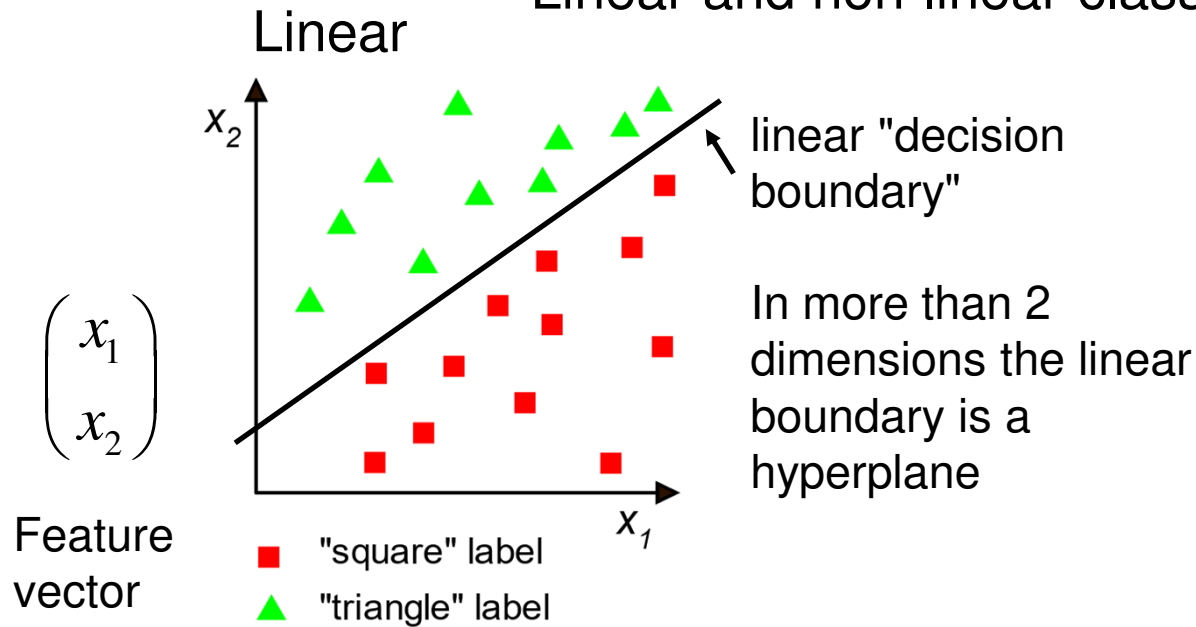


These categories are independent of the type of input data (i.e. type of features), which can be continuous even in the case of binary or discrete label

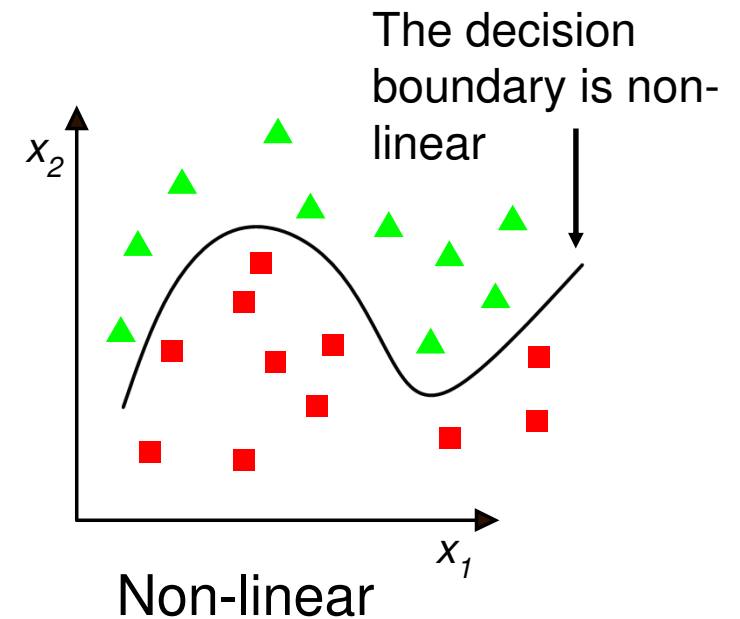
More on classification

- Non-probabilistic classification:** the membership to a given class is predicted
- Probabilistic classification:** the probability of belonging to a class is predicted

Linear and non-linear classification



Graphical representation of the training set



Unsupervised learning

$$\begin{pmatrix} x_1 \\ x_2 \\ \dots \\ x_k \end{pmatrix}$$

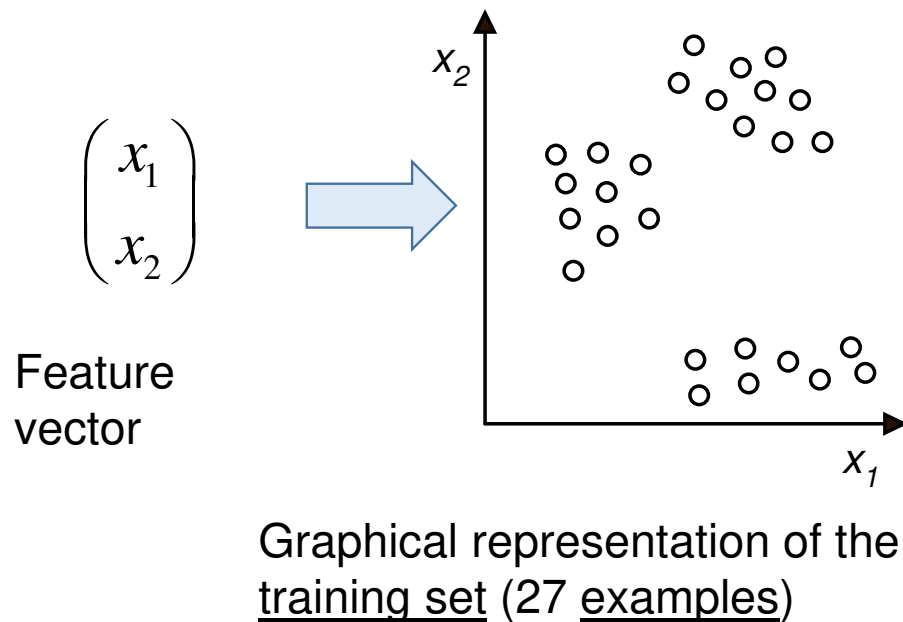
Single input data
"example"

Data are
"unlabeled"


The goal of the training is finding hidden characteristics that can be used to find a criterion to classify the data (data clustering) or to eliminate useless information (dimensionality reduction)

To illustrate the typical goals of unsupervised learning it is better to show a few examples.

Example of Data Clustering



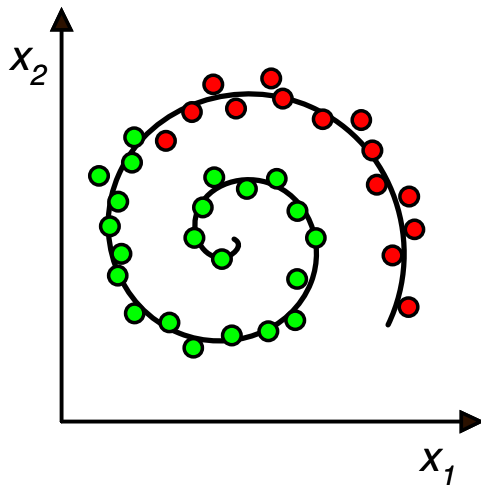
The clustering algorithm understands that data can be grouped into three "clusters". After clustering, a label can be created to signify belonging to a given cluster.



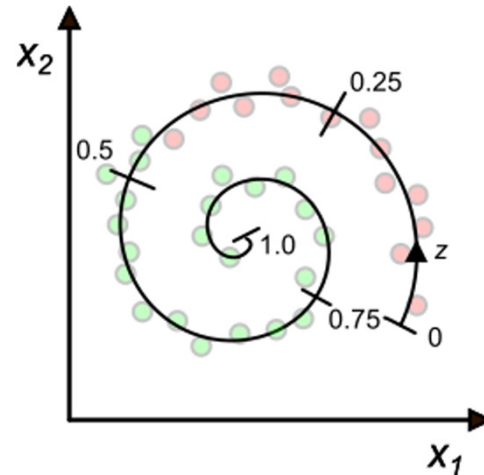
In this example the label can assume three values (red, green, blue)

Once we have learnt this property of the data (that we did not know before) we still have to give an interpretation to it and if it is useful, use this information to simplify classification algorithms.

Example of dimensionality reduction



We have a set of labelled data (green / red) with a 2-D vector. Depending on the algorithm, classification can be difficult with data arranged in this way.




It is clearly possible to arrange data according their position in the spiral (coordinate z).
For this operation, the algorithm do not use the label

The dimensionality reduction task consists in finding coordinate transformation that allow reducing the number of significant features



With this transformation, classification is much easier

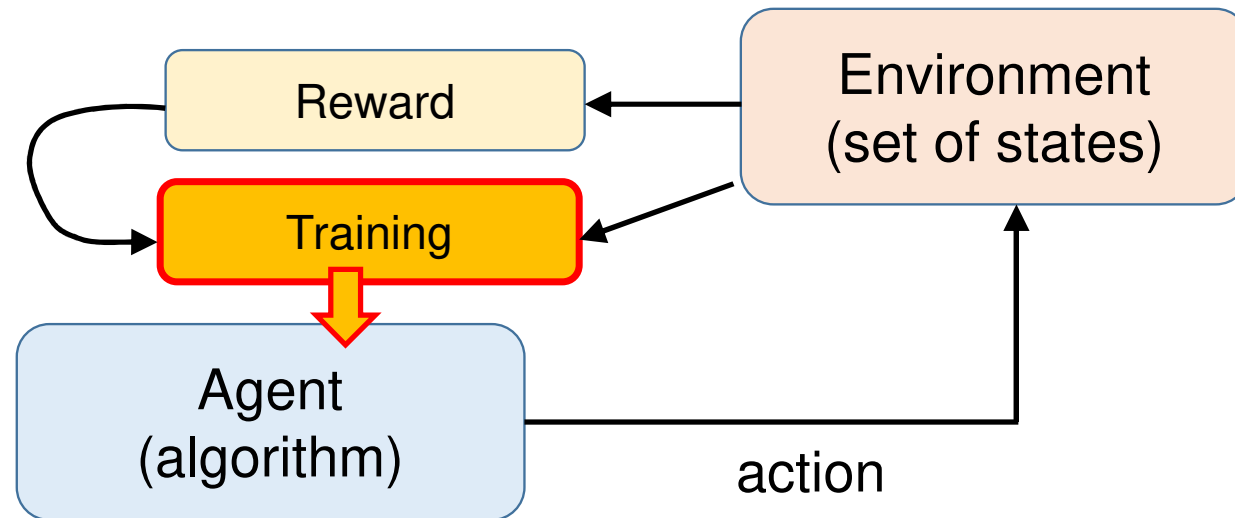


Reinforcement learning

Reinforcement algorithms learn how to perform actions in response to particular situations with the goal to obtain a certain result.

Examples of goals are maximizing the profit of a company, winning a chess match etc ..

The algorithm learns which actions are convenient on the basis of the reward, which is a function of the state. The reward can be positive (win) or negative (lose).

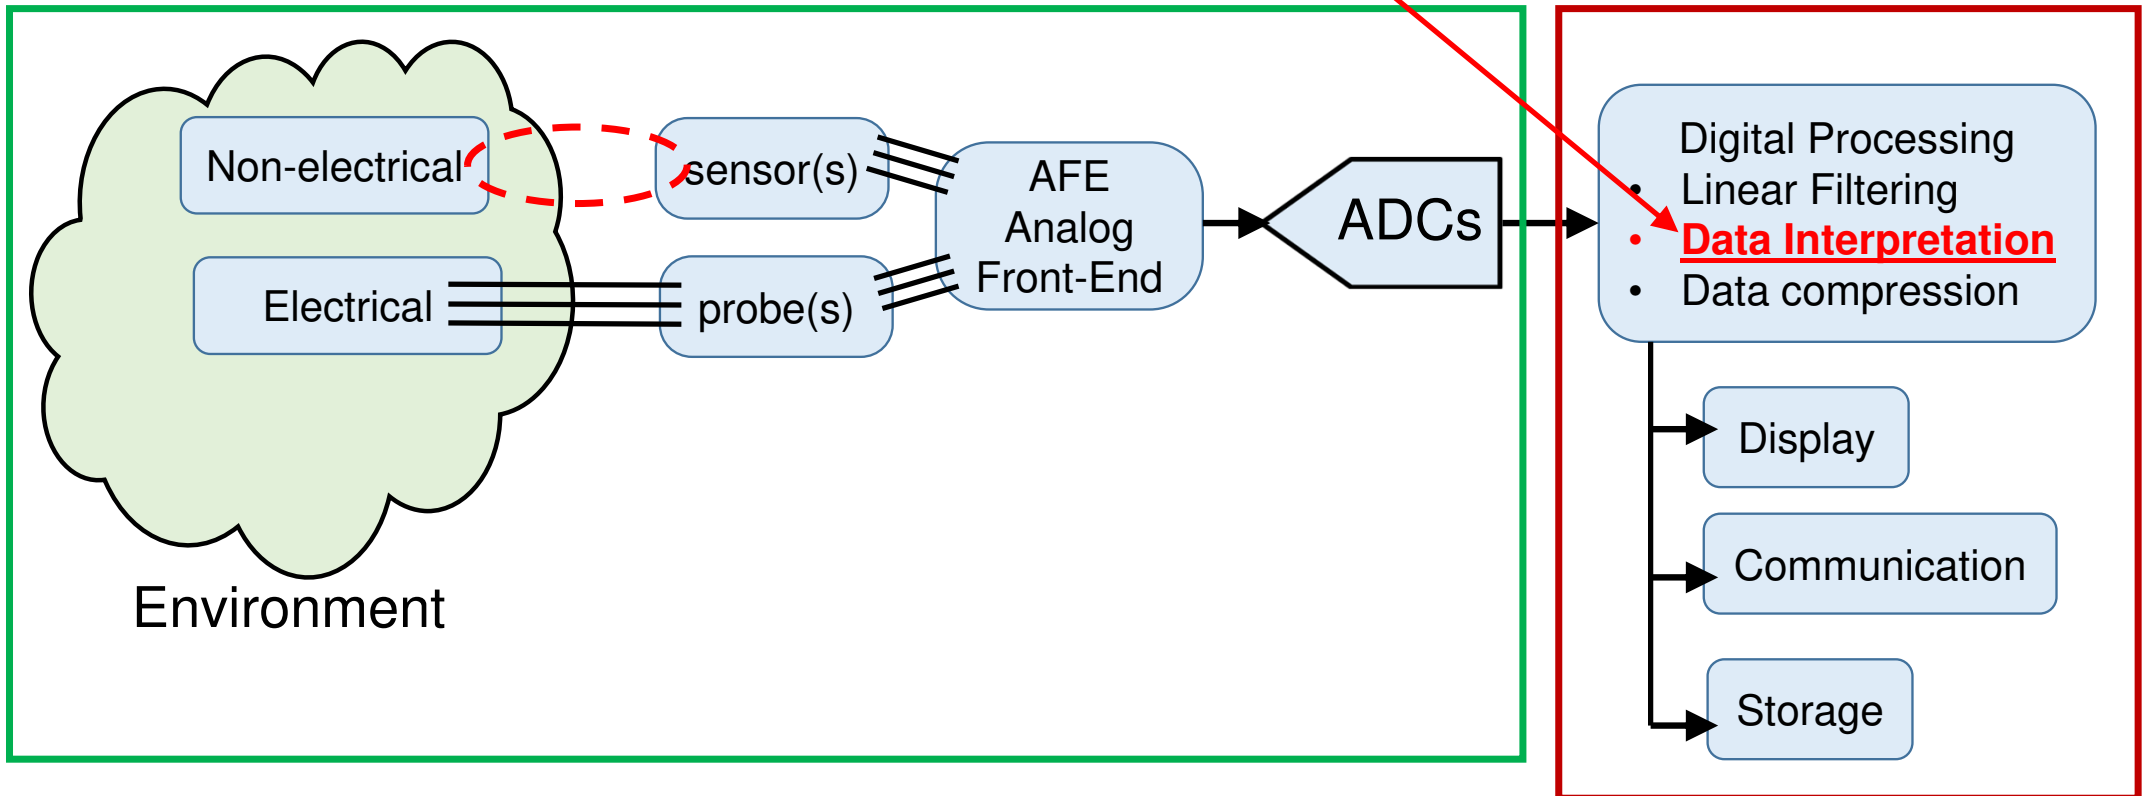


Machine learning in sensor applications

Machine learning can be used at this stage

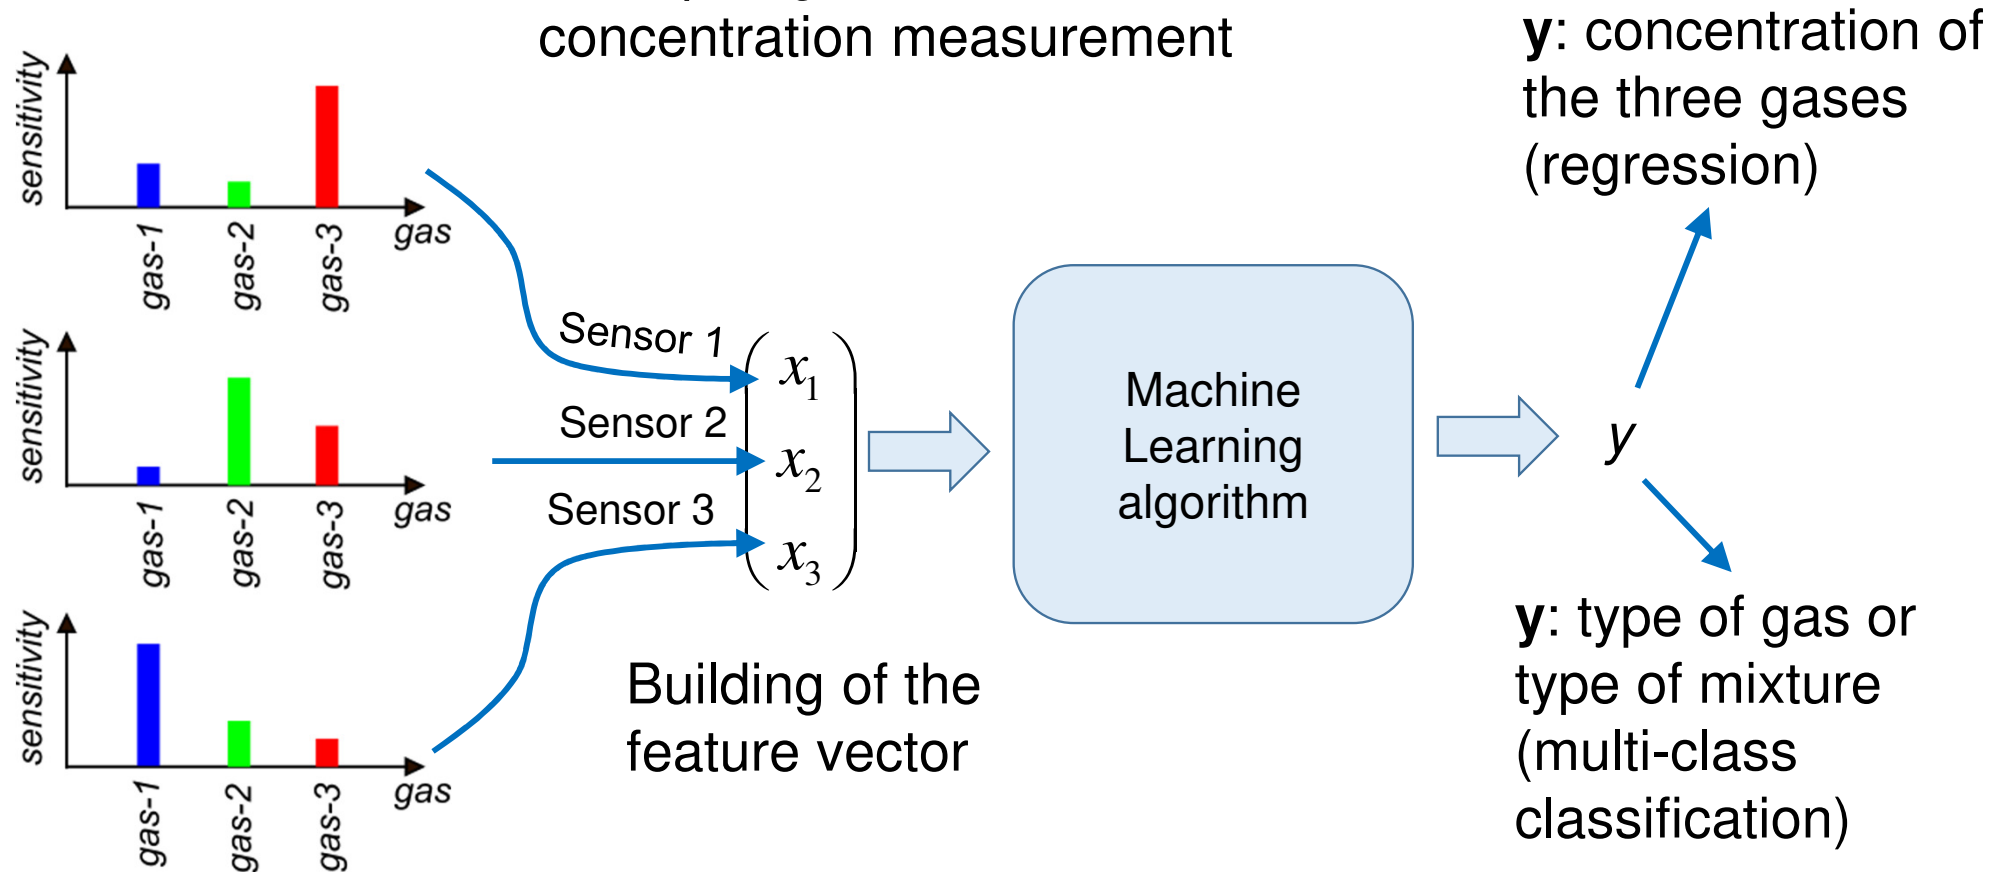
Analog domain

Digital Domain



Machine learning in sensor applications

Example: gas classification and/or concentration measurement



When is ML useful for data interpretation in sensors

Data interpretation can be accomplished in two ways:

A priori model tuning
No Machine learning

1. **Model-driven**: we know the deterministic relationship between the quantities that we want to sense and the sensor outputs ("features"). If the inverse relationship (from the features to the quantities) can be easily found, then it is directly implemented in the algorithm.

2. **Data-driven**: if the relationship between from the quantities to the sensor outputs is **difficult-to-model**, due to a large number of features, presence of strong non-linearities, unreliable properties, complicated dependencies, then its better to use a generic model and let the data tune it to perform the required data interpretation.

Machine learning (= statistical learning)

General characteristics of the training process

Hyperparameters: these are not set by the learning process but have to be chosen before the training begins-. Hyperparameters affect:

- a) The structure of the model
- b) The way the training process is performed

For example, in a polynomial regression process, an hyperparameter can be the degree of the polynomial.

Cost or loss function: measure the difference between the desired outputs and the predicted values

\hat{y} :predicted by the model

Example of
loss function $\frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i)^2$

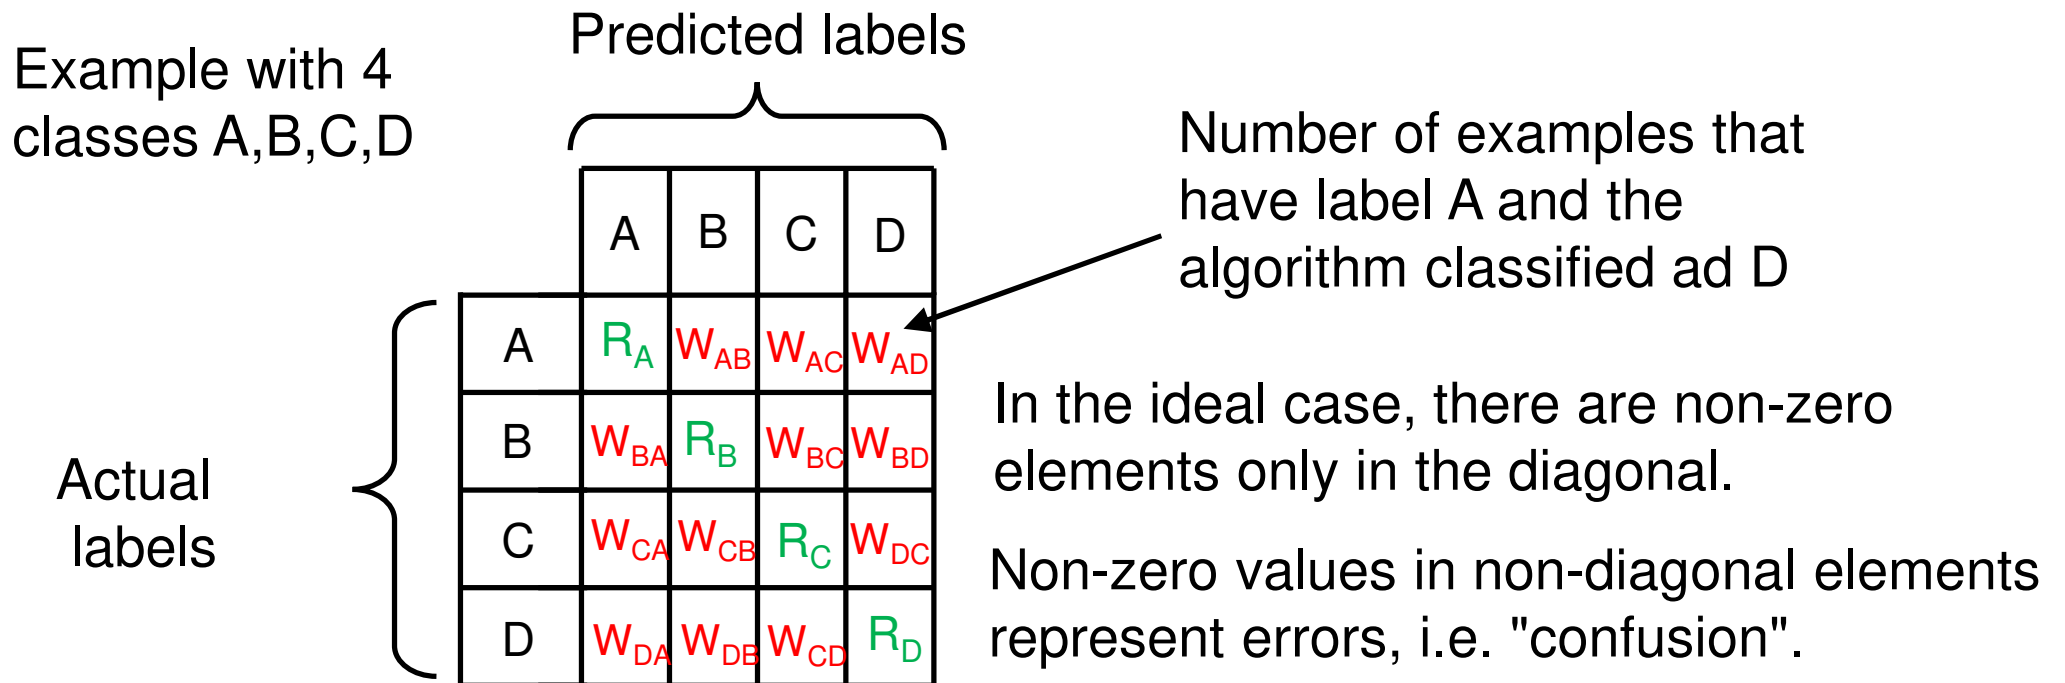
In all phases, proper cost functions are used to assess the degree of accuracy.

Data used in the training process

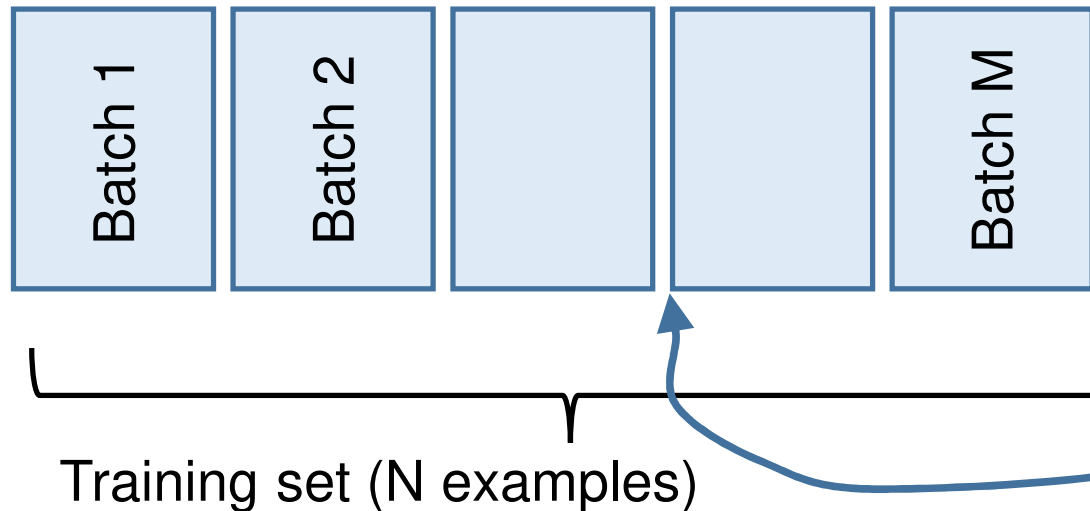
Training data	Used to really train the algorithm (i.e. set the model parameters)
Development data	Used to choose the hyperparameters
Test data	Used at the end of the learning process to test the accuracy of the model

Effectiveness metrics: the confusion matrix

The confusion matrix is used to represent the accuracy of a classification algorithm



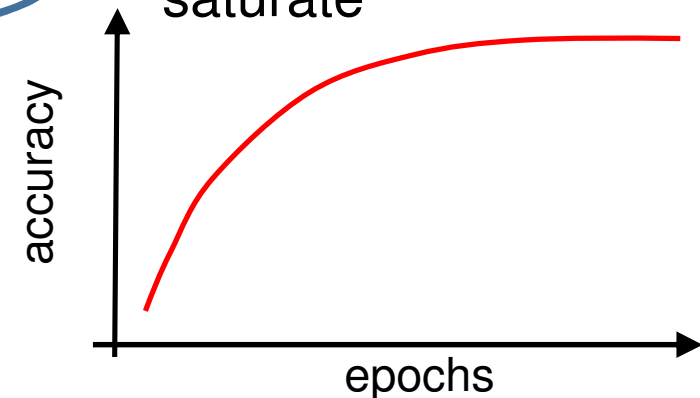
The training phase



Every batch is formed by N / M examples

The model parameters are updated at the end of each batch

The accuracy increases with the number of epochs but tends to saturate



Generally, the same training set is passed through the learning algorithm several times. A single pass of the whole training set is called epoch

Types of ML algorithms

- Decision trees (classification)
- Regression Analysis (regression, continuous target y)
- K-Nearest-Neighbors (k-NN) (classification)
- Logistic Regression (classification)
- Principal Component Analysis (PCA) (clustering)
- Support Vector Machine (SVM): classification
- **Artificial Neural Networks (ANN) (classification and /or regression)**

Artificial Neural Networks: the origins

The predecessor: **The Perceptron**
(Rosenblatt, Cornell Aeronautical Lab. 1958)


Originally implemented on a IBM 704 computer and connected to an array of photodetector to obtain a visual classifier.

Inspired by:

1943: First models of the natural neuron
(McCulloch-Pitt)


Followed by:

1960: Adaline (ADaptive LInear NEuron)
(B. Widrow et al., Stanford)



The perceptron classifier

Inputs (features)



Parameters (weights)

$$\mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_D \end{pmatrix}, \quad \mathbf{w} = \begin{pmatrix} w_1 \\ \vdots \\ w_D \end{pmatrix}$$

$$\hat{y} = \begin{cases} 1 & \text{if } \mathbf{x} * \mathbf{w} + b \geq 0 \\ -1 & \text{if } \mathbf{x} * \mathbf{w} + b < 0 \end{cases}$$

The perceptron operation: 2D case


$$\hat{y} = \begin{cases} 1 & \text{if } \mathbf{x} * \mathbf{w} + b \geq 0 \\ -1 & \text{if } \mathbf{x} * \mathbf{w} + b < 0 \end{cases}$$

$$\mathbf{x} * \mathbf{w} \geq -b$$

$$\mathbf{x} * \mathbf{w} < -b$$


$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$$

$$\mathbf{w} = \begin{pmatrix} w_1 \\ w_2 \end{pmatrix}$$



$$\mathbf{x} * \mathbf{w} = \text{constant} > 0$$

$$\mathbf{x} * \mathbf{w} = \text{constant} < 0$$




Perceptron classifier

Labels y (0 or 1)

$$\begin{pmatrix} x_1^{(1)} & x_2^{(1)} \\ x_1^{(1)} & x_2^{(1)} \\ x_1^{(3)} & x_2^{(3)} \\ \dots & \dots \\ x_1^{(N)} & x_2^{(N)} \end{pmatrix} \begin{pmatrix} 1 \\ -1 \\ -1 \\ \dots \\ 1 \end{pmatrix}$$

Training data

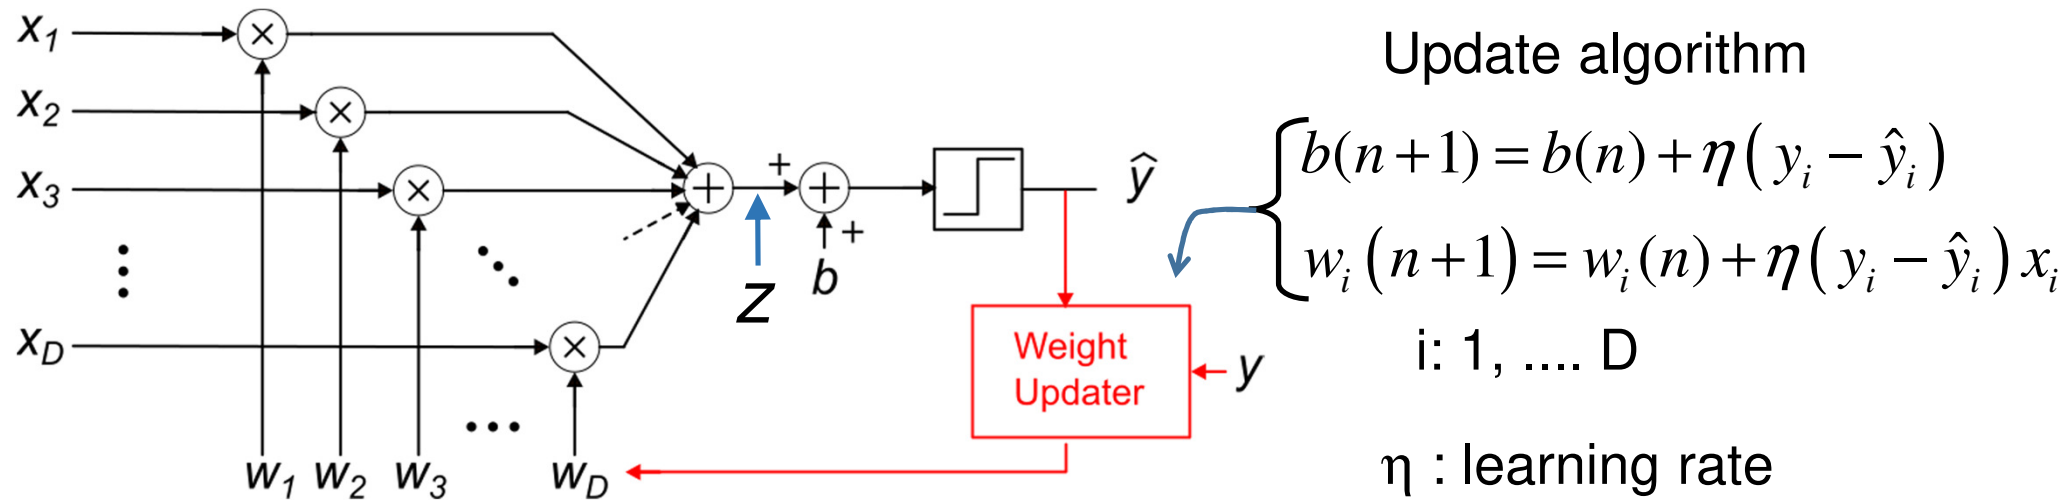


Target: find the correct weight set (w_1, w_2) and bias b that allow the perceptron to predict the label on a new data. In the case shown above, the data are "linearly separable", i.e. a line exists that divide the x_1, x_2 plane into two regions where points with homogeneous label are present

Since the plane exists, this classification problem can be solved by a perceptron

The problem is: how to make the perceptron set the weights and the bias autonomously.

Setting the weights from the training data = learning



$$z = \sum_{i=0}^D x_i w_i$$

$$w_i(n+1) = w_i(n) + \eta (y_i - \hat{y}_i) x_i$$

$$i: 0, \dots, D$$

$$w_0 = b \quad x_0 = 1$$

Artificial Neural Networks

The perceptron implements a "single layer" ANN and can solve only binary classification problems for linearly separable data

An ANN uses more perceptrons (modified to produce a continuous output) to solve a much wider variety of problems