Computer Vision – TP13
Statistical Classifiers

Miguel Tavares Coimbra

Outline

- Statistical Classifiers
- Neural Networks
- Support Vector Machines

Topic: Statistical Classifiers

- Statistical Classifiers
- Neural Networks
- Support Vector Machines

Statistical PR

- I use **statistics** to make a decision
  - I can make **decisions** even when I don't have full a priori knowledge of the whole process
  - I can make **mistakes**
- How did I **recognize** this pattern?
  - I **learn** from previous observations where I know the classification result
  - I **classify** a new observation

Features

- Feature ***F_i*** $F_i = [f_i]$
- Feature ***F_i*** with ***N*** values.

$$F_i = [f_{i1}, f_{i2}, \ldots, f_{iN}]$$

- Feature vector F with M features.

$$F = [F_1 \mid F_2 \mid \ldots \mid F_M]$$

- Naming conventions:
  - Elements of a **feature vector** are called **coefficients**
  - **Features** may have one or more **coefficients**
  - **Feature vectors** may have one or more **features**

Classifiers

- A **Classifier C** maps a class into the feature space

$$C_{\text{Spain}}(x, y) = \begin{cases} true & , y > K \\ false & , otherwise \end{cases}$$

- Various types of classifiers
  - Nearest-Neighbours
  - Neural Networks
  - Support Vector Machines
  - Etc...

Distance to Mean

- I can represent a class by its mean feature vector

$$C = \overline{F}$$

- To classify a new object, I choose the class with the closest mean feature vector
- Different distance measures!

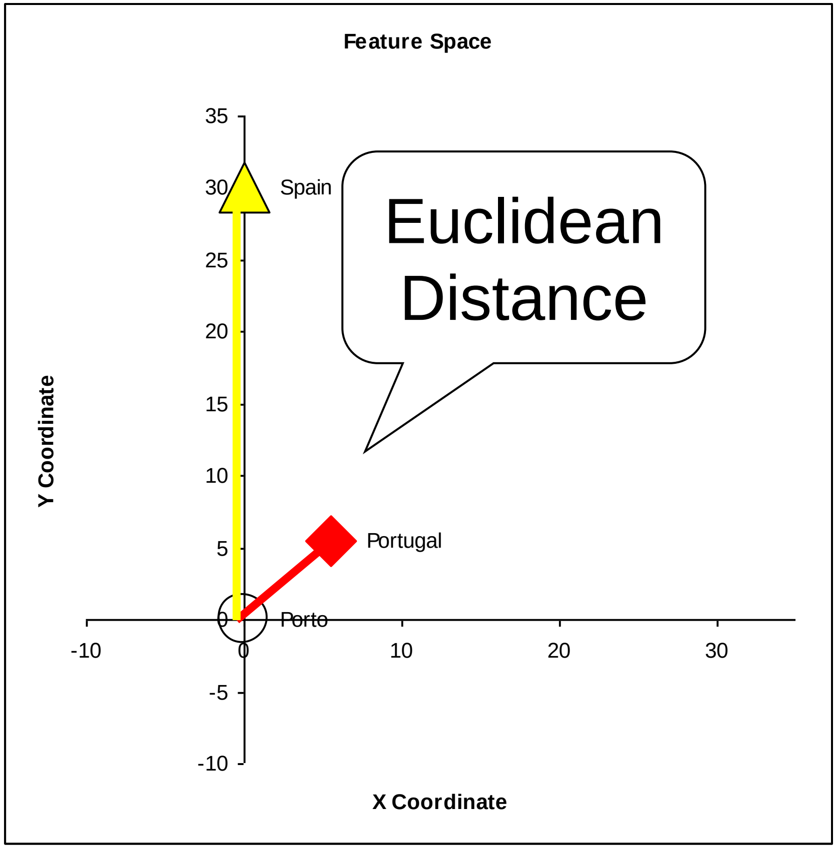

Possible Distance Measures

- L1 Distance

$$L1 = \frac{1}{N}\sum_{x=1}^{N}\left|S(x) - v(x)\right|$$

- Euclidean Distance (L2 Distance)

$$L2 = \frac{1}{N}\sum_{x=1}^{N}\left(S(x) - v(x)\right)^2$$

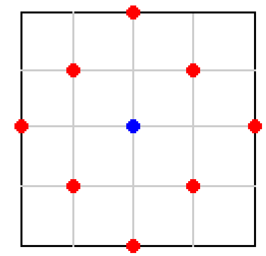

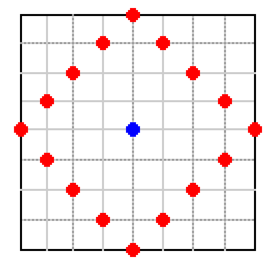

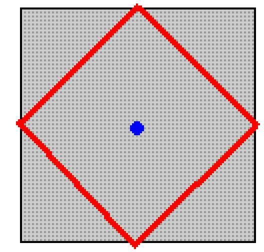

L1 or Taxicab Distance

Gaussian Distribution

- Defined by two parameters:
  - Mean: μ
  - Variance: σ^2
- Great approximation to the distribution of many phenomena.
  - *Central Limit Theorem*

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-u)^2}{2\sigma^2}\right)$$

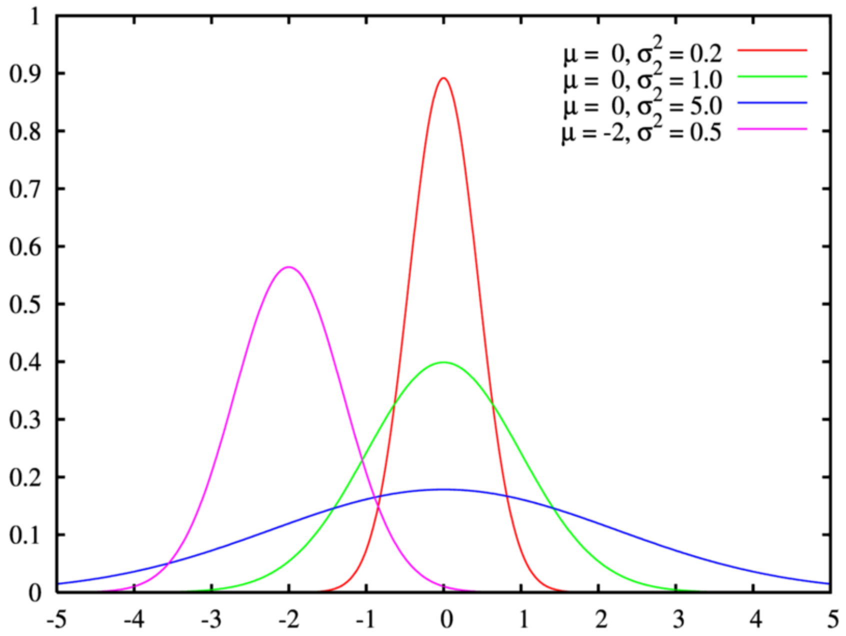

Multivariate Distribution

- For N dimensions:

$$f_X(x_1,\ldots,x_N) = \frac{1}{(2\pi)^{N/2}\,|\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(x-\mu)^\top \Sigma^{-1}(x-\mu)\right)$$

- Mean feature vector:

$$\mu = \overline{F}$$

- Covariance Matrix:

$$X = \begin{bmatrix} X_1 \\ \vdots \\ X_n \end{bmatrix} \qquad \mu_i = \mathrm{E}(X_i) \qquad \Sigma_{ij} = \mathrm{E}\left[(X_i - \mu_i)(X_j - \mu_j)\right]$$

Mahalanobis Distance

- Based on the covariance of coefficients
- Superior to the Euclidean distance

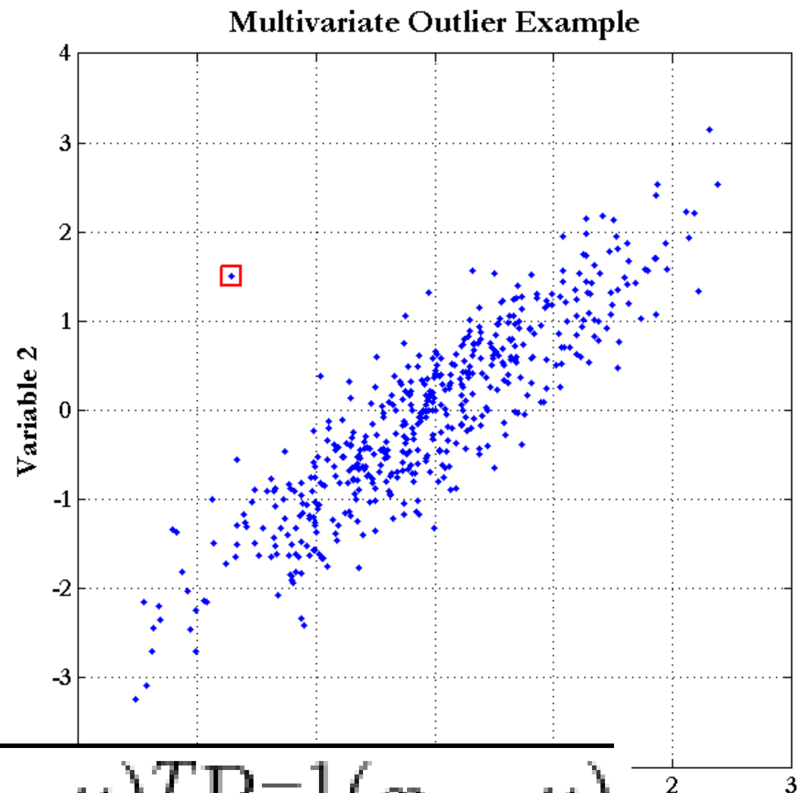

$$D_M(x) = \sqrt{(x-\mu)^T \mathrm{P}^{-1}(x-\mu)}.$$

K-Nearest Neighbours

- Algorithm
  – Choose the closest K neighbours to a new observation
  – Classify the new object based on the **class** of these K objects
- Characteristics
  – Assumes no model
  – Does not scale very well...

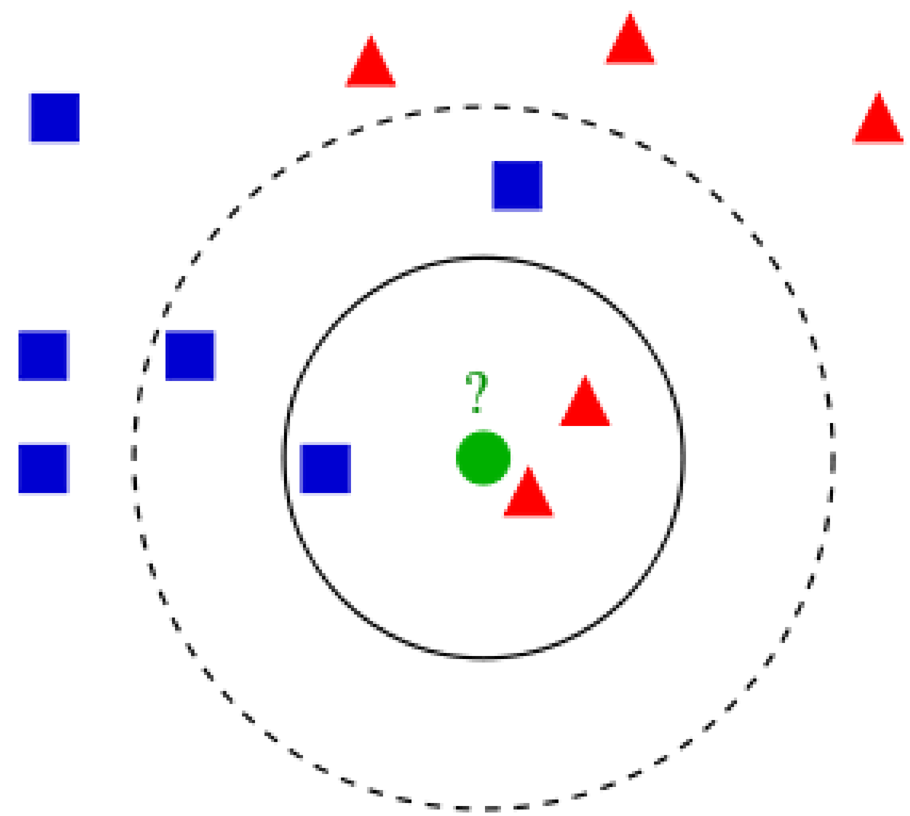

Topic: Neural Networks

- Statistical Classifiers
- Neural Networks
- Support Vector Machines

If you can't beat it.... Copy it!

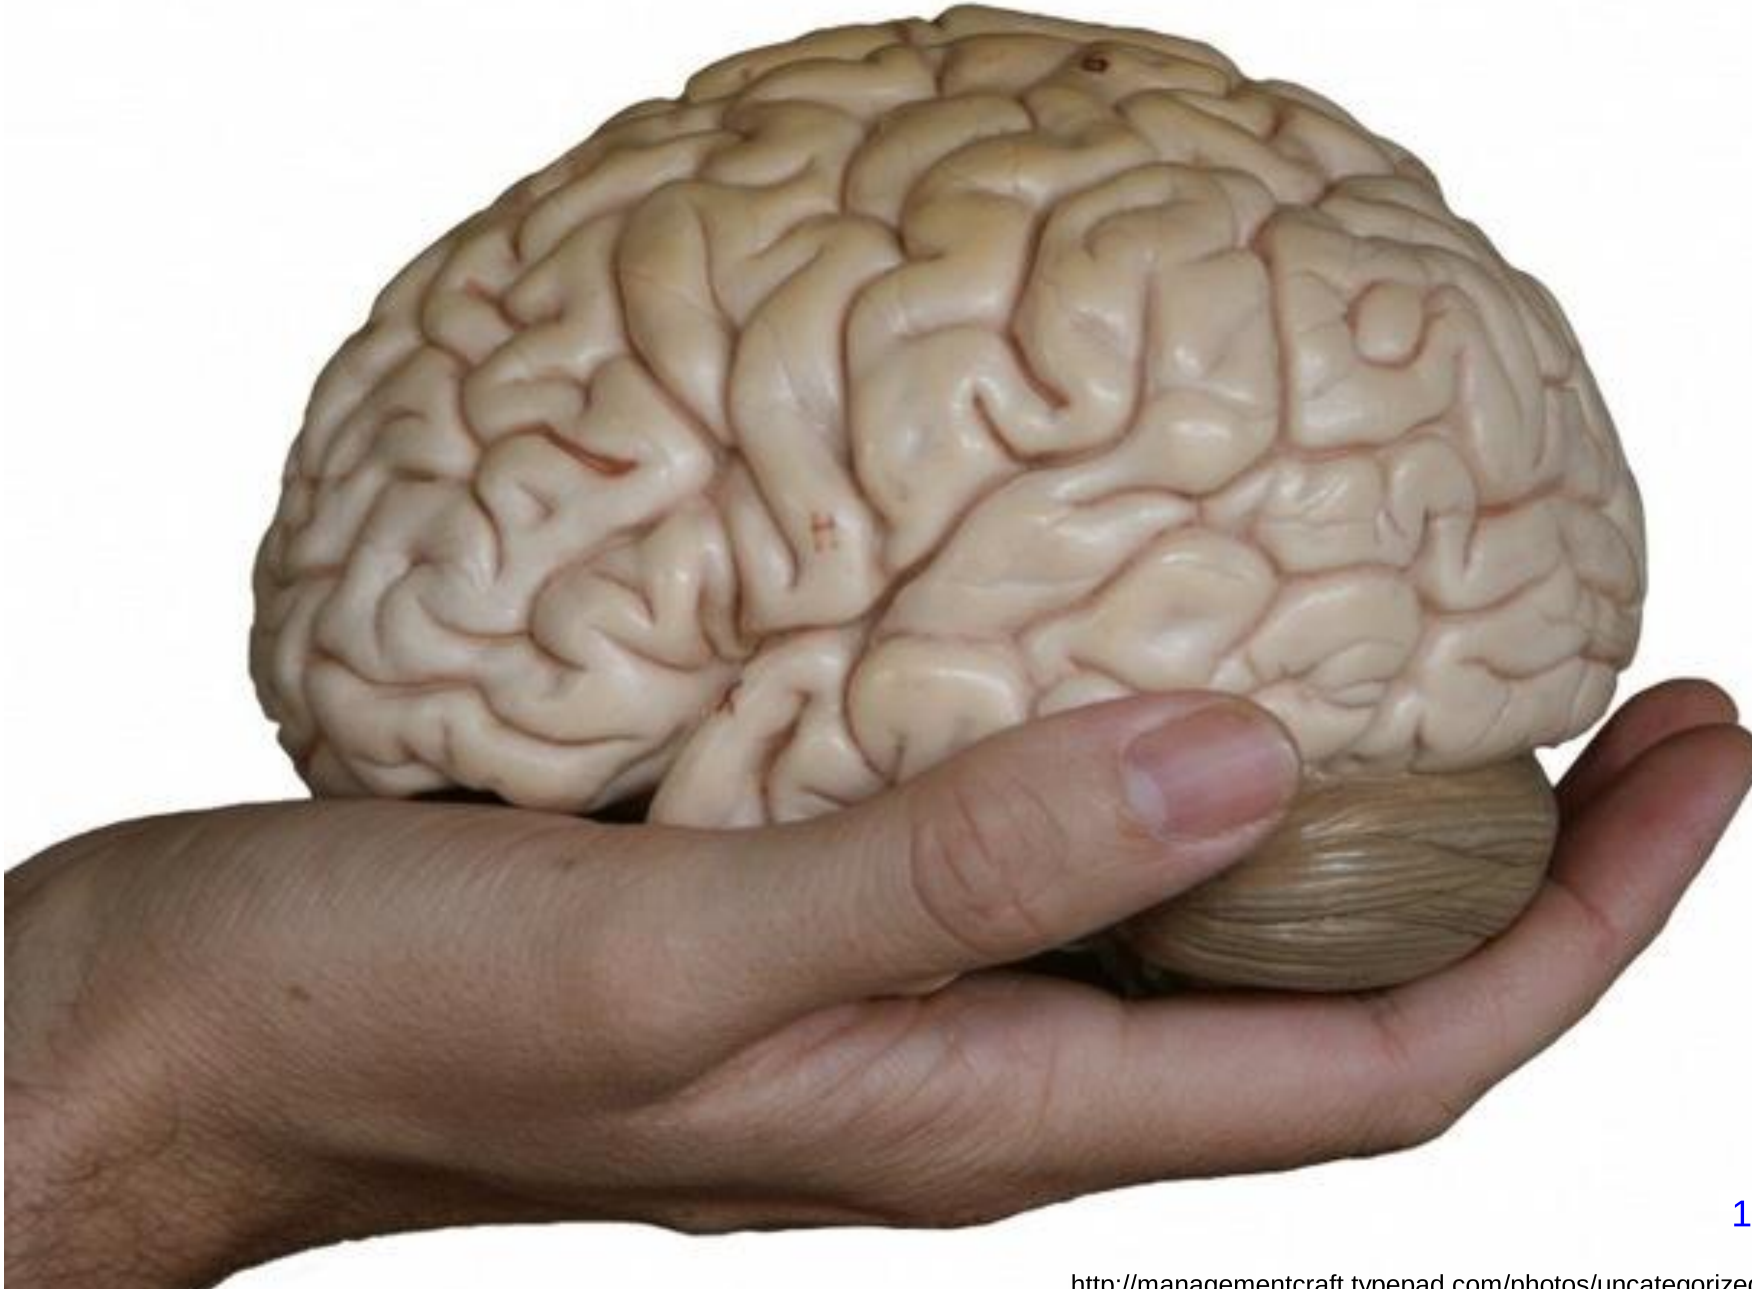

http://managementcraft.typepad.com/photos/uncategorized/brain.jpg

Biological Neural Networks

- Neuroscience:
  - Population of physically inter-connected neurons
- Includes:
  - Biological **Neurons**
  - Connecting **Synapses**
- The human brain:
  - 100 billion neurons
  - 100 trillion synapses

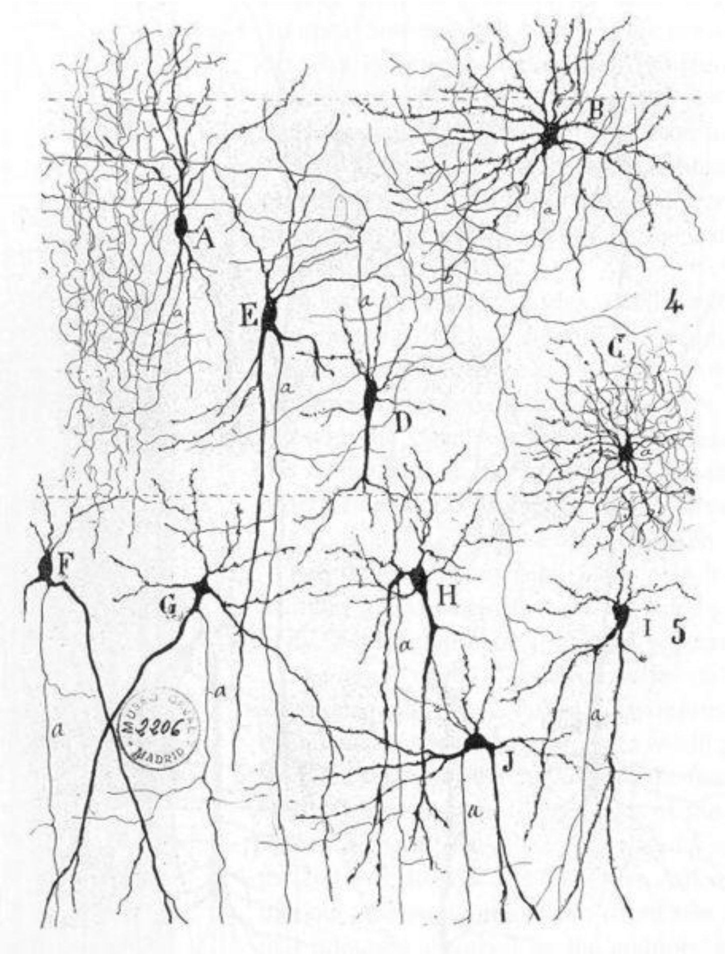

Biological Neuron

- Neurons:
  - Have *K* inputs (*dendrites*)
  - Have 1 output (*axon*)
  - If the sum of the input signals surpasses a *threshold*, sends an *action potential* to the axon
- Synapses
  - Transmit electrical signals between neurons

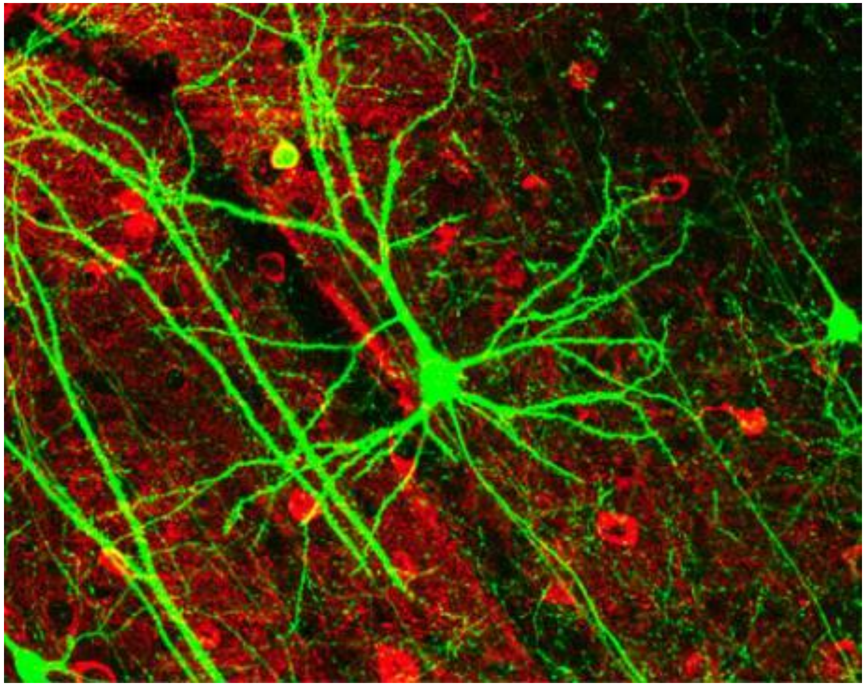

Artificial Neuron

- Also called the **McCulloch-Pitts neuron**
- Passes a **weighted sum of inputs**, to an **activation function**, which produces an **output** value

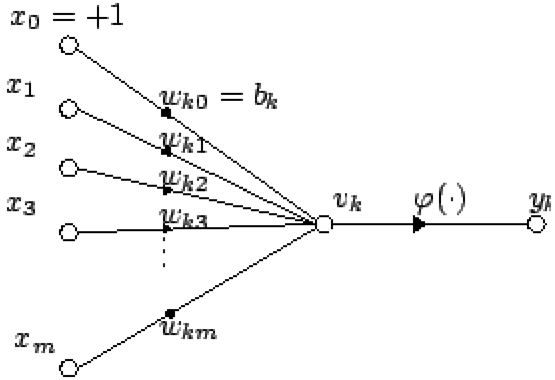

$$y_k = \varphi\left(\sum_{j=0}^{m} w_{kj} x_j\right)$$

McCulloch, W. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. Bulletin of Mathematical Biophysics, 7:115 - 133.

Sample activation functions

- Rectified Linear Unit (ReLU)

$$y = \begin{cases} u, & \text{if } u \geq 0 \\ 0, & \text{if } u < 0 \end{cases}, \quad u = \sum_{i=1}^{n} w_i x_i$$

- Sigmoid function

$$y = \frac{1}{1+e^{-u}}$$

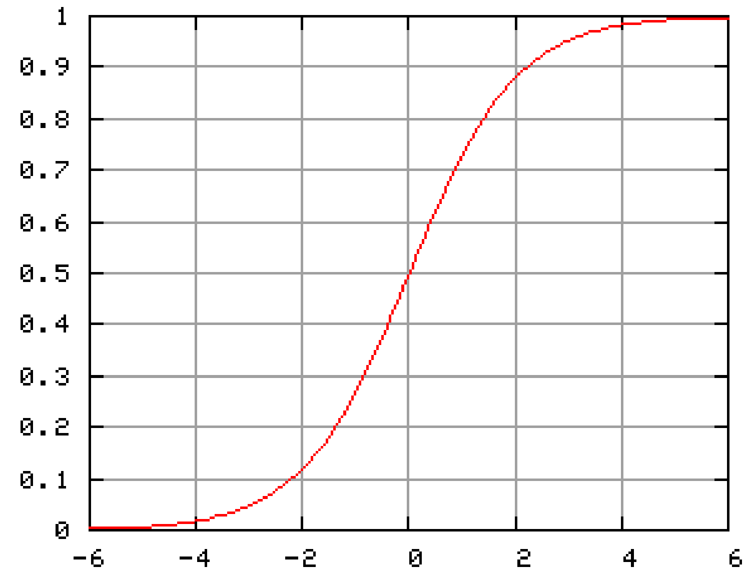

Artificial Neural Network

- Commonly refered as **Neural Network**
- Basic principles:
  - One neuron can perform a simple decision
  - Many **connected** neurons can make more **complex decisions**

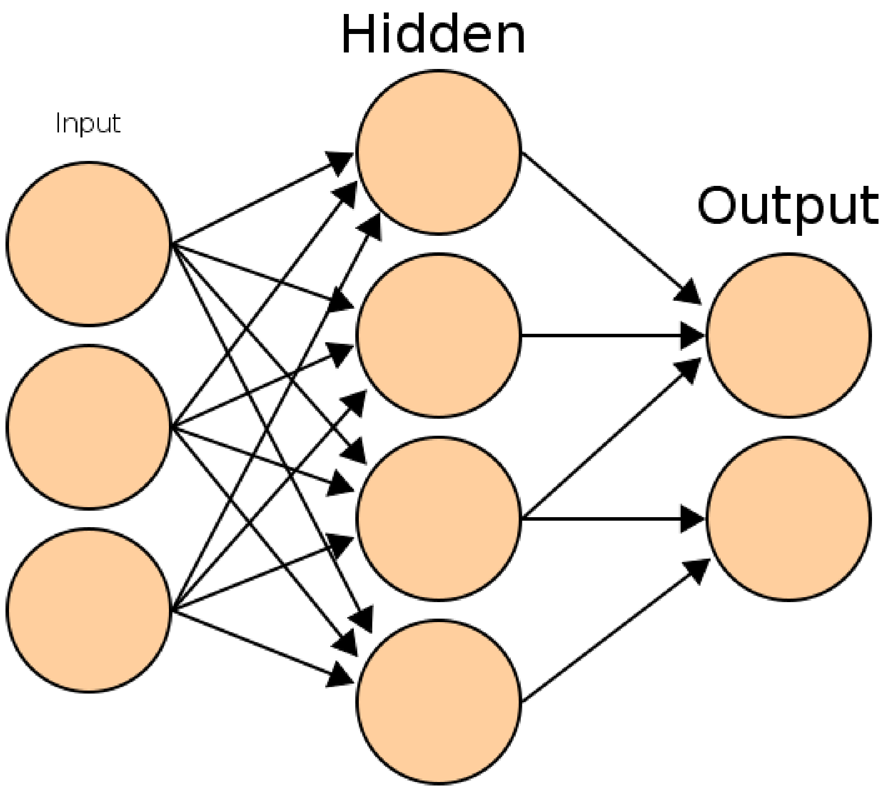

Characteristics of a NN

- Network configuration
  - How are the neurons inter-connected?
  - We typically use *layers* of neurons (input, output, hidden)
- Individual Neuron parameters
  - Weights associated with inputs
  - Activation function
  - Decision *thresholds*

How do we find these values?

Learning paradigms

- We can define the network configuration
- How do we define neuron ***weights*** and ***decision thresholds***?
  - **Learning** step
  - We **train** the NN to classify what we want
- Different learning paradigms
  - Supervised learning
  - Unsupervised learning
  - Reinforcement learning

Appropriate for **Pattern Recognition**.

Learning

- We want to obtain an **optimal solution** given a set of **observations**
- A **cost function** measures how close our solution is to the **optimal solution**
- Objective of our learning step:
  - Minimize the **cost function**

Backpropagation Algorithm

Backpropagation

Backpropagation

$$\text{Out}(\mathbf{x}) = g\left(\sum_j W_j\, g\left(\sum_k w_{jk} x_k\right)\right)$$

Find a set of weights $\{W_j\}, \{w_{jk}\}$

to minimize

$$\sum_i \left(y_i - \text{Out}(\mathbf{x}_i)\right)^2$$

by gradient descent.

That's it!
That's the backpropagation algorithm.

For more details please study Dr. Andrew Moore's excellent tutorials

http://www.cs.cmu.edu/~awm/tutorials.html

Feedforward neural network

- Simplest type of NN
- Has no *cycles*
- Input layer
  - Need as many neurons as coefficients of my *feature vector*
- Hidden layers
- Output layer
  - Classification results

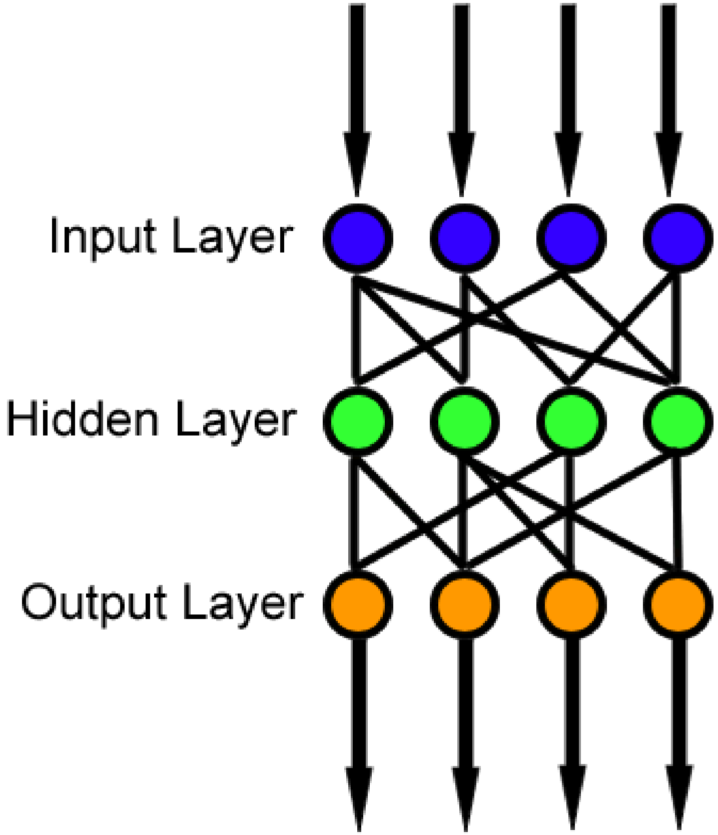

Topic: Support Vector Machines

- Statistical Classifiers
- Neural Networks
- Support Vector Machines

Maximum-margin hyperplane

- There are many planes that can separate our **classes** in **feature space**
- Only one **maximizes the separation margin**
- Of course that classes need to be separable in the first place...

$$f(x) : \Re^M \rightarrow \{+1, -1\}$$

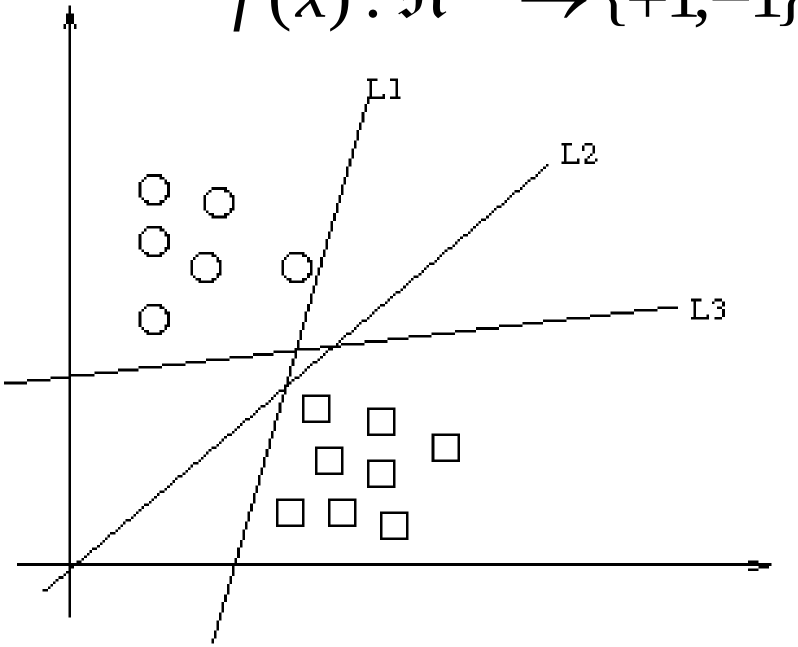

$$C \in \{+1, -1\}$$

$$V = \{v_1, v_2, \ldots, v_M\}, v_i \in \Re^M$$

Support vectors

- The **maximum-margin hyperplane** is limited by some vectors
- These are called **support vectors**
- Other vectors are irrelevant for my decision

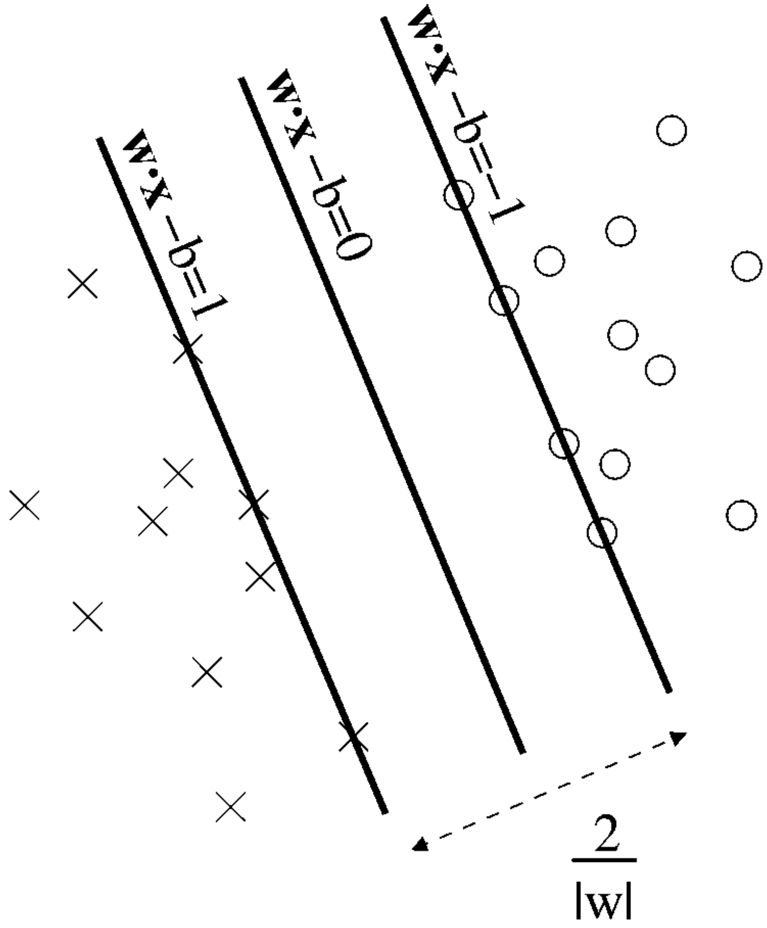

Decision

- I map a new **observation** into my **feature space**
- Decision hyperplane:

$$(w.x)+b=0, w\in\Re^N, b\in\Re$$

- Decision function:

$$f(x)=sign((w.x)+b)$$

A vector is either **above** or **below** the hyperplane

Slack variables

- Most **feature spaces** cannot be segmented so easily by a hyperplane
- Solution:
  - Use slack variables
  - ‘Wrong’ points ‘pull’ the margin in their direction
  - Classification errors!

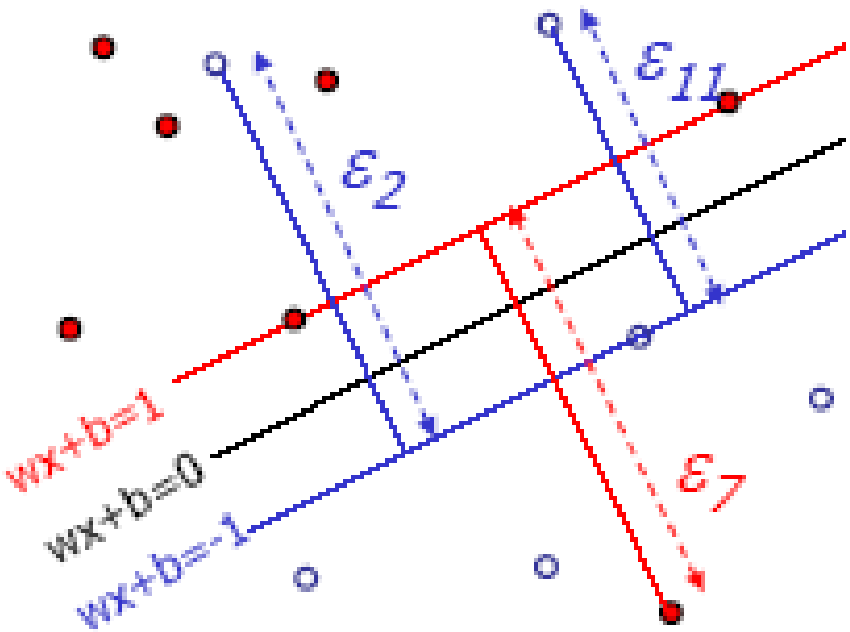

But this doesn't work in most situations...

- Still, how do I find a **Maximum-margin hyperplane** for some situations?

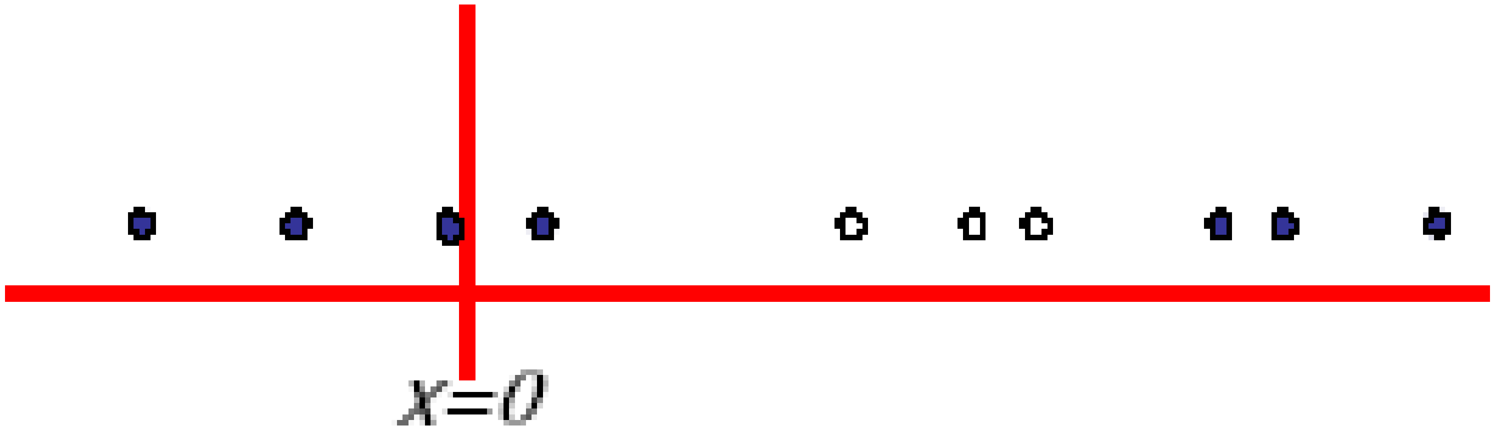

- Most real situations face this problem...

Solution: Send it to hyperspace!

- Take the previous case: **f(x) = x**
- Create a new higher-dimensional function: **$g(x^2) = (x, x^2)$**
- A **kernel function** is responsible for this transformation

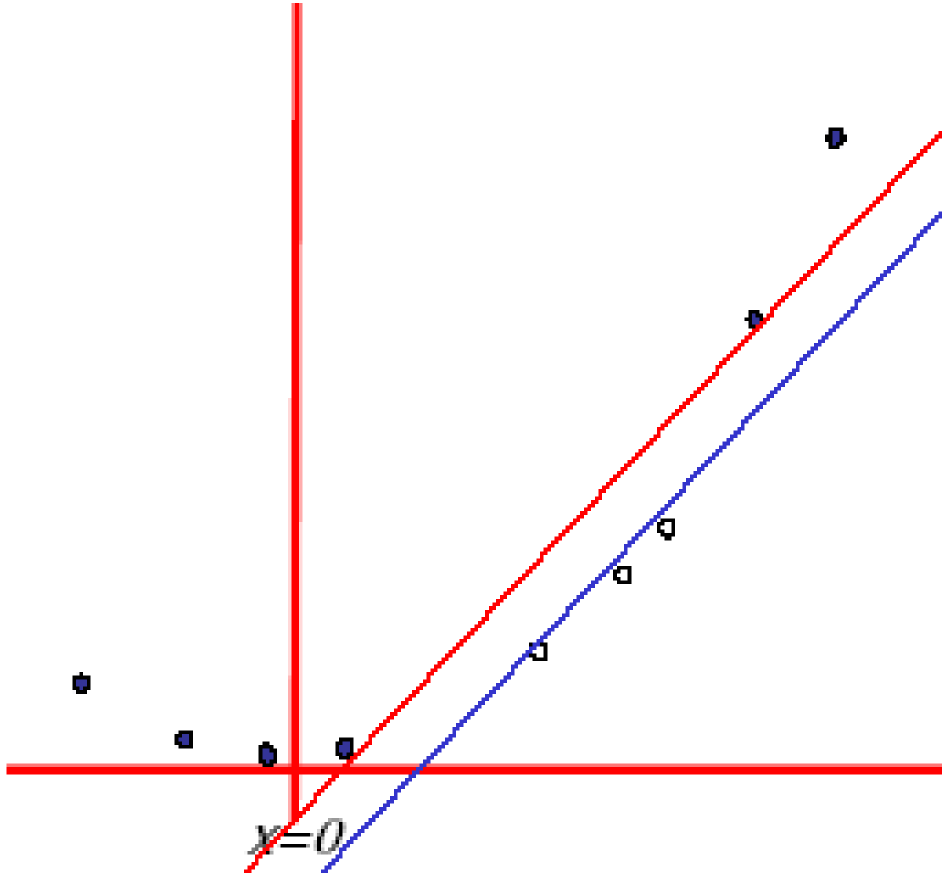

https://www.youtube.com/watch?v=3liCbRZPrZA

Typical kernel functions

Linear	$K(x,y) = x.y + 1$
Polynomial	$K(x,y) = (x.y + 1)^p$
Radial-Base Functions	$K(x,y) = e^{-\lVert x-y \rVert^2 / 2\sigma^2}$
Sigmoid	$K(x,y) = \tanh(kx.y - \delta)$

Classification

- Training stage:
  - Obtain kernel parameters
  - Obtain maximum-margin hyperplane
- Given a new **observation**:
  - Transform it using the kernel
  - Compare it to the hyperspace

Resources

- Andrew Moore, “Statistical Data Mining Tutorials”, http://www.cs.cmu.edu/~awm/tutorials.html
- C.J. Burges, “A tutorial on support vector machines for pattern recognition”, in Knowledge Discovery Data Mining, vol.2, no.2, 1998, pp.1-43.