TECNICHE DI DIAGNOSI AUTOMATICA DEI GUASTI

Reti Neurali per l'Identificazione di Sistemi non Lineari e Pattern Recognition

Silvio Simani

silvio.simani@unife.it

References

Textbook (*suggested*):

- *Neural Networks for Identification, Prediction, and Control*, by Duc Truong Pham and Xing Liu. Springer Verlag; (December 1995). ISBN: 3540199594
- *Nonlinear Identification and Control: A Neural Network Approach*, by G. P. Liu. Springer Verlag; (October 2001). ISBN: 1852333421.
- *Fuzzy Modeling for Control*, by Robert Babuska. Springer; 1st edition (May 1, 1998) ISBN-10: 0792381548, ISBN-13: 978-0792381549.

Course Overview

1. Introduction
   1. Course introduction
   2. Introduction to neural network
   3. Issues in neural network
2. Simple neural network
   1. Perceptron
   2. Adaline
3. Multilayer Perceptron
   1. Basics
4. Radial basis networks: overview
5. Fuzzy Systems: overview
6. Application examples

Machine Learning

- Improve automatically with experience
- Imitating human learning
  - Human learning

    Fast recognition and classification of complex classes of objects and concepts and fast adaptation
  - Example: neural networks
- Some techniques assume statistical source

  Select a statistical model to model the source
- Other techniques are based on reasoning or inductive inference (e.G. Decision tree)

Machine Learning Definition

A computer program is said to **learn** from *experience* **E** with respect to some class of *tasks* **T** and *performance measure* **P**, if its performance at tasks in **T**, as measured by **P**, improves with experience.

Examples of Learning Problems

Example 1: handwriting recognition:

- T: recognizing and classifying handwritten words within images.
- P: percentage of words correctly classified.
- E: a database of handwritten words with given classification.

Example 2: learn to play checkers:

- T: play checkers.
- P: percentage of games won in a tournament.
- E: opportunity to play against itself (war games…).

Type of Training Experience

- **Direct or indirect**?
  - Direct: board state -> correct move
  - Indirect: credit assignment problem (degree of credit or blame for each move to the final outcome of win or loss)
- **Teacher or not** ?
  - Teacher selects board states and provide correct moves **or**
  - Learner can select board states
- **Is training experience representative of performance goal**?
  - Training playing against itself
  - Performance evaluated playing against world champion

Issues in *Machine Learning*

- What algorithms can approximate functions well and when?
- How does the number of training examples influence accuracy?
- How does the complexity of hypothesis representation impact it?
- How does noisy data influence accuracy?
- *How do you reduce a learning problem to a set of function approximation* ?

Summary

- *Machine learning* is useful for data mining, poorly understood domain (face recognition) and programs that must dynamically adapt.
- Draws from many diverse disciplines.
- Learning problem needs well-specified task, performance metric and training experience.
- Involve searching space of possible hypotheses. Different learning methods search different hypothesis space, such as numerical functions, *neural networks*, decision trees, symbolic rules.

Introduction to Neural Networks

Brain

- 10^{11} neurons (processors)
- On average 1000-10000 connections

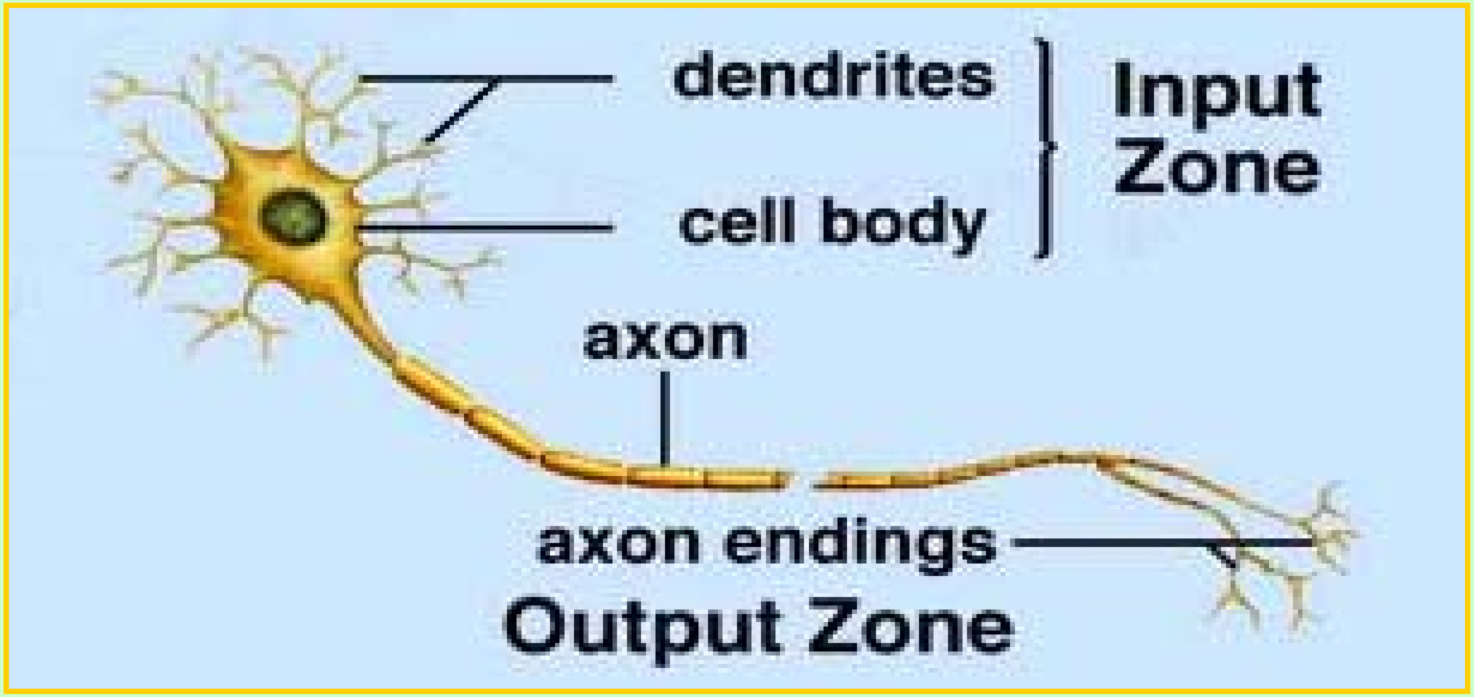

Artificial Neuron

bias

$$net_i = \sum_j w_{ij} y_j + b$$

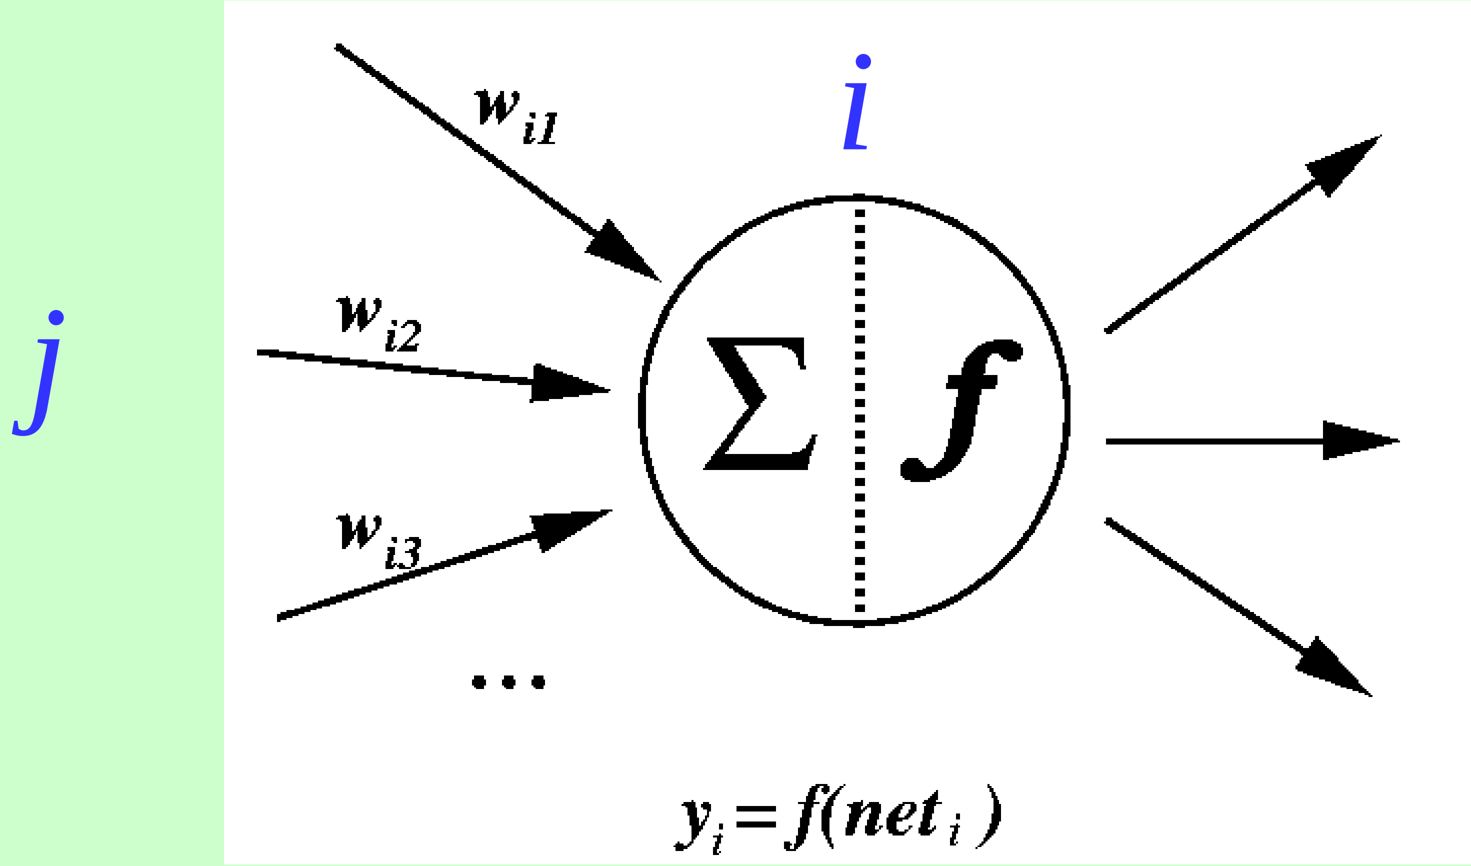

Artificial Neuron

- Input/Output Signal may be.
  - Real value.
  - Unipolar {0, 1}.
  - Bipolar {-1, +1}.
- Weight : w_{ij} – strength of connection.

Note that w_{ij} refers to the weight from **unit *j* to unit *i*** (not the other way round).

Artificial Neuron

- The bias *b* is a constant that can be written as $w_{i0}y_0$ with $y_0 = b$ and $w_{i0} = 1$ such that

$$net_i = \sum_{j=0}^{n} w_{ij} y_j$$

- The function *f* is the unit's activation function. In the simplest case, *f* is the identity function, and the unit's output is just its net input. This is called a *linear unit*.
- Other activation functions are : step function, sigmoid function and Gaussian function.

Activation Functions

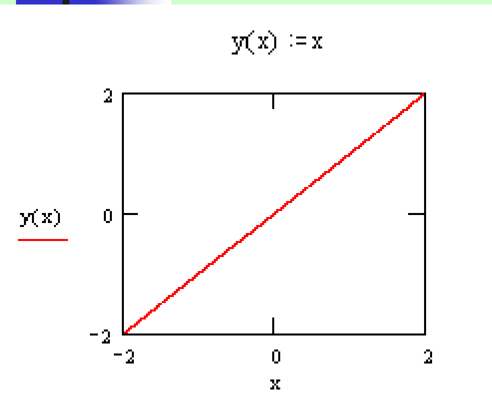

Identity function

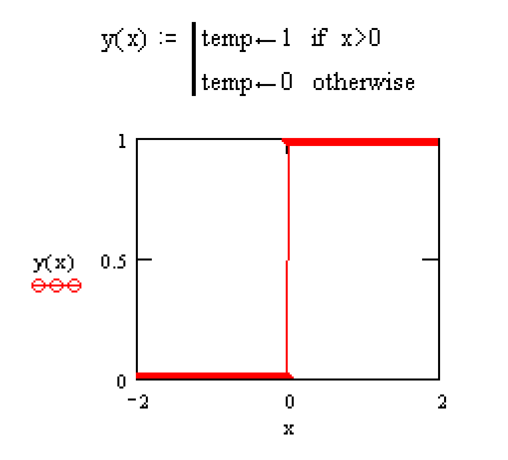

Binary Step function

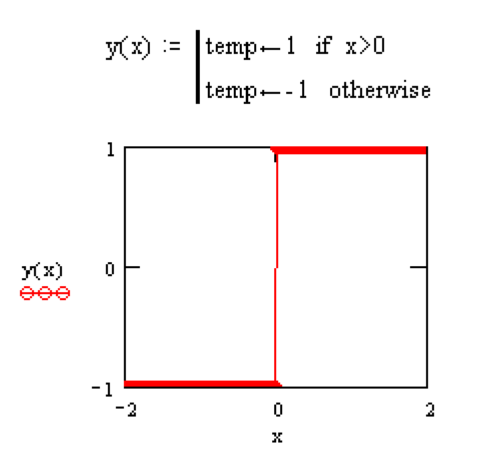

Bipolar Step function

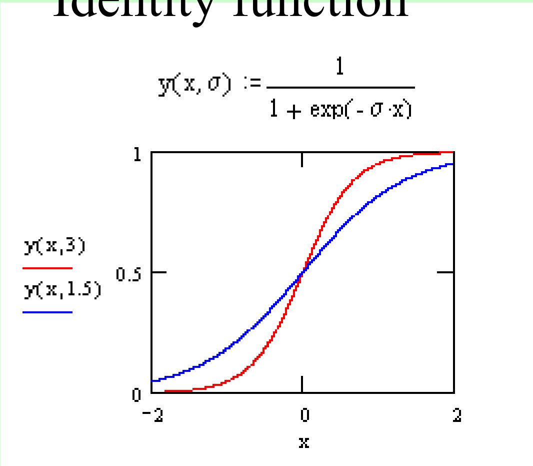

Sigmoid function

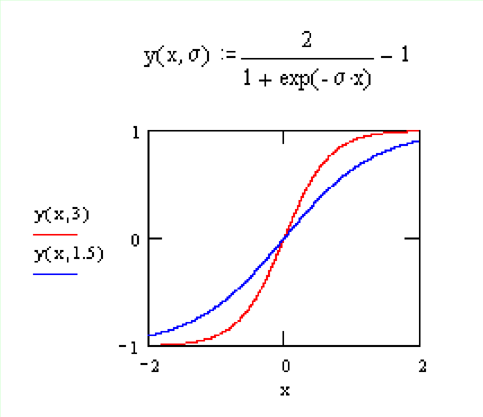

Bipolar Sigmoid function

$$y(x) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

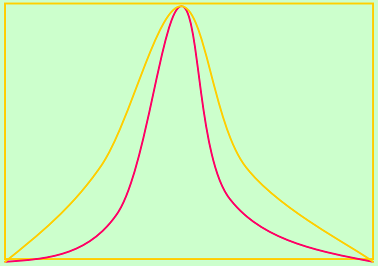

Gaussian function

Artificial Neural Networks (ANN)

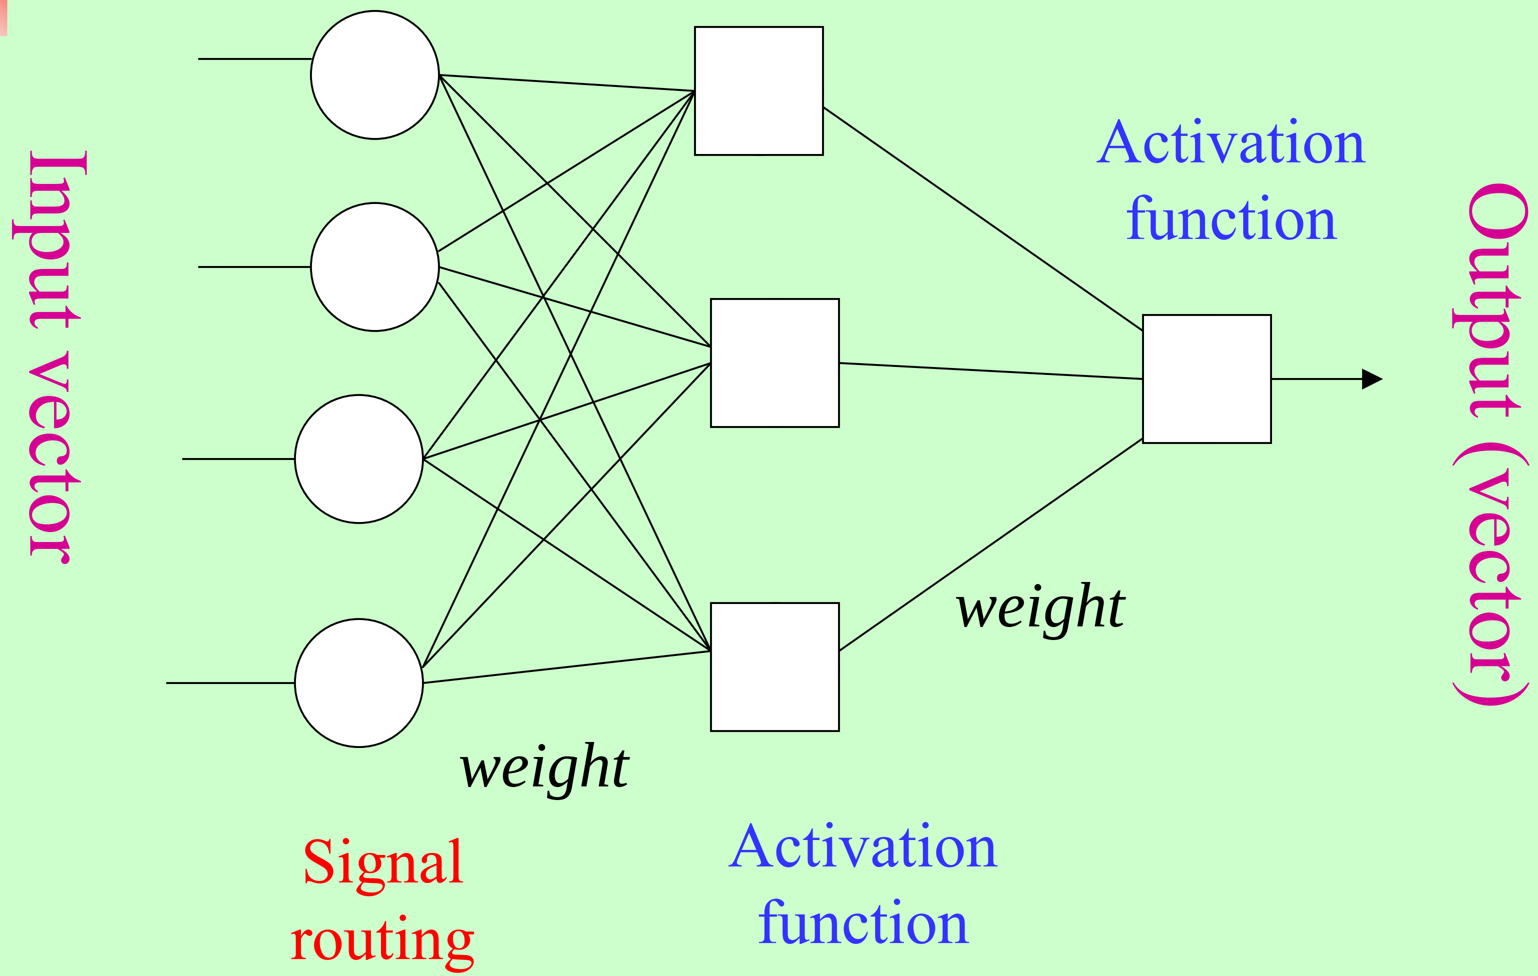

When Should ANN Solution Be Considered ?

- The solution to the problem cannot be explicitly described by an algorithm, a set of equations, or a set of rules.
- There is some evidence that an input-output mapping exists between a set of input and output variables.
- There should be a large amount of data available to train the network.

Problems That Can Lead to Poor Performance ?

- The network has to distinguish between very similar cases with a very high degree of accuracy.
- The train data does not represent the ranges of cases that the network will encounter in practice.
- The network has a several hundred inputs.
- The main discriminating factors are not present in the available data. *E.g.* Trying to assess the loan application without having knowledge of the applicant's salaries.
- The network is required to implement a very complex function.

Applications of Artificial Neural Networks

- Manufacturing : fault diagnosis, fraud detection.
- Retailing : fraud detection, forecasting, data mining.
- Finance : fraud detection, forecasting, data mining.
- Engineering : fault diagnosis, signal/image processing.
- Production : fault diagnosis, forecasting.
- Sales & marketing : forecasting, data mining.

Data Pre-processing

Neural networks very **rarely** operate on the raw data. An initial **pre-processing** stage is essential.

Some examples are as follows:

- Feature extraction of images: for example, the analysis of x-rays requires pre-processing to extract features which may be of interest within a specified region.
- Representing input variables with numbers. For example "+1" is the person is married, "0" if divorced, and "-1" if single. Another example is representing the pixels of an image: 255 = bright white, 0 = black. To ensure the generalization capability of a neural network, the data should be encoded in form which allows for interpolation.

Data Pre-processing

- **Categorical Variable**
  - A categorical variable is a variable that can belong to one of a number of discrete categories. For example, red, green, blue.
  - Categorical variables are usually encoded using 1 out-of n coding. e.g. for three colors, red = (1 0 0), green =(0 1 0) Blue =(0 0 1).
  - If we used red = 1, green = 2, blue = 3, then this type of encoding imposes an ordering on the values of the variables which does not exist.

Data Pre-processing

- **CONTINUOUS VARIABLES**
  - A continuous variable can be directly applied to a neural network. However, if the dynamic range of input variables are not approximately the same, it is better to *normalize* all input variables of the neural network.

Simple Neural Networks

Simple Perceptron

Outlines

- The Perceptron
- Linearly separable problem
- Network structure
- Perceptron learning rule
- Convergence of Perceptron

THE PERCEPTRON

- The perceptron was a simple model of ANN introduced by Rosenblatt of MIT in the 1960' with the idea of learning.

- Perceptron is designed to accomplish a simple pattern recognition task: after learning with real value training data

$$\{\, \underline{x}(i), d(i),\ i = 1,2, \ldots, p \,\} \qquad \text{where } d(i) = 1 \text{ or } -1$$

- For a new signal (pattern) $\underline{x}(i+1)$, the perceptron is capable of telling you to which class the new signal belongs

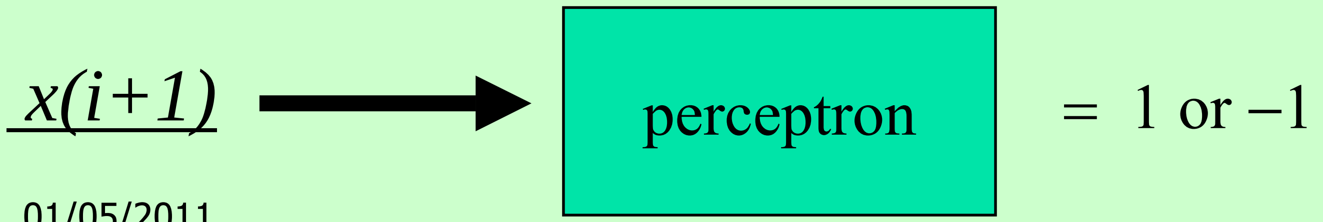

Perceptron

- Linear Threshold Unit (LTU)

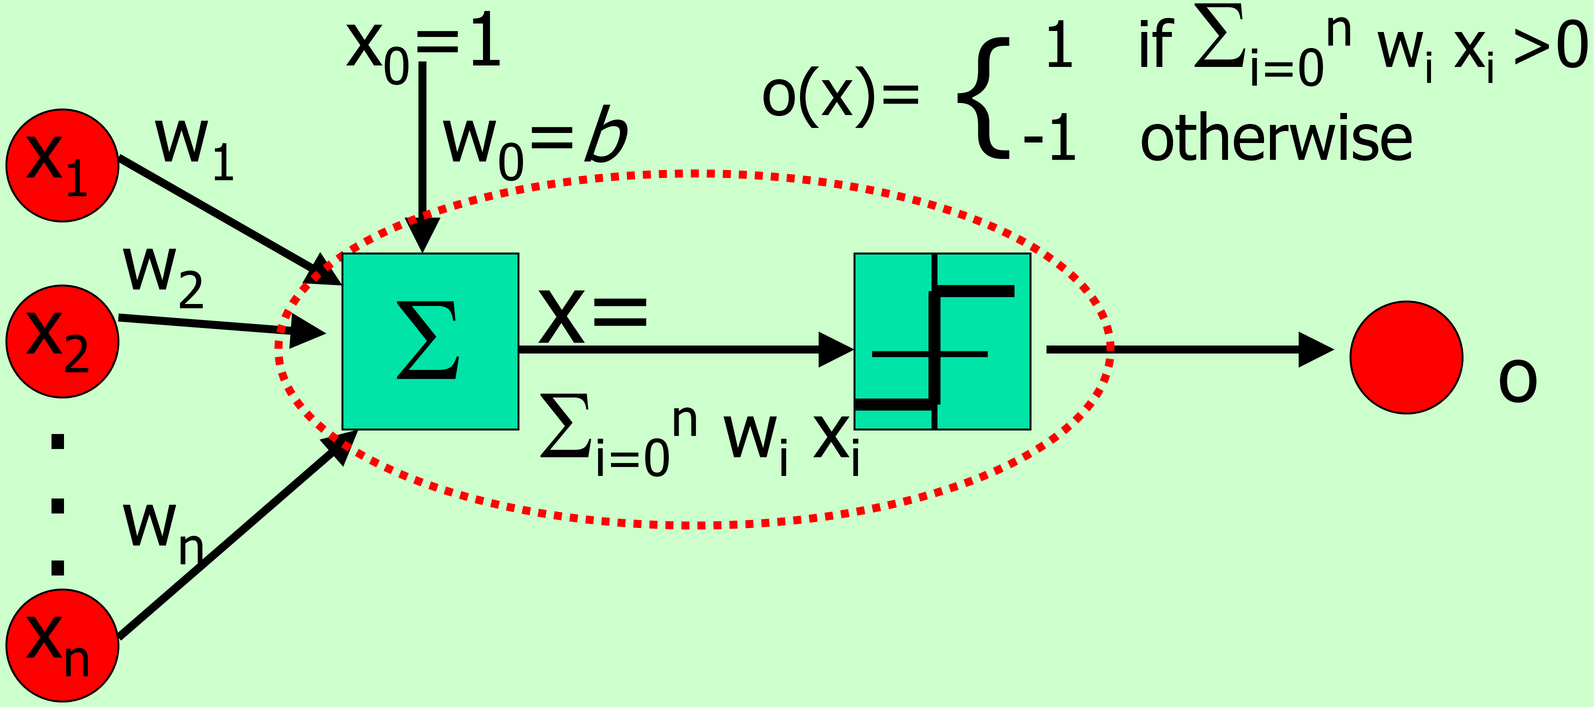

Mathematically the Perceptron is

$$y = f\left(\sum_{i=1}^{m} w_i x_i + b\right) = f\left(\sum_{i=0}^{m} w_i x_i\right)$$

We can always treat the bias *b* as another weight with inputs equal 1

where f is the hard limiter function i.e.

$$y = \begin{cases} 1 \; if \; \sum_{i=1}^{m} w_i x_i + b > 0 \\ -1 \; if \; \sum_{i=1}^{m} w_i x_i + b \leq 0 \end{cases}$$

Why is the network capable of solving linearly separable problem ?

$$\sum_{i=1}^{m} w_i x_i + b = 0$$

$$\sum_{i=1}^{m} w_i x_i + b < 0$$

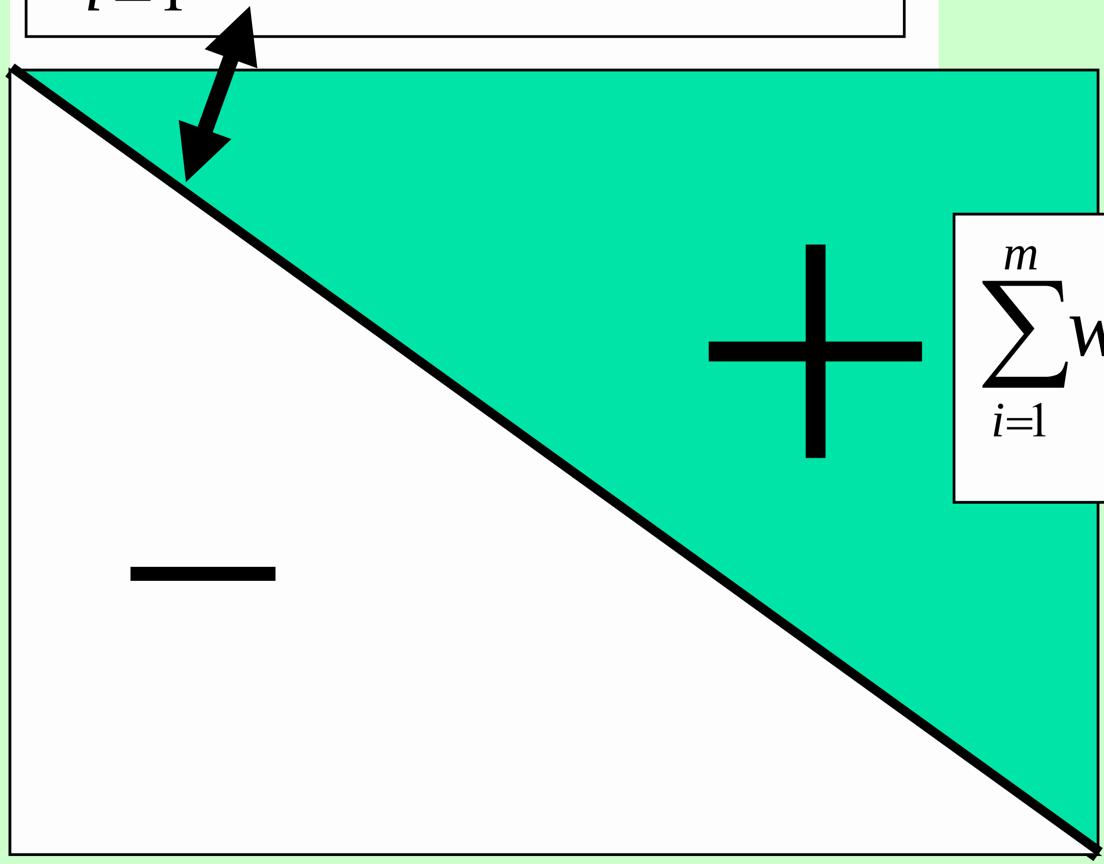

$$\sum_{i=1}^{m} w_i x_i + b > 0$$

Learning rule

An algorithm to update the weights $\underline{w}$ so that finally the input patterns lie on both sides of the line decided by the perceptron

Let t be the time, at $t = 0$, *we have*

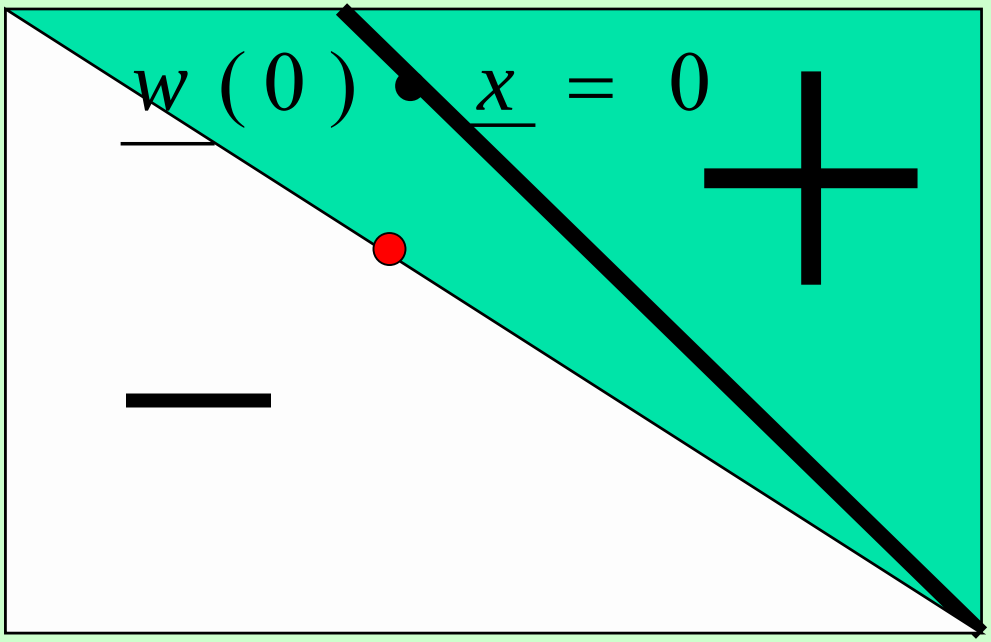

In Math

$$d(t) = \begin{cases} +1 \; if \;\; x(t) \, in \;\; class \;\;\; + \\ -1 \; if \;\; x(t) \, in \;\; class \;\;\; - \end{cases}$$

Perceptron learning rule

$$\underline{w}(t+1) = \underline{w}(t) + \eta(t)[d(t) - sign(\underline{w}(t) \bullet \underline{x}(t))] \, \underline{x}(t)$$

Where $\eta(t)$ is the learning rate >0,

$$sign(x) = \begin{cases} +1 \text{ if } x>0 \\ -1 \text{ if } x<=0, \end{cases}$$

hard limiter function

<u>NB : d(t) is the same as d(i) and x(t) as x(i)</u>

In words:

- If the classification is right, do not update the weights

- If the classification is not correct, update the weight towards the opposite direction so that the output move close to the right directions.

Perceptron convergence theorem (Rosenblatt, 1962)

Let the subsets of training vectors be linearly separable. Then after finite steps of learning we have

lim $\underline{w}(t) = \underline{w}$ **which correctly separate the samples.**

The idea of proof is that to consider $||\underline{w}(t+1)-\underline{w}||-||\underline{w}(t)-\underline{w}||$ which is a decrease function of t

Summary of Perceptron learning …

Variables and parameters

$\underline{x}(t)$ = $(m+1)$ *dim.* input vectors at time t

$= (b, x_1(t), x_2(t), \ldots , x_m(t))$

$\underline{w}(t) = (m+1)$ *dim.* weight vectors

$= (1 , w_1(t), \ldots , w_m(t))$

b = bias

$y(t)$ = actual response

$\eta(t)$ = learning rate parameter, a +ve constant < 1

$d(t)$ = desired response

Summary of Perceptron learning …

Data { (x(i), d(i)), i=1,…,p}

- Present the data to the network once a point
- could be cyclic :

(x(1), d(1)), (x(2), d(2)),…, (x(p), d(p)), (x(p+1), d(p+1)),…

- *or randomly*

(Hence we mix time t with i here)

Summary of Perceptron learning (algorithm)

1. Initialisation Set $\underline{w}(0)=0$. Then perform the following computation for time step t=1,2,...

2. Activation At time step t, activate the perceptron by applying input vector $\underline{x}(t)$ and desired response $d(t)$

3. Computation of actual response Compute the actual response of the perceptron

$$y(t) = sign\left(\underline{w}(t) \cdot \underline{x}(t) \right)$$

where ***sign*** is the sign function

4. Adaptation of weight vector Update the weight vector of the perceptron

$$\underline{w}(t+1) = \underline{w}(t) + \eta(t) \left[d(t) - y(t) \right] \underline{x}(t)$$

5. Continuation

Questions remain

Where or when to stop?

By minimizing the generalization error

For training data $\{(\underline{x}(i), d(i)), i=1,\ldots p\}$

How to define training error after t steps of learning?

$$E(t)= \sum_{i=1}^{p} [d(i)-\mathrm{sign}(\underline{w}(t) \,.\, \underline{x}(i)]^2$$

We next turn to ADALINE learning, from which we can understand the learning rule, and more general the Back-Propagation (BP) learning

Simple Neural Network

ADALINE Learning

Outlines

- ADALINE
- Gradient descending learning
- Modes of training

Unhappy Over Perceptron Training

- When a perceptron gives the right answer, no learning takes place
- Anything below the threshold is interpreted as ‘no’, even it is just below the threshold.
- It might be better to train the neuron based on how far below the threshold it is.

ADALINE

- ADALINE is an acronym for ADAptive LINear Element (or ADAptive LInear NEuron) developed by Bernard Widrow and Marcian Hoff (1960).
- There are several variations of Adaline. One has threshold same as perceptron and another just a bare linear function.
- The Adaline learning rule is also known as the least-mean-squares (LMS) rule, the delta rule, or the Widrow-Hoff rule.
- It is a training rule that minimizes the output error using (approximate) gradient descent method.

- Replace the step function in the perceptron with a continuous (differentiable) function *f*, e.g the simplest is linear function

- With or without the threshold, the Adaline is trained based on the output of the function *f* rather than the final output.

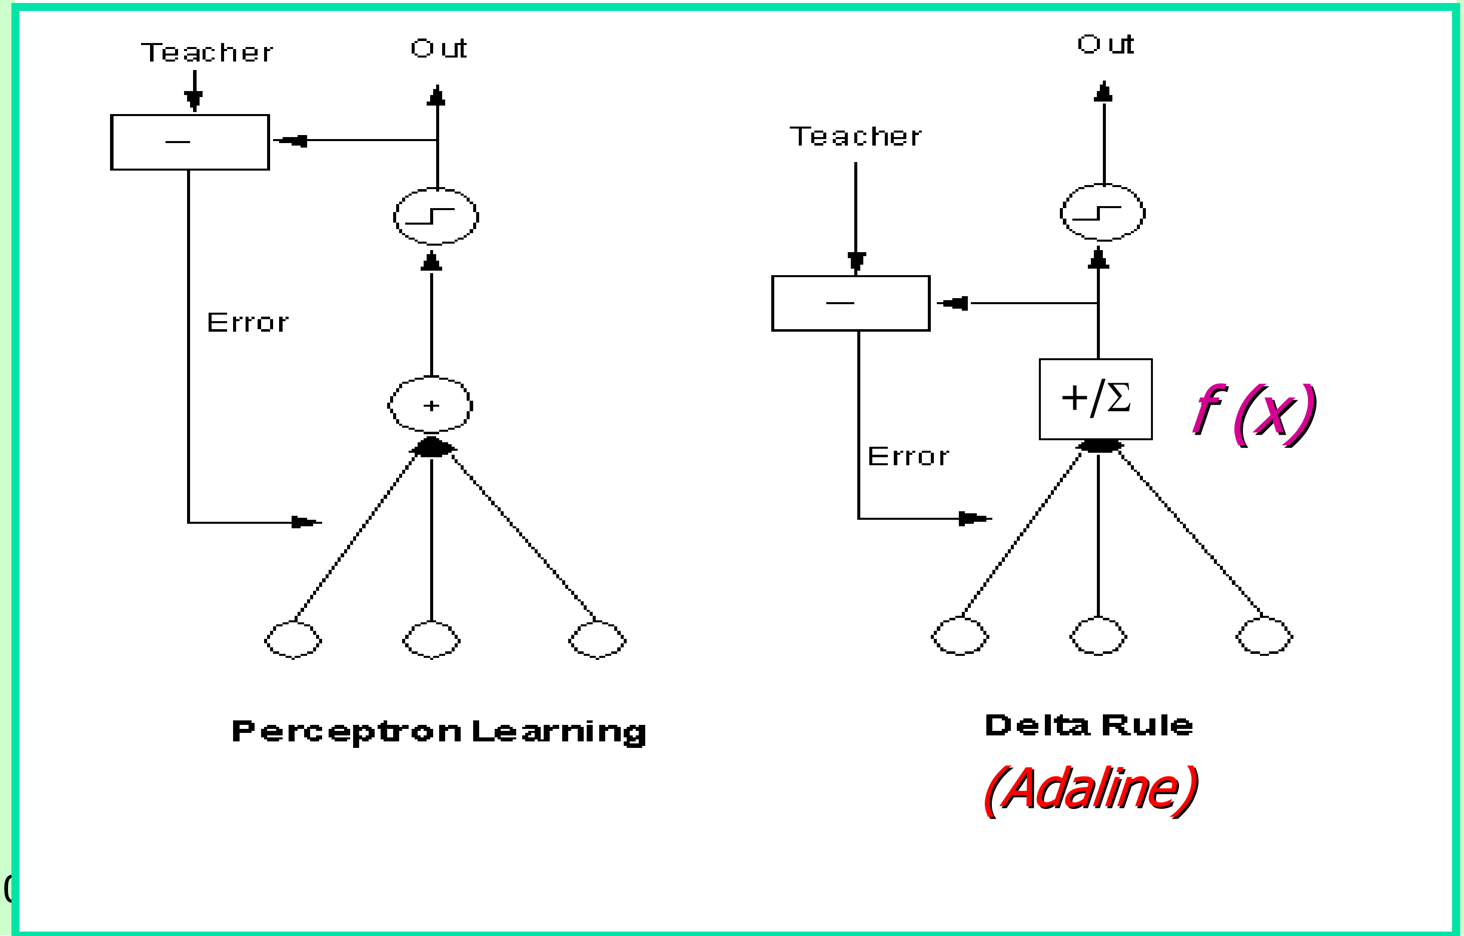

After each training pattern $\underline{x}(i)$ is presented, the correction to apply to the weights is proportional to the error.

$$E(i,t) = \tfrac{1}{2} \left[d(i) - f(\underline{w}(t) \cdot \underline{x}(i)) \right]^2 \qquad i=1,\ldots,p$$

N.B. If f is a linear function $f(\underline{w}(t) \cdot \underline{x}(i)) = \underline{w}(t) \cdot \underline{x}(i)$

Summing together, our purpose is to find $\underline{W}$ which minimizes

$$E(t) = \sum_i E(i,t)$$

General Approach gradient descent method

To find g

$$\underline{w}(t+1) = \underline{w}(t) + g(\, E(\underline{w}(t))\,)$$

so that $\underline{w}$ automatically tends to the global minimum of E(w).

$$\underline{w}(t+1) = \underline{w}(t) - E'(\underline{w}(t))\eta(t)$$

(see figure below)

- Gradient direction is the direction of uphill for example, in the Figure, at position 0.4, the gradient is uphill (F is E, consider one dim case)

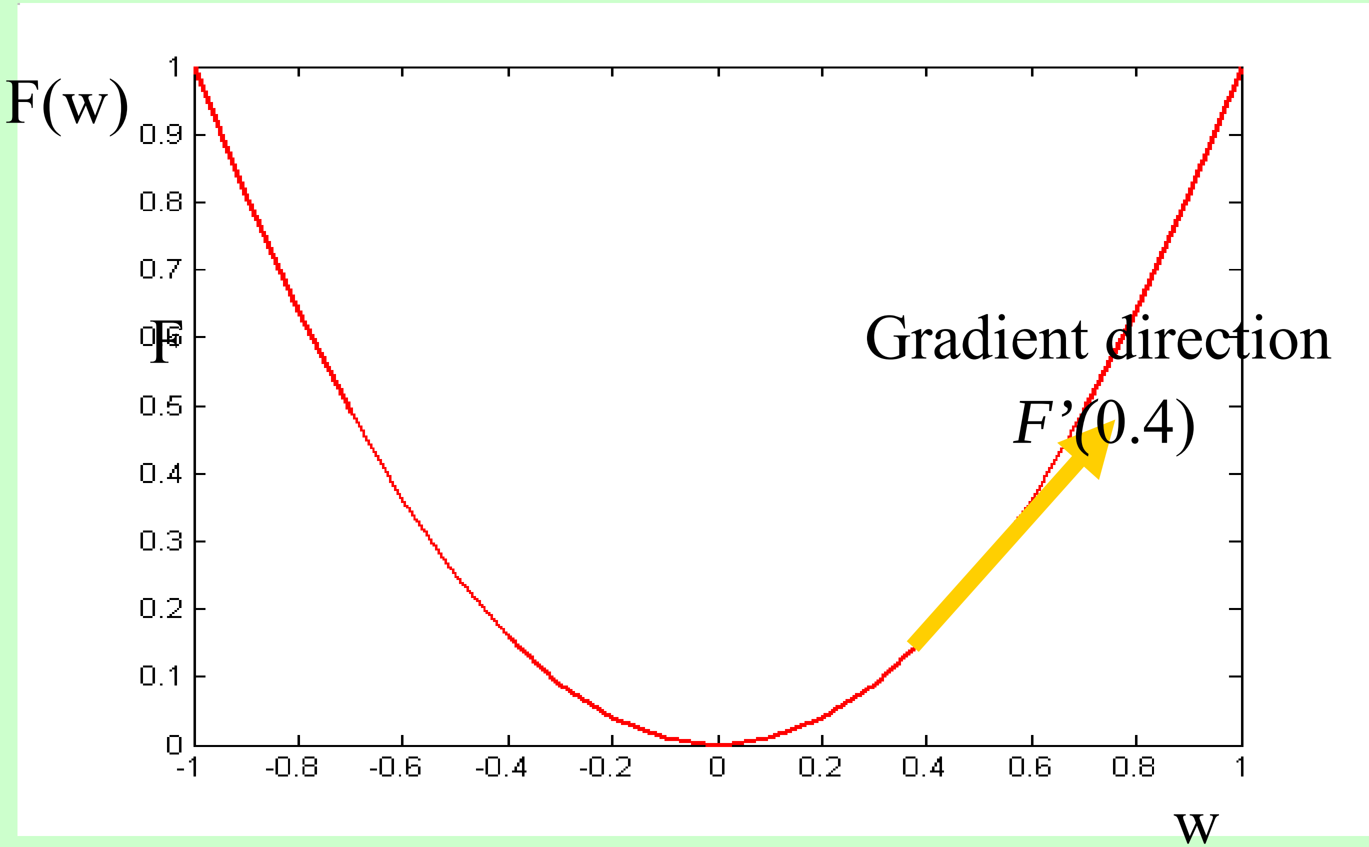

- In gradient descent algorithm, we have

$$\underline{w}(t+1) = \underline{w}(t) - F'(w(t))\,\eta(\tau)$$

therefore the ball goes downhill since – F'(w(t)) is downhill direction

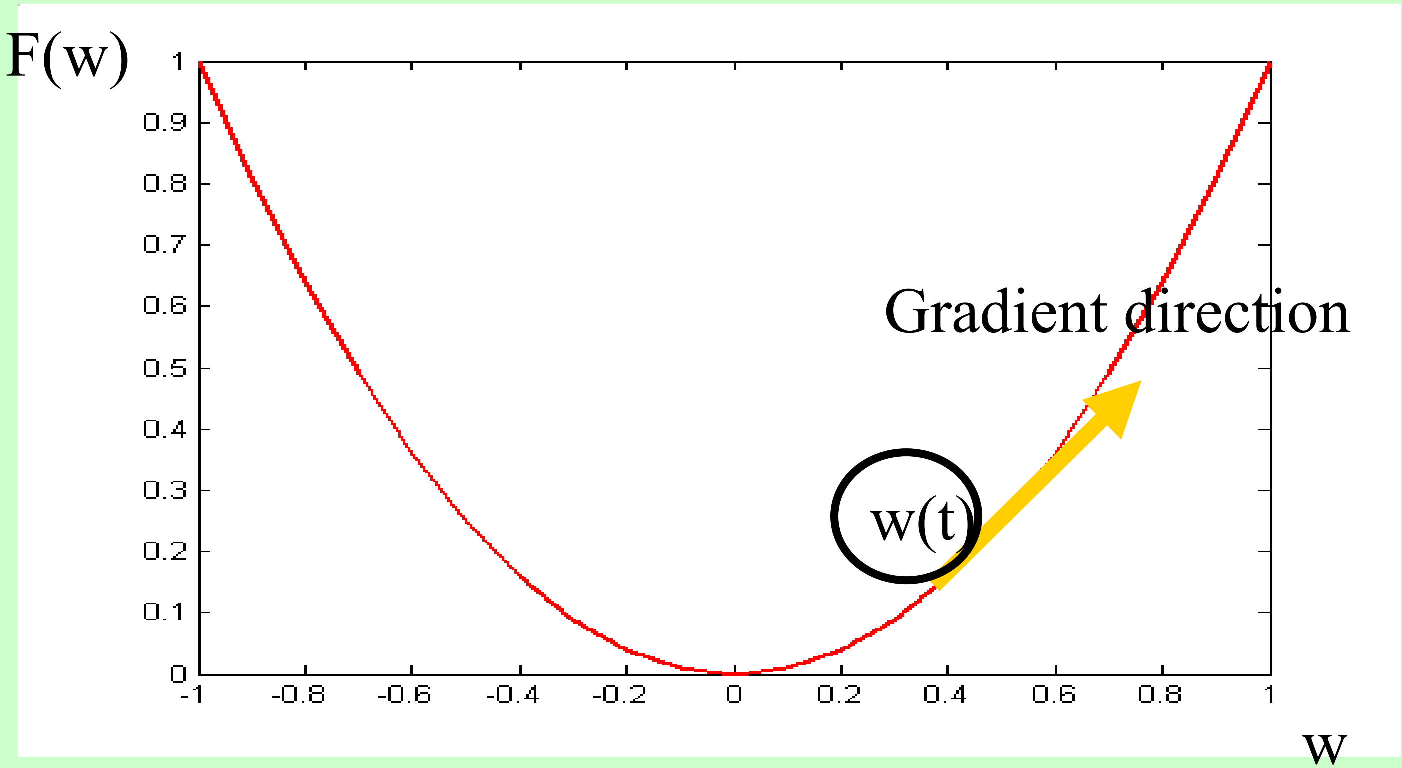

- In gradient descent algorithm, we have

$$w(t+1) = w(t) - F'(w(t))\,\eta(\tau)$$

therefore the ball goes downhill since – F′(w(t)) is downhill direction

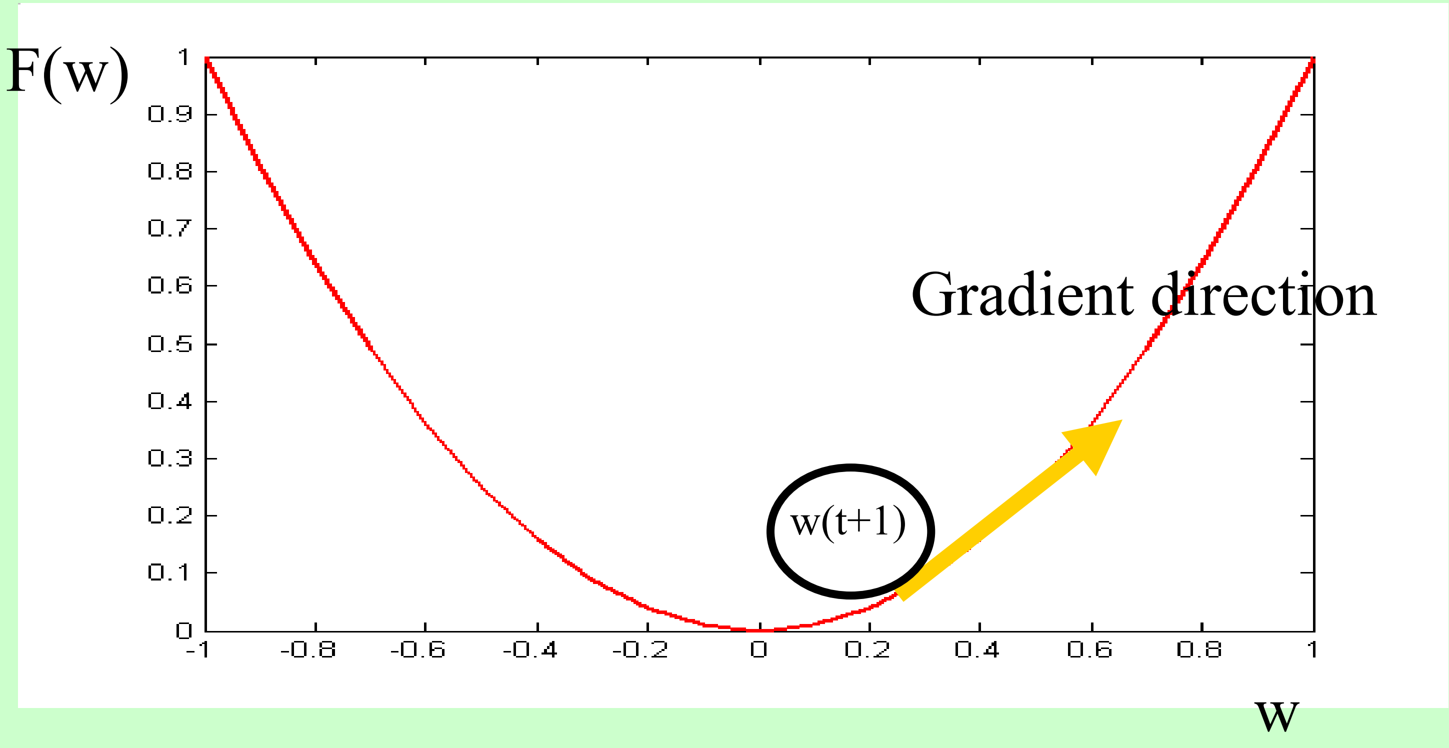

- Gradually the ball will stop at a local minima where the gradient is zero

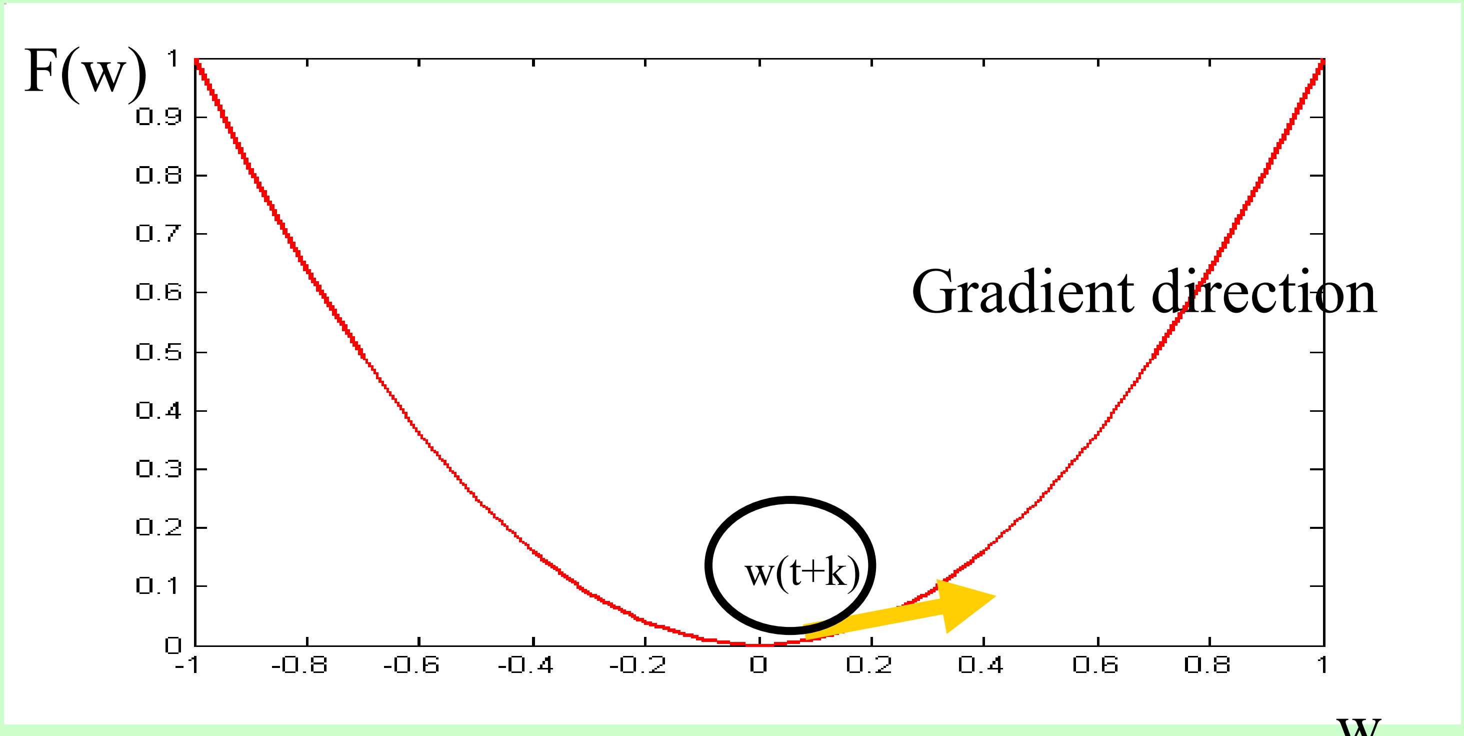

• In words

Gradient method could be thought of as a ball rolling down from a hill: the ball will roll down and finally stop at the valley

Thus, the weights are adjusted by

$$w_j(t+1) = w_j(t) + \eta(t) \Sigma \; [d(i) - f(\underline{w}(t) \cdot \underline{x}(i))] \, x_j(i) \, f'$$

This corresponds to gradient descent on the quadratic error surface E

When f' =1, we have the perceptron learning rule (we have in general f'>0 in neural networks). The ball moves in the right direction.

Two types of network training:

Sequential mode (on-line, stochastic, or per-pattern) :
Weights updated after each pattern is presented (Perceptron is in this class)

Batch mode (off-line or per-epoch) :
Weights updated after all patterns are presented

Comparison Perceptron and Gradient Descent Rules

- Perceptron learning rule guaranteed to succeed if
  - Training examples are linearly separable
  - Sufficiently small learning rate η
- Linear unit training rule uses gradient descent guaranteed to converge to hypothesis with minimum squared error given sufficiently small learning rate η
  - Even when training data contains noise
  - Even when training data not separable by hyperplanes

Summary

Perceptron

$\underline{W}(t+1) = \underline{W}(t) + \eta(t) \, [\, d(t) - \text{sign}\,(\underline{w}(t) \,.\, \underline{x})]\, \underline{x}$

Adaline (Gradient descent method)

$\underline{W}(t+1) = \underline{W}(t) + \eta(t) \, [\, d(t) - f(\underline{w}(t) \,.\, \underline{x})]\, \underline{x}\, f'$

Multi-Layer Perceptron (MLP)

Idea: Credit assignment problem

- Problem of assigning ‘credit’ or ‘blame’ to individual elements involving in forming overall response of a learning system (hidden units)

- In neural networks, problem relates to dividing which weights should be altered, by how much and in which direction.

Example: *Three-layer* networks

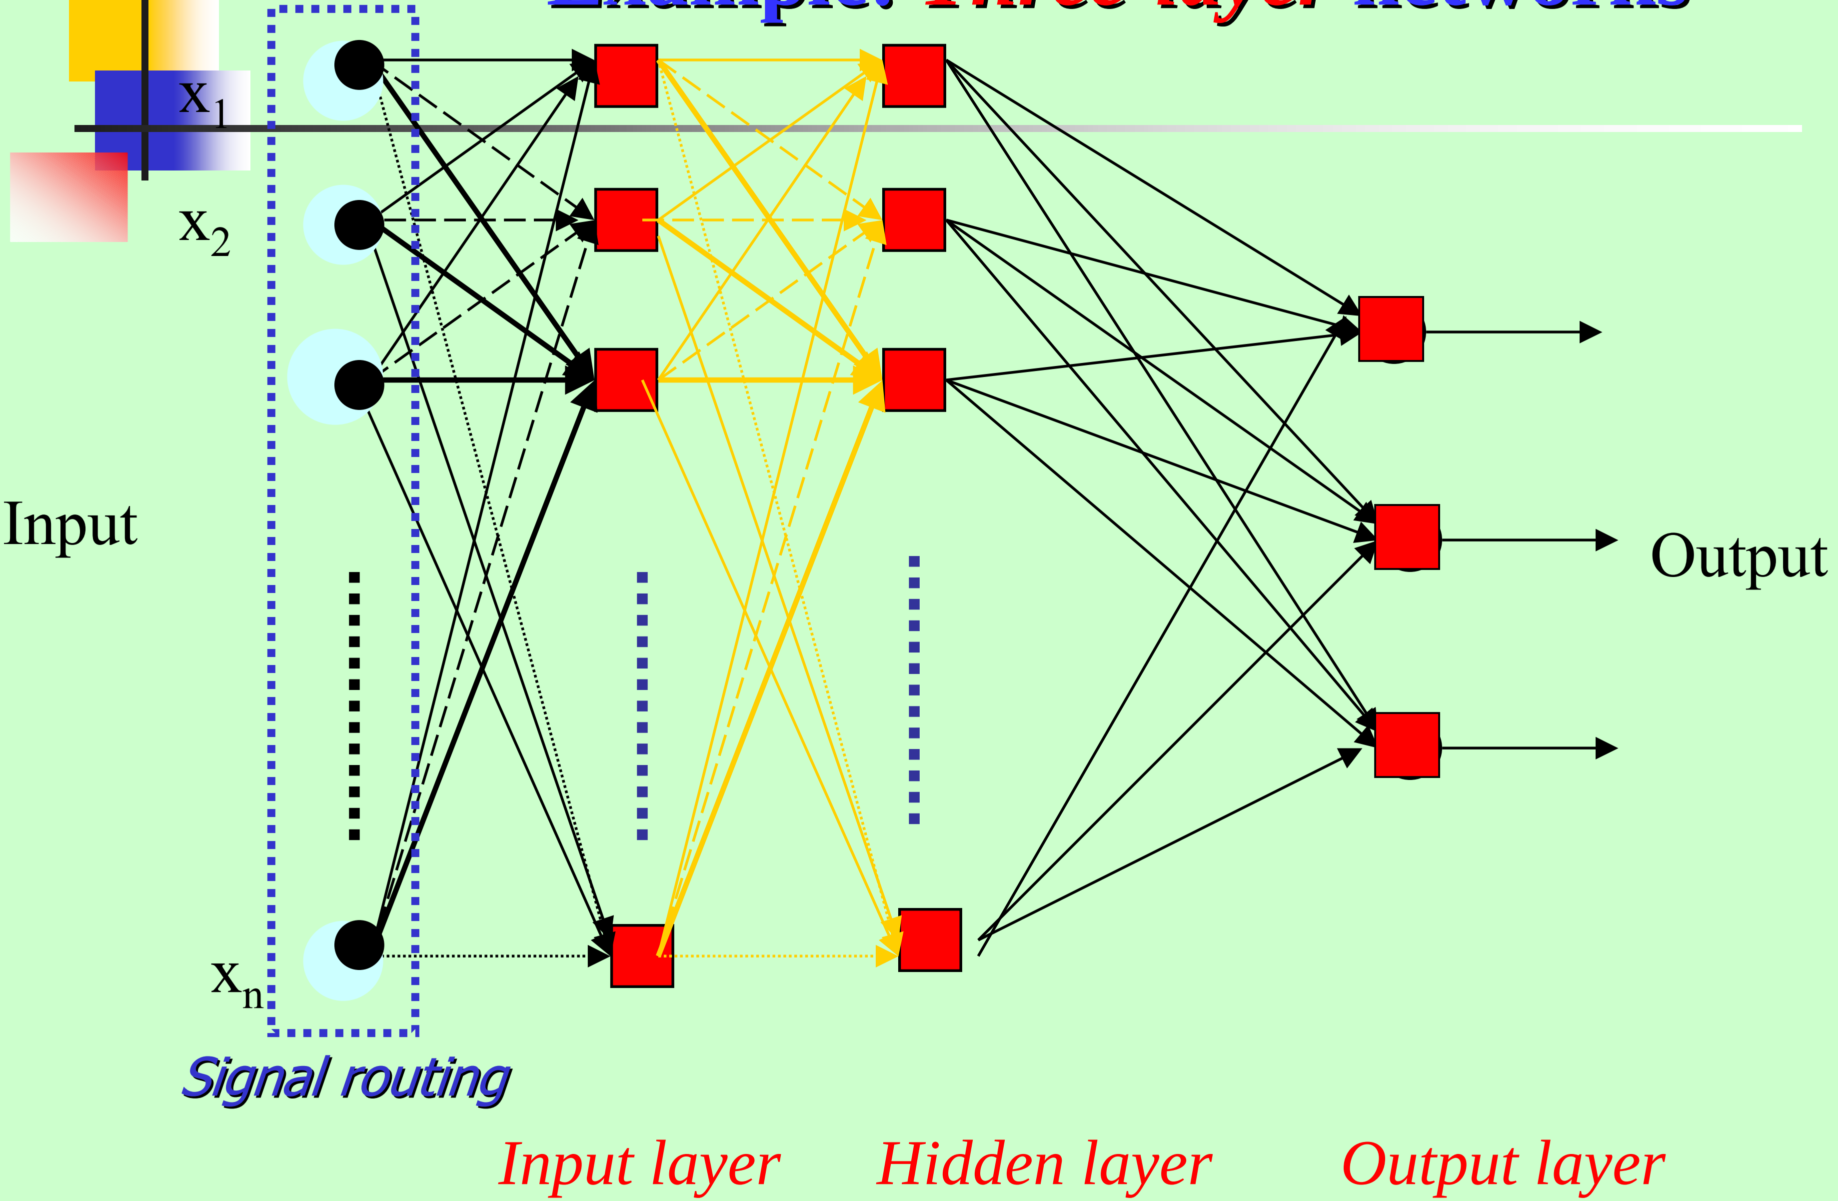

Properties of architecture

- No connections within a layer
- No direct connections between input and output layers
- Fully connected between layers
- Often more than 2 layers
- Number of output units need not equal number of input units
- Number of hidden units per layer can be more or less than input or output units

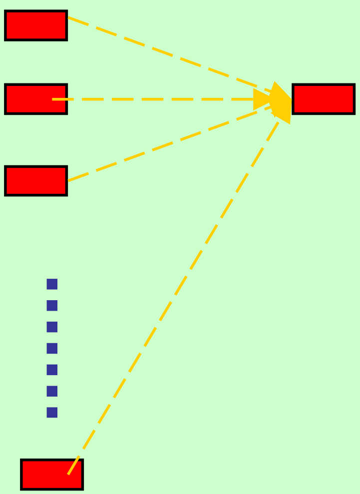

Each unit ‘‘ is a perceptron

$$y_i = f\left(\sum_{j=1}^{m} w_{ij} x_j + b_i\right)$$

BP (Back Propagation)

gradient descent method

\+

multilayer networks

MultiLayer Perceptron I

Back Propagating Learning

BP learning algorithm

Solution to "credit assignment problem" in MLP

Rumelhart, Hinton and Williams (1986)

BP has two phases:

Forward pass phase: computes 'functional signal', feedforward propagation of input pattern signals through network

Backward pass phase: computes 'error signal', propagation of error (difference between actual and desired output values) backwards through network starting at output units

BP Learning for Simplest MLP

Task : Data {I, d} to minimize

$$E = (d - o)^2 / 2$$
$$= [d - f(W(t)y(t))]^2 / 2$$
$$= [d - f(W(t)f(w(t)I))]^2 / 2$$

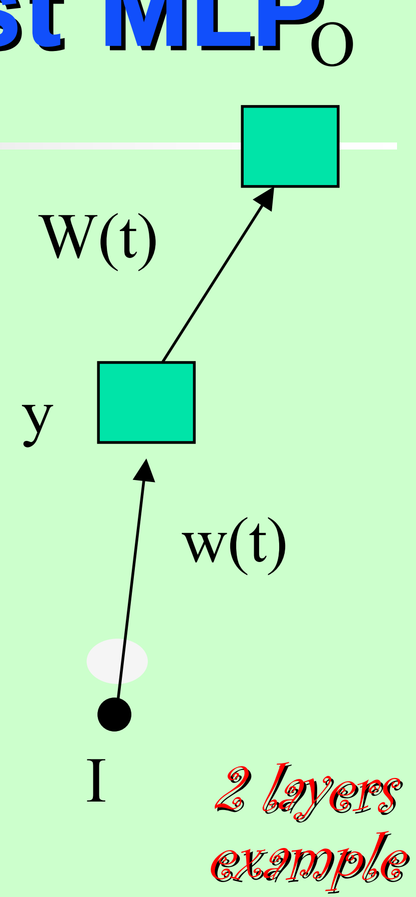

Error function at the output unit

Weight at time t is w(t) and W(t),
intend to find the weight w and W at time t+1

Where y = f(w(t)I), output of the input unit

Forward pass phase

Suppose that we have w(t), W(t) of time t

For given input I, we can calculate

$y = f(w(t)I)$

and

$o = f(W(t)\,y)$
$= f(W(t)\,f(w(t)\,I))$

Error function of output unit will be

$$E = (d - o)^2/2$$

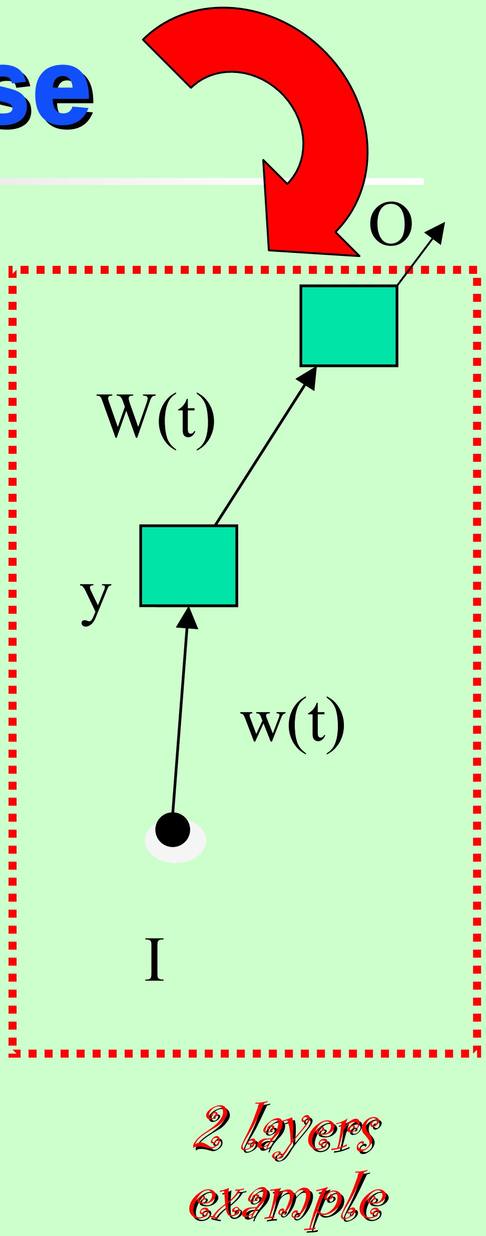

Backward Pass Phase

$$W(t+1) = W(t) - \eta \frac{dE}{dW(t)}$$
$$= W(t) - \eta \frac{dE}{df} \frac{df}{dW(t)}$$
$$= W(t) + \eta (d - o) f'(W(t) y) y$$

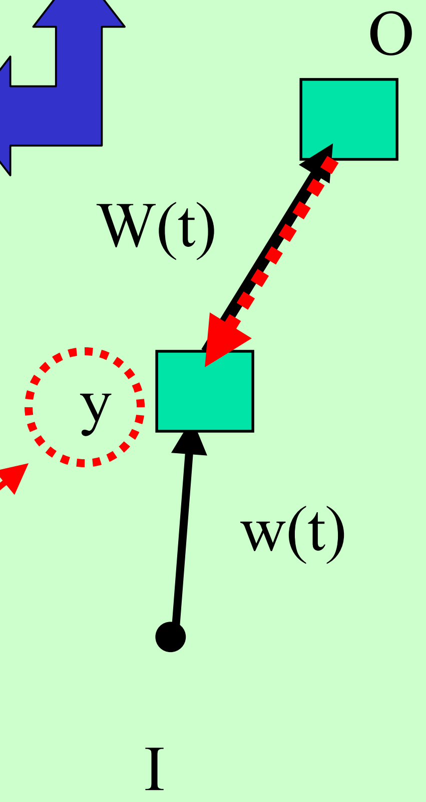

E = *(d - o)² /2*   o = f (W(t) y)

Backward pass phase

$$W(t+1) = W(t) - \eta \frac{dE}{dW(t)}$$

$$= W(t) - \eta \frac{dE}{df} \frac{df}{dW(t)}$$

$$= W(t) + \eta (d - o) f'(W(t) y) y$$

$$= W(t) + \eta \Delta y$$

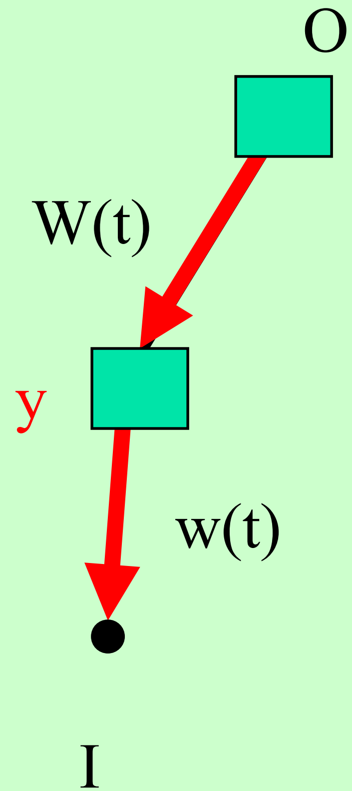

where $\Delta = (d\text{-}o)\, f'$

Backward pass phase

$$w(t+1) = w(t) - \eta \frac{dE}{dw(t)}$$

$$= w(t) - \eta \frac{dE}{dy} \frac{dy}{dw(t)}$$

$$= w(t) + \eta (d - o) f'(W(t) y) W(t) \frac{dy}{dw(t)}$$

$$= w(t) + \eta \Delta W(t) f'(w(t) I) I$$

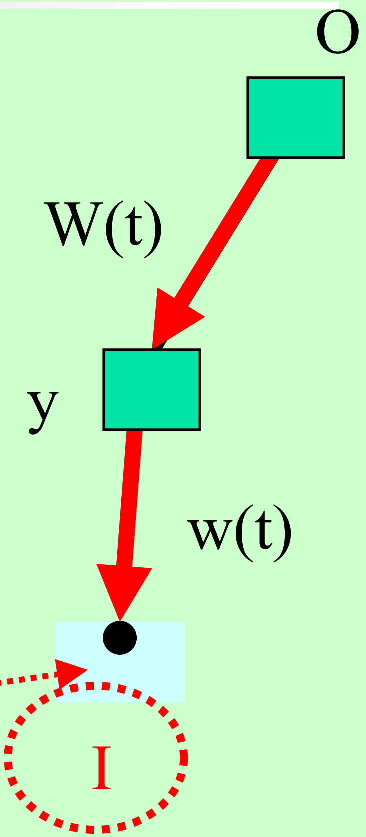

o = f (W(t) y)

= f (W(t) f(w(t) I))

Summary

weight updates are local

$$w_{ji}(t+1) - w_{ji}(t) = \eta\delta_j(t) I_i(t)$$ (input unit)

$$W_{kj}(t+1) - W_{kj}(t) = \eta\Delta_k(t) y_j(t)$$ (output unit)

output unit

$$W_{kj}(t+1) - W_{kj}(t) = \eta\Delta_k(t) y_j(t)$$
$$= \eta(d_k(t) - O_k(t)) f'(Net_k(t)) y_j(t)$$

input unit

$$w_{ji}(t+1) - w_{ji}(t) = \eta\delta_j(t) I_i(t)$$
$$= \eta f'(net_j(t)) \sum_k \Delta_k(t) W_{kj} I_i(t)$$

Once weight changes are computed for all units, weights are updated at same time (bias included as weights here)

We now compute the derivative of the activation function $f(\,)$.

Activation Functions

- to compute δ_j and Δ_k we need to find the derivative of activation function *f*
- to find derivative the activation function must be smooth

Sigmoidal (logistic) function-common in MLP

$$f\ (net_i(t)) = \frac{1}{1+\exp(-k\,net_i(t))}$$

where k is a positive constant. The sigmoidal function gives value in range of 0 to 1

Input-output function of a neuron (rate coding assumption)

Shape of sigmoidal function

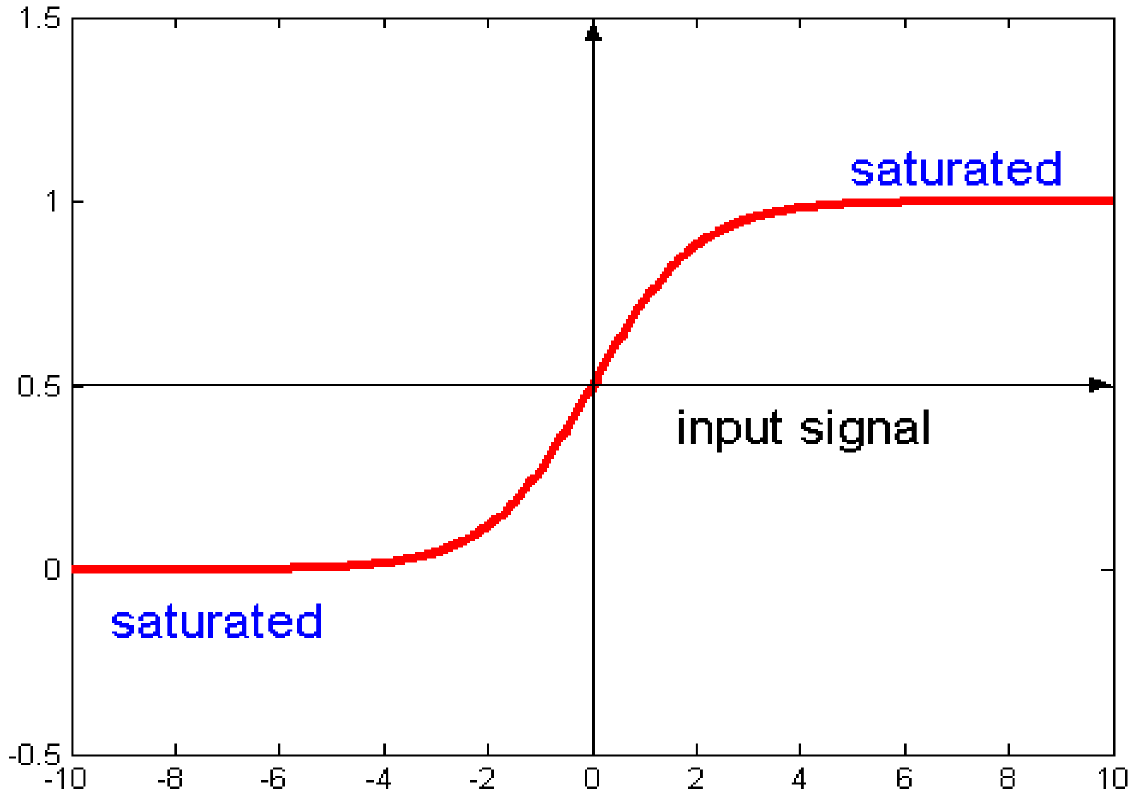

Note: when net = 0, f = 0.5

Shape of sigmoidal function derivative

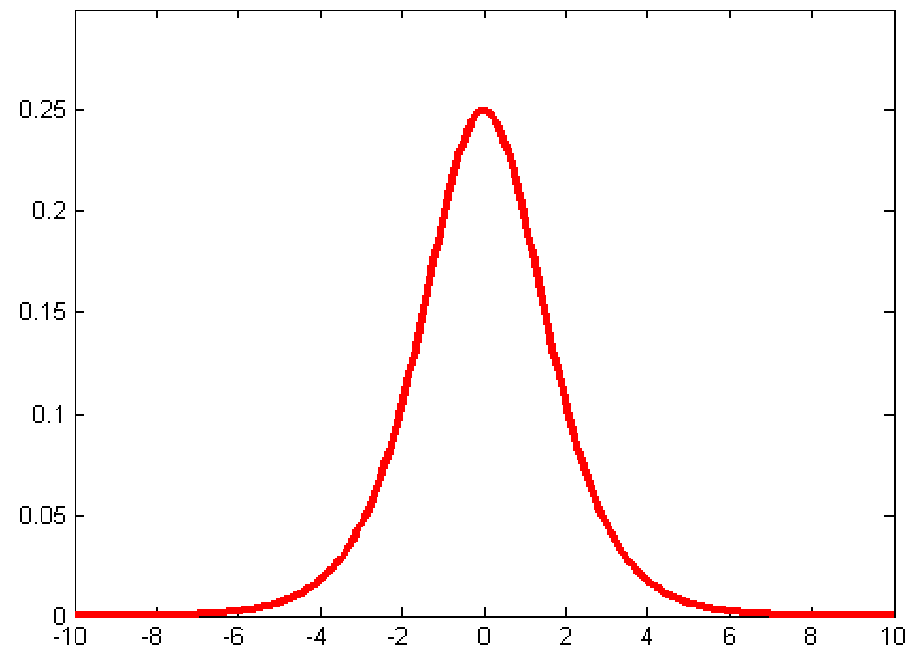

Derivative of sigmoidal function has max at x= 0, is symmetric about this point falling to zero as sigmoidal approaches extreme values

Returning to **local error gradients** in BP algorithm we have for output units

$$\Delta_i(t) = (d_i(t) - O_i(t))\, f'(Net_i(t))$$
$$= (d_i(t) - O_i(t))\, kO_i(t)(1 - O_i(t))$$

For input units we have

$$\delta_i(t) = f'(net_i(t)) \sum_k \Delta_k(t) W_{ki}$$
$$= ky_i(t)(1 - y_i(t)) \sum_k \Delta_k(t) W_{ki}$$

Since degree of weight change is proportional to derivative of activation function, weight changes will be greatest when units receives mid-range functional signal than at extremes

Summary of BP learning algorithm

Set learning rate η

Set initial weight values (incl.. biases): w, W

Loop until stopping criteria satisfied:

present input pattern to NN inputs
compute functional signal for input units
compute functional signal for output units

present Target response to output units
compute error signal for output units
compute error signal for input units
update all weights at same time
increment n to n+1 and select next I and d

end loop

Network training:

- Training set shown repeatedly until stopping criteria are met
- Each full presentation of all patterns = ‘epoch’
- Randomise order of training patterns presented for each epoch in order to avoid correlation between consecutive training pairs being learnt (order effects)

Two types of network training:

- **Sequential mode** (on-line, stochastic, or per-pattern) Weights updated after each pattern is presented

- **Batch mode** (off-line or per -epoch)

Advantages and disadvantages of different modes

Sequential mode:

- Less storage for each weighted connection
- Random order of presentation and updating per pattern means search of weight space is stochastic-reducing risk of local minima able to take advantage of any redundancy in training set (*i.e.* same pattern occurs more than once in training set, esp. for large training sets)
- Simpler to implement

Batch mode:

- Faster learning than sequential mode

MultiLayer Perceptron II

Dynamics of MultiLayer Perceptron

Summary of Network Training

Forward phase: $\underline{I}(t), \underline{w}(t), \underline{net}(t), \underline{y}(t), \underline{W}(t), \underline{Net}(t), \underline{O}(t)$

Backward phase:

Output unit

$$
\begin{aligned}
W_{kj}(t+1) - W_{kj}(t) &= \eta \Delta_k(t) y_j(t) \\
&= \eta (d_k(t) - O_k(t)) f'(Net_k(t)) y_j(t)
\end{aligned}
$$

Input unit

$$
\begin{aligned}
w_{ji}(t+1) - w_{ij}(t) &= \eta \delta_j(t) I_i(t) \\
&= \eta f'(net_j(t)) \sum_k \Delta_k(t) W_{kj}(t) I_i(t)
\end{aligned}
$$

Network training:

Training set shown repeatedly until stopping criteria are met.

Possible convergence criteria are

- Euclidean norm of the gradient vector reaches a sufficiently small denoted as θ.
- When the absolute rate of change in the average squared error per epoch is sufficiently small denoted as θ.
- Validation for generalization performance : stop when generalization reaching the peak (illustrate in this lecture)

Goals of Neural Network Training

To give the correct output for input training vector (Learning)

To give good responses to new unseen input patterns (Generalization)

Training and Testing Problems

- **Stuck neurons**: Degree of weight change is proportional to derivative of activation function, weight changes will be greatest when units receives mid-range functional signal than at extremes neuron. To avoid stuck neurons weights initialization should give outputs of all neurons approximate 0.5

- **Insufficient number of training patterns**: In this case, the training patterns will be learnt instead of the underlying relationship between inputs and output, i.e. network just memorizing the patterns.

- **Too few hidden neurons**: network will not produce a good model of the problem.

- **Over-fitting**: the training patterns will be learnt instead of the underlying function between inputs and output because of too many of hidden neurons. This means that the network will have a poor generalization capability.

Dynamics of BP learning

Aim is to minimise an error function over all training patterns by adapting weights in MLP

Recalling the typical error function is the mean squared error as follows

E(t)= $$\frac{1}{2}\sum_{k=1}^{p}(d_k(t)-O_k(t))^2$$

The idea is to reduce E(t) to global minimum point.

Dynamics of BP learning

In single layer perceptron with linear activation functions, the error function is simple, described by a smooth parabolic surface with a single minimum

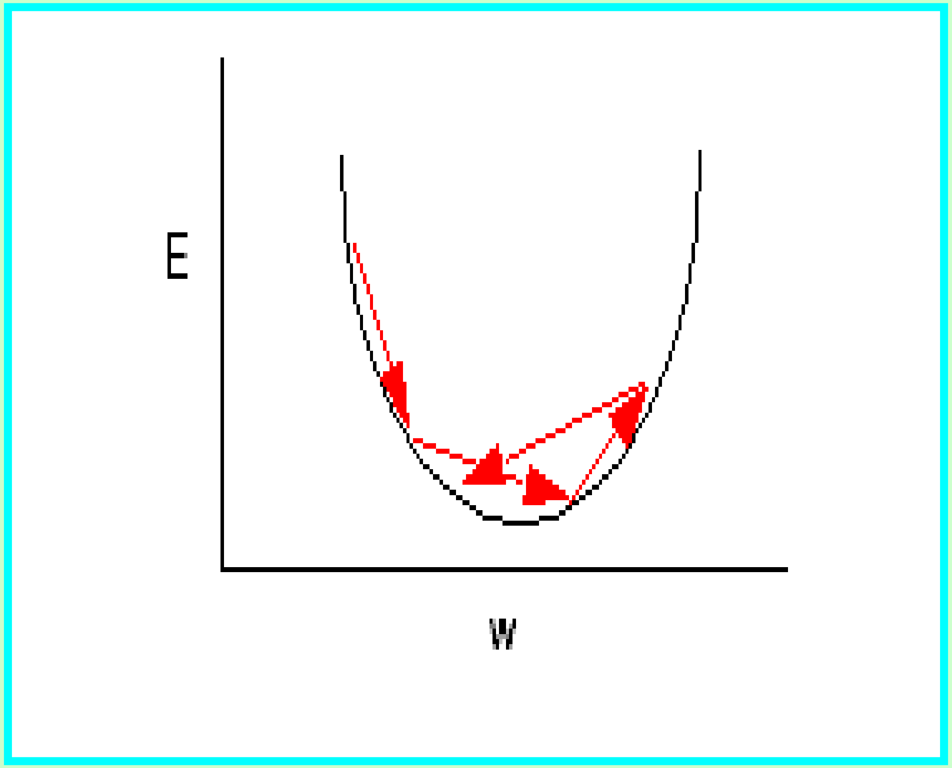

Dynamics of BP learning

- MLP with non-linear activation functions have complex error surfaces (e.g. plateaus, long valleys etc.) with no single minimum

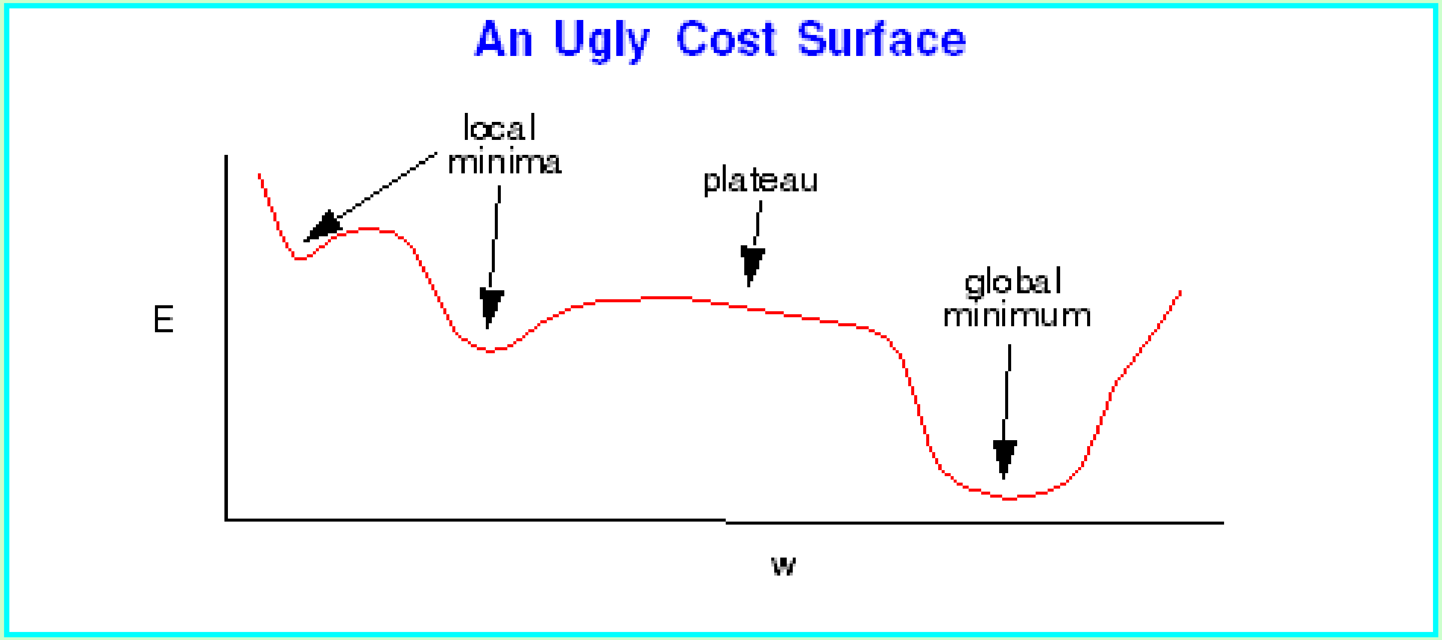

For complex error surfaces the problem is learning rate must keep small to prevent divergence. Adding momentum term is a simple approach dealing with this problem.

Momentum

- Reducing problems of instability while increasing the rate of convergence
- Adding term to weight update equation can effectively holds as exponentially weight history of previous weights changed

Modified weight update equation is

$$w_{ij}(n+1) - w_{ij}(n) = \eta\delta_j(n)y_i(n) + \\ + \alpha[w_{ij}(n) - w_{ij}(n-1)]$$

Effect of momentum term

- If weight changes tend to have same sign, momentum term increases and gradient decrease speed up convergence on shallow gradient
- If weight changes tend have opposing signs, momentum term decreases and gradient descent slows to reduce oscillations (stabilizes)
- Can help escape being trapped in local minima

Selecting Initial Weight Values

- Choice of initial weight values is important as this decides starting position in weight space. That is, how far away from global minimum
- Aim is to select weight values which produce midrange function signals
- Select weight values randomly from uniform probability distribution
- Normalise weight values so number of weighted connections per unit produces midrange function signal

Convergence of Backprop

Avoid local minumum with fast convergence:

- Add momentum
- Stochastic gradient descent
- Train multiple nets with different initial weights

Nature of convergence

- Initialize weights 'near zero' or initial networks near-linear
- Increasingly non-linear functions possible as training progresses

Use of Available Data Set for Training

The available data set is normally split into three sets as follows:

- Training set – use to update the weights. Patterns in this set are repeatedly in random order. The weight update equation are applied after a certain number of patterns.
- Validation set – use to decide when to stop training only by monitoring the error.
- Test set – Use to test the performance of the neural network. It should not be used as part of the neural network development cycle.

Earlier Stopping - Good Generalization

- Running too many epochs may overtrain the network and result in overfitting and perform poorly in generalization.
- Keep a hold-out validation set and test accuracy after every epoch. Maintain weights for best performing network on the validation set and stop training when error increases increases beyond this.

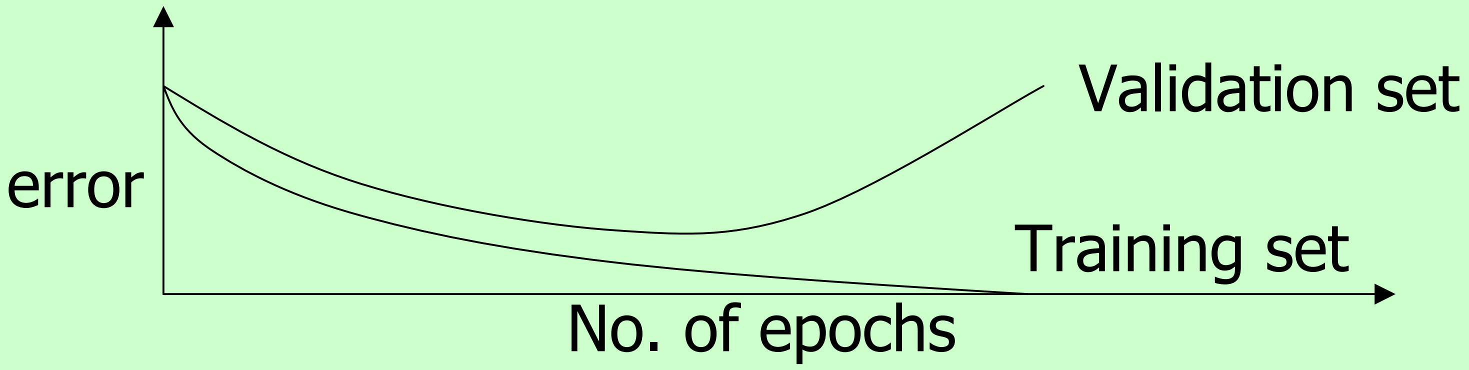

Model Selection by Cross-validation

- Too few hidden units prevent the network from learning adequately fitting the data and learning the concept (more than two layer networks).
- Too many hidden units leads to overfitting.
- Similar cross-validation methods can be used to determine an appropriate number of hidden units by using the optimal test error to select the model with optimal number of hidden layers and nodes.

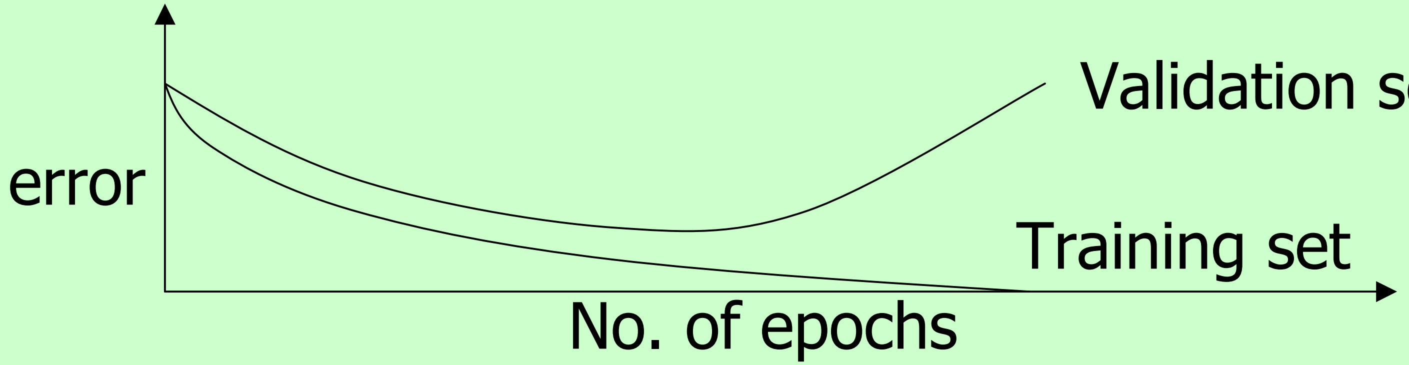

Radial Basis Functions

Radial Basis Functions

Overview

Radial-basis function (RBF) networks

RBF = radial-basis function: a function which depends only on the radial distance from a point

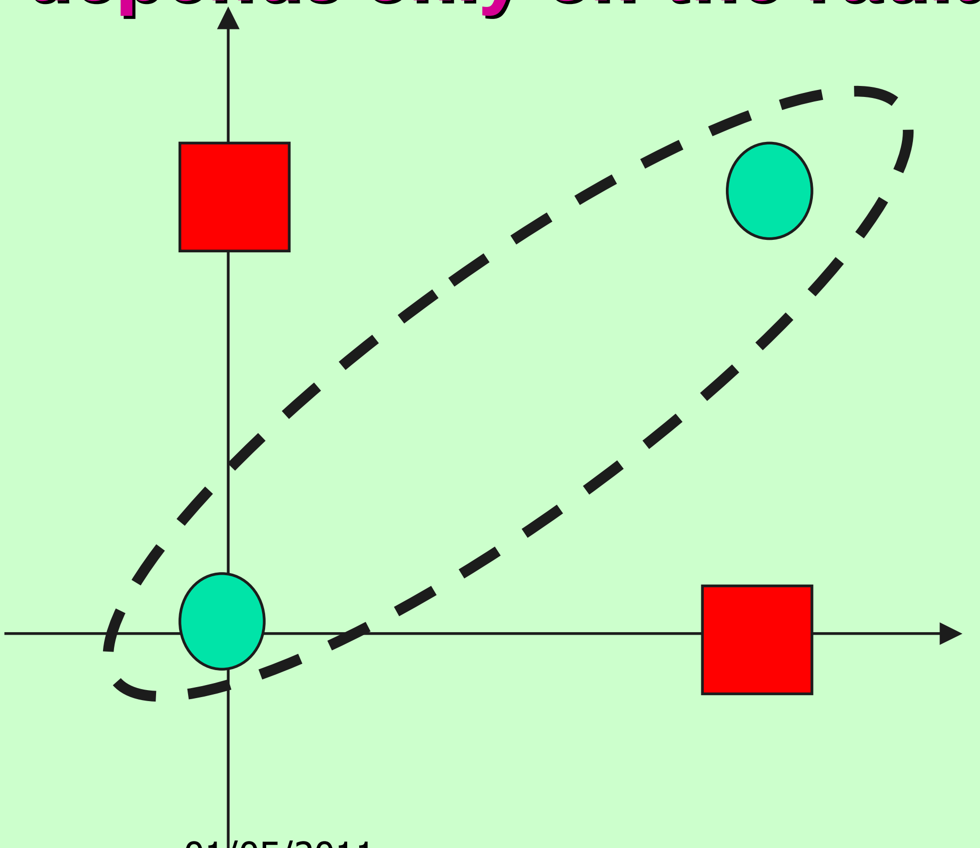

XOR problem

quadratically separable

Radial-basis function (RBF) networks

So RBFs are functions taking the form

$$\phi(\| \underline{x} - \underline{x}_i \|)$$

where ϕ is a non-linear activation function, $\underline{x}$ is the input and $\underline{x}_i$ is the *i'th* position, prototype, *basis* or *centre* vector.

The idea is that points near the centres will have similar outputs (i.e. if $\underline{x} \sim \underline{x}i$ then $f(\underline{x}) \sim f(\underline{x}i)$) since they should have similar properties.

The simplest is the linear RBF : $\phi(x) = \| \underline{x} - \underline{x}_i \|$

Typical RBFs include

(a) Multiquadrics

$$\phi(r) = (r^2 + c^2)^{1/2}$$

for some c>0

(b) Inverse multiquadrics

$$\phi(r) = (r^2 + c^2)^{-1/2}$$

for some c>0

(c) Gaussian

$$\phi(r) = \exp(-\frac{r^2}{2\sigma^2})$$

for some σ >0

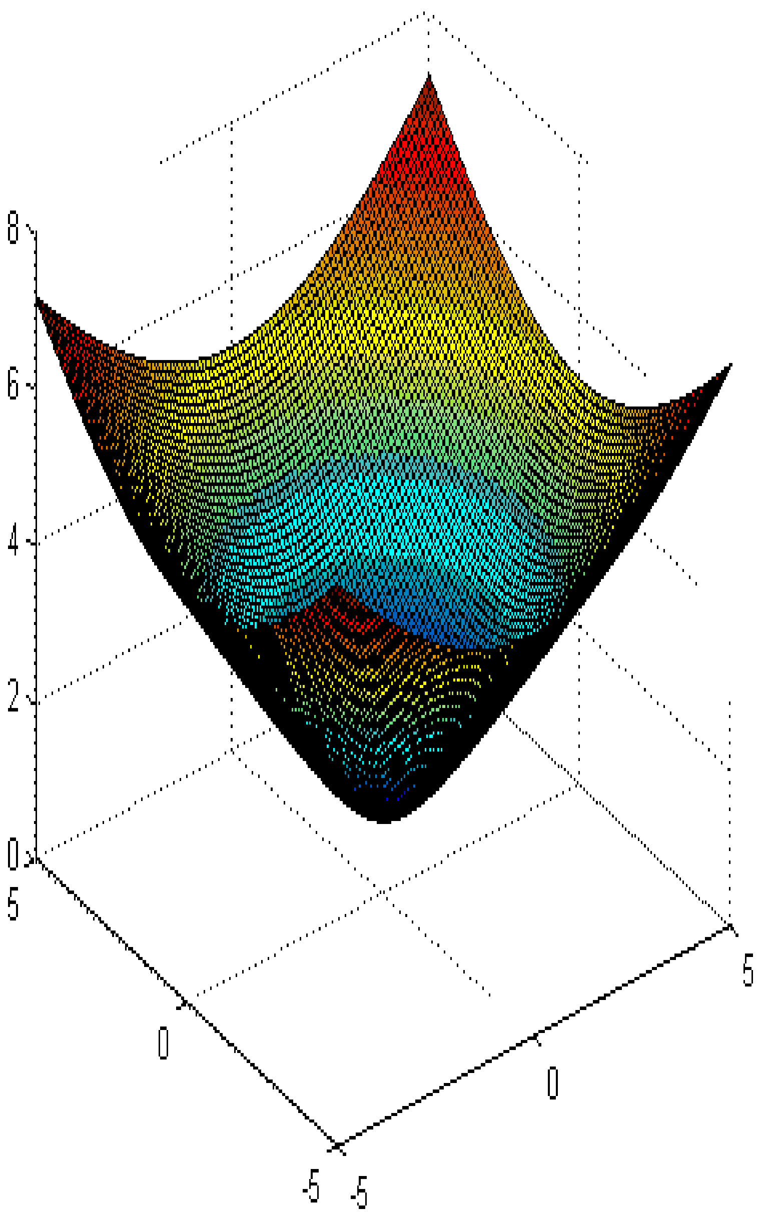

‘nonlocalized’ functions

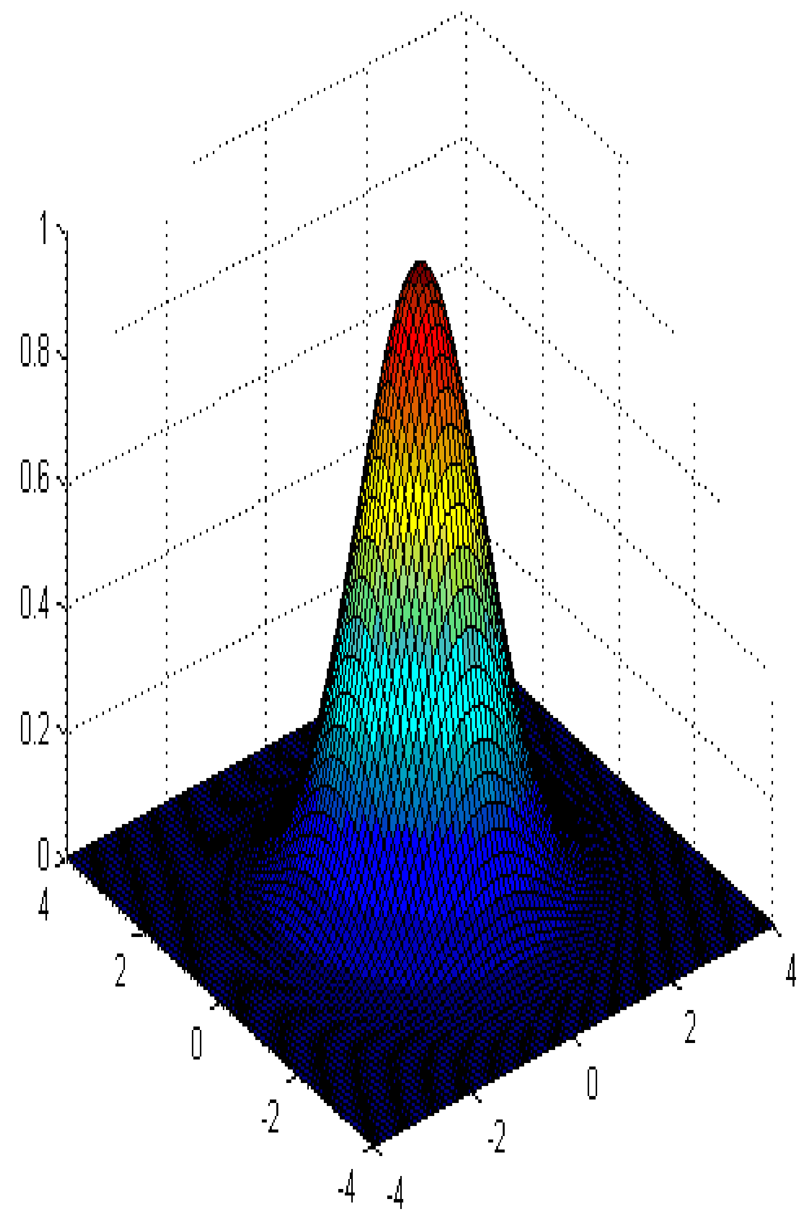

‘localized’ functions

➢ Idea is to use a weighted sum of the outputs from the basis functions to represent the data.

➢ Thus centers can be thought of as prototypes of input data.

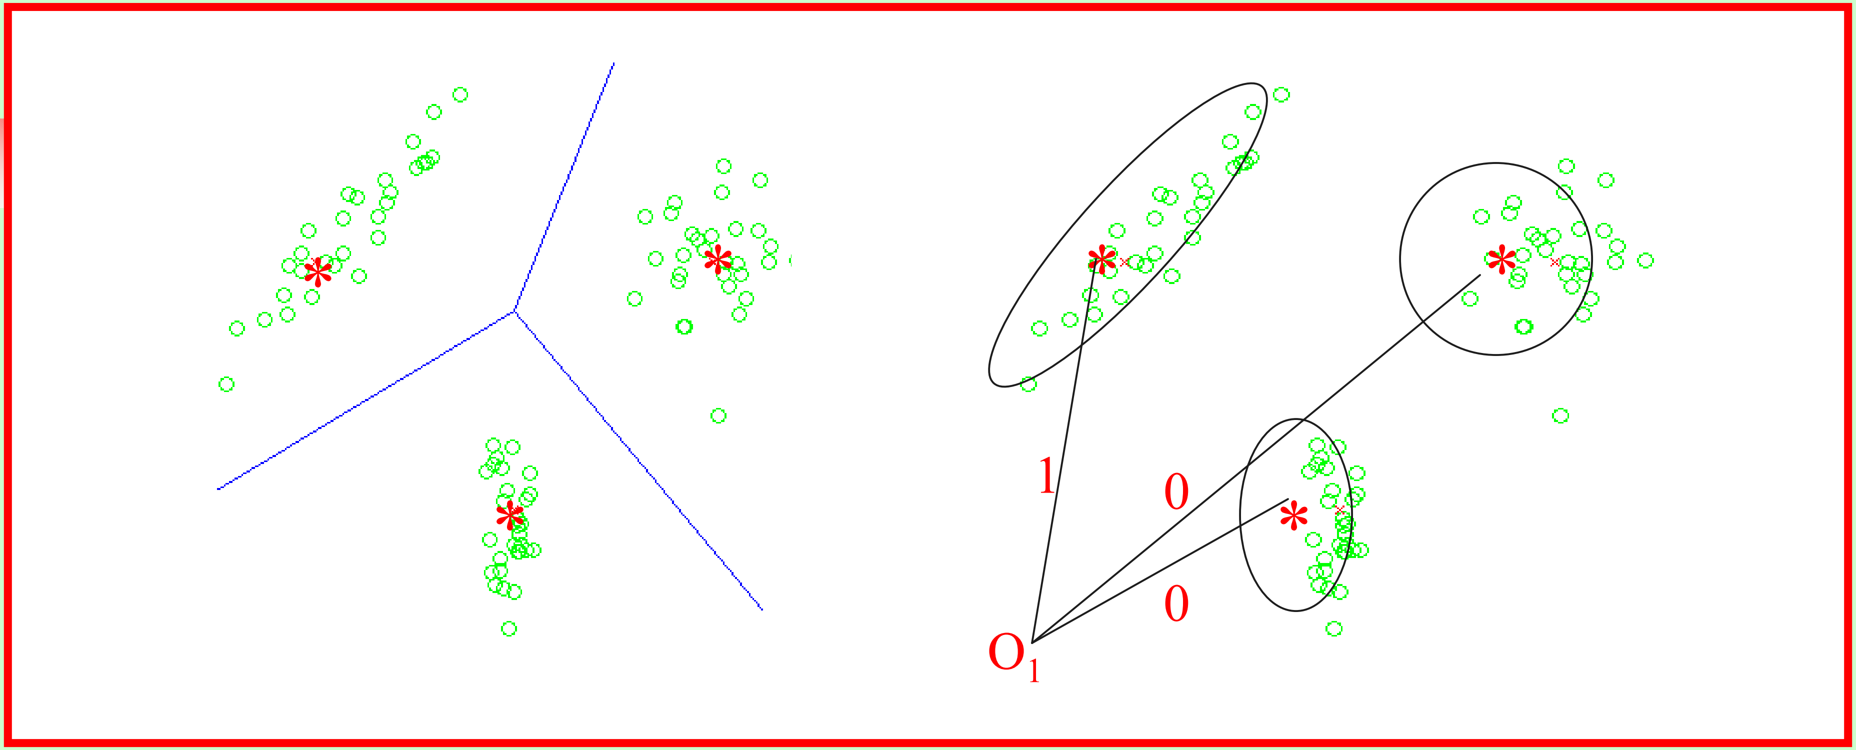

MLP distributed vs RBF local

Starting point: exact interpolation

Each input pattern x must be mapped onto a target value d

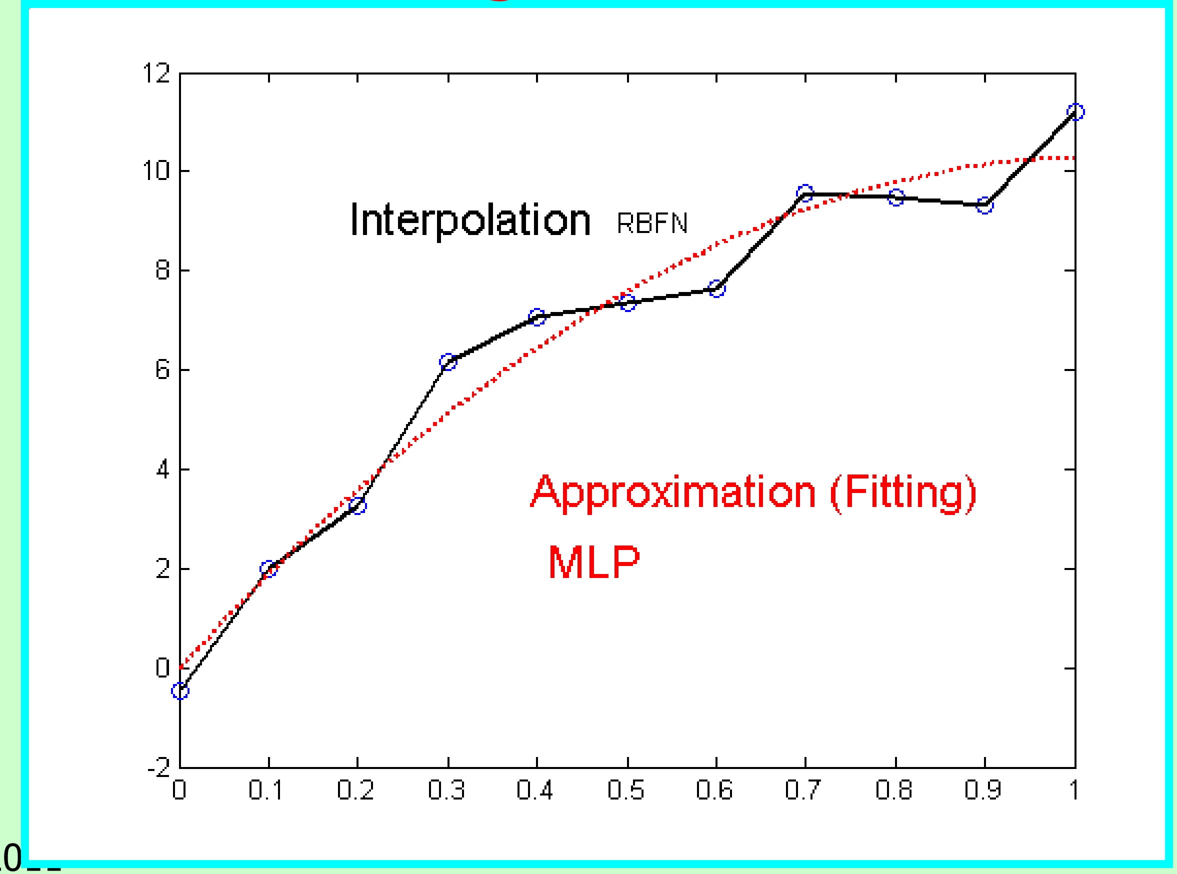

That is, given a set of N vectors $\underline{x}_i$ and a corresponding set of N real numbers, d_i (the targets), find a function F that satisfies the interpolation condition:

$$F(\underline{x}_i) = d_i \quad \text{for } i = 1,\ldots,N$$

or more exactly find:

$$F(\underline{x}) = \sum_{j=1}^{N} w_j \phi(\|\underline{x} - \underline{x}_j\|)$$

satisfying:

$$F(\underline{x}_i) = \sum_{j=1}^{N} w_j \phi(\|\underline{x}_i - \underline{x}_j\|) = d_i$$

Single-layer networks

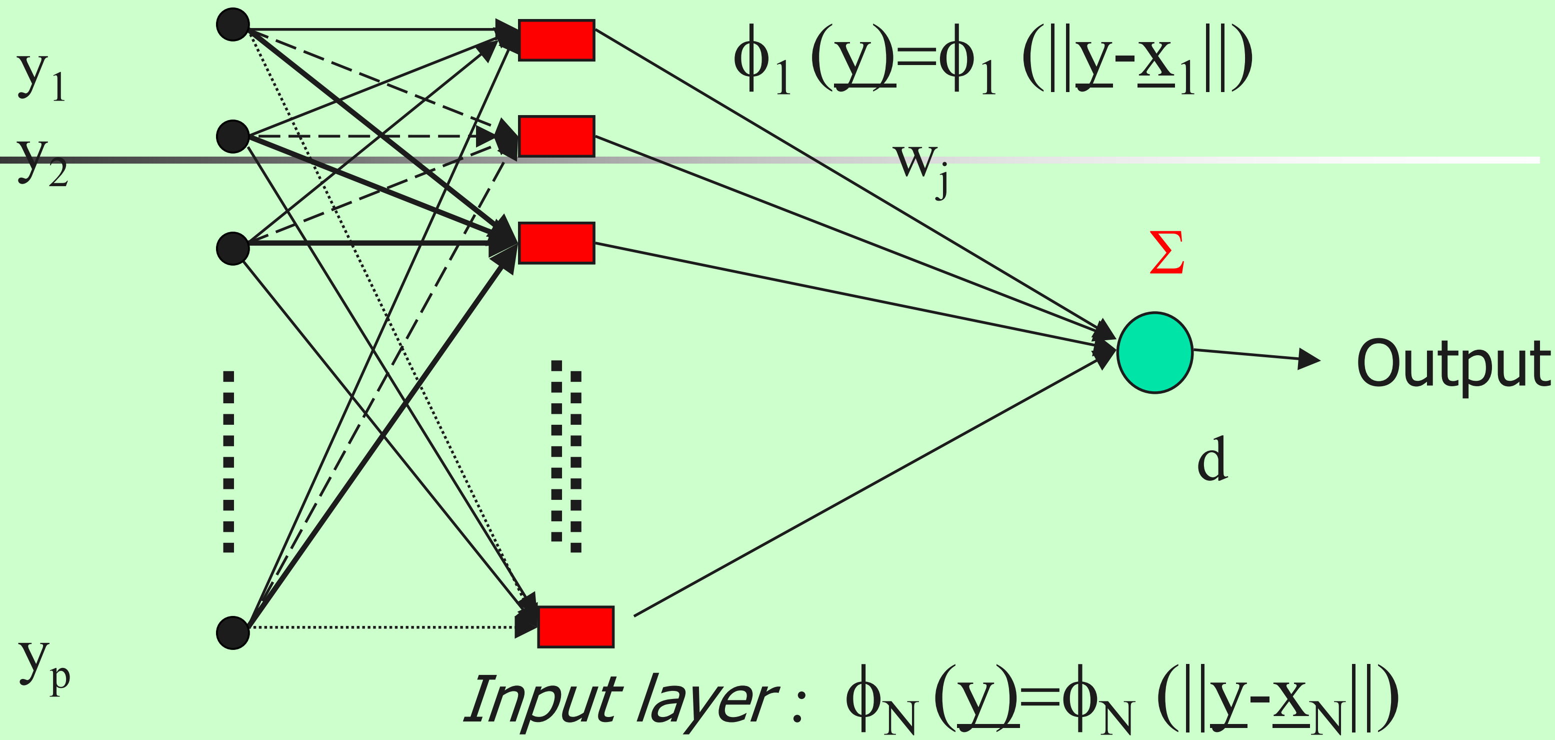

- output = $\Sigma\, w_i\, \phi_i (\underline{y} - \underline{x}_i)$
- adjustable parameters are weights w_j
- number of input units $\leq$ number of data points
- Form of the basis functions decided in advance

To summarize:

- For a given data set containing N points $(\underline{x}_i, d_i)$, i=1,…,N
- Choose a RBF function ϕ
- Calculate $\phi(\underline{x}_j - \underline{x}_i)$
- Solve the linear equation $\Phi\ \underline{W} = \underline{D}$
- Get the unique solution
- Done

- Like MLP's, RBFNs can be shown to be able to approximate any function to arbitrary accuracy (using an arbitrarily large numbers of basis functions).
- Unlike MLP's, however, they have the property of 'best approximation' i.e. there exists an RBFN with minimum approximation error.

Large σ = 1

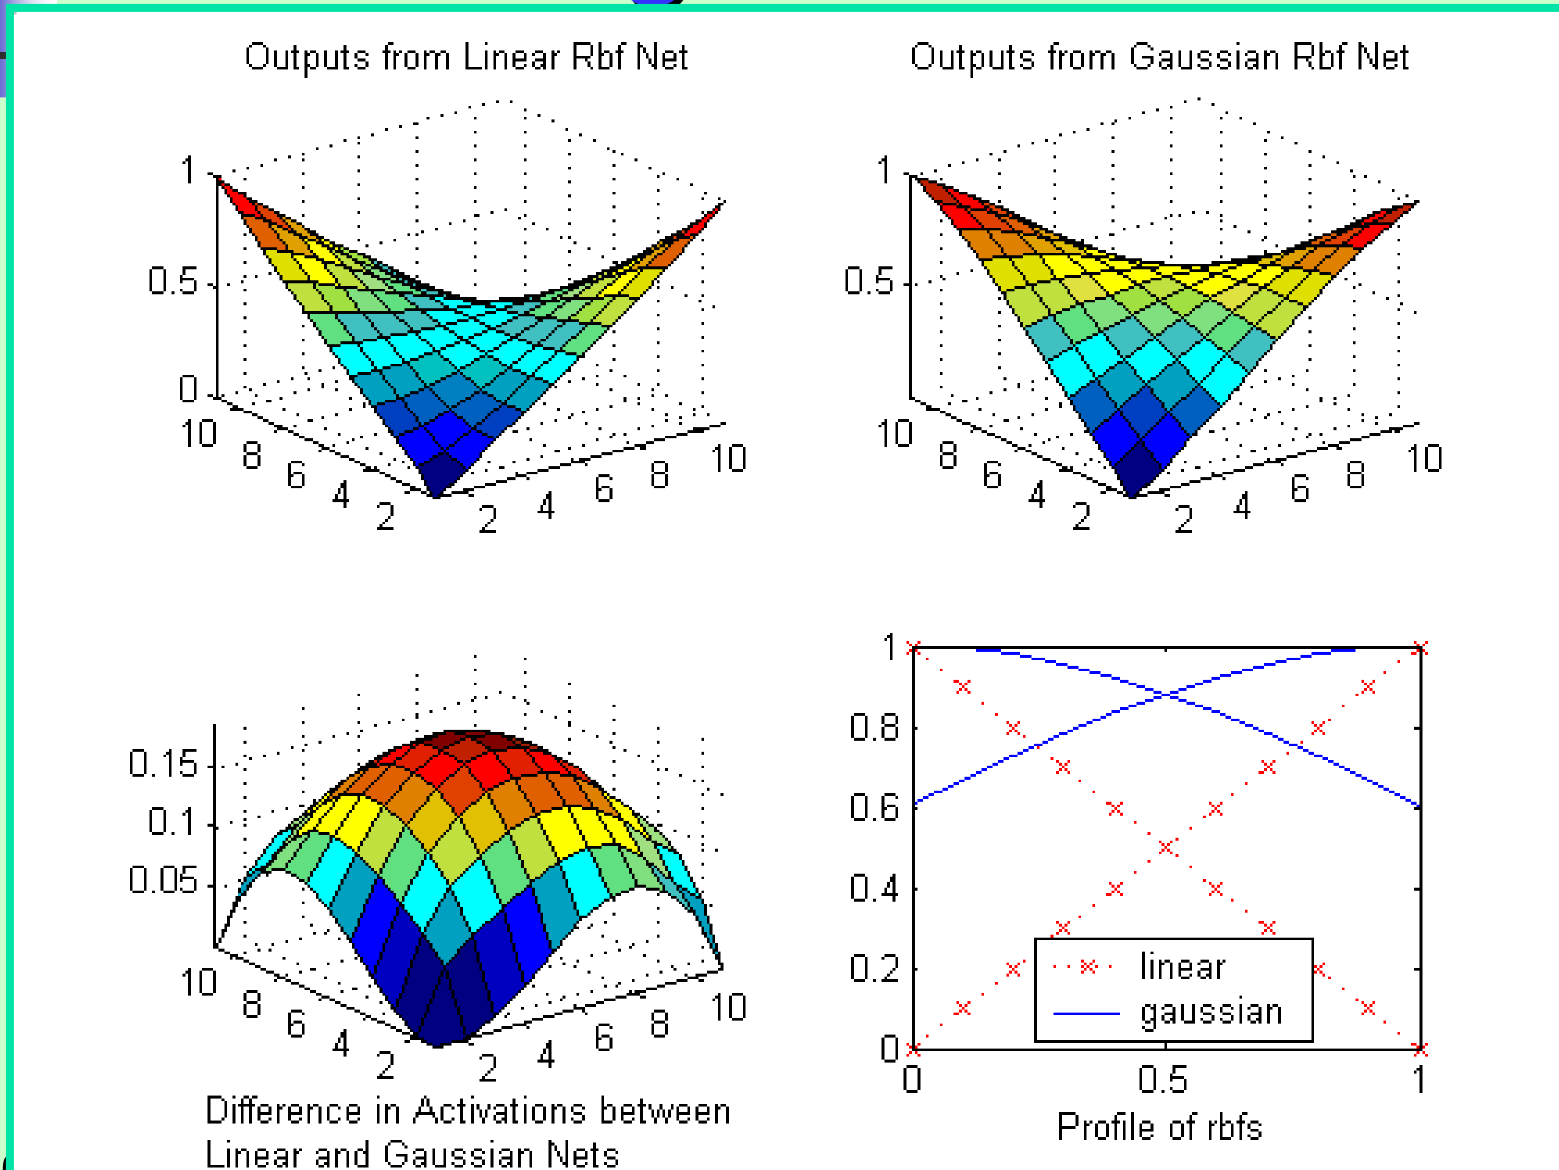
Outputs from Linear Rbf Net
Outputs from Gaussian Rbf Net
Difference in Activations between Linear and Gaussian Nets
linear
gaussian
Profile of rbfs

Small σ = 0.2

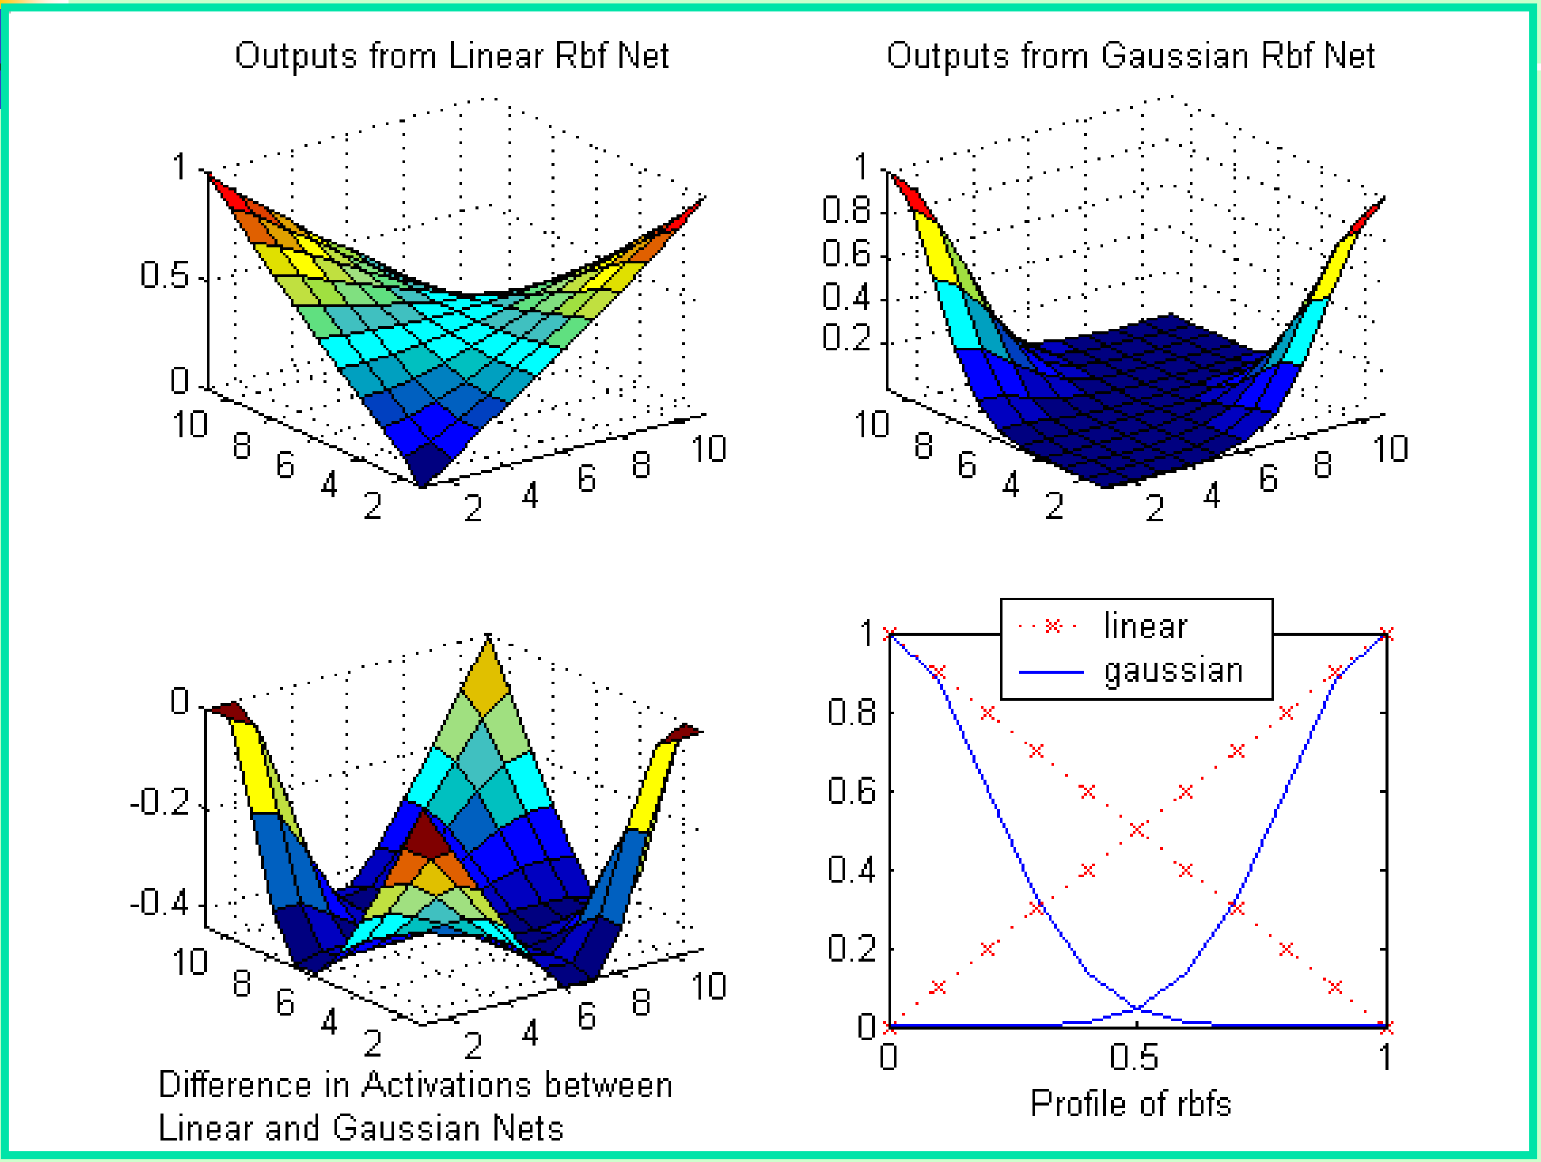
Outputs from Linear Rbf Net
Outputs from Gaussian Rbf Net
Difference in Activations between
Linear and Gaussian Nets
linear
gaussian
Profile of rbfs

Problems with exact interpolation

can produce poor generalisation performance as only data points constrain mapping

Overfitting problem

Bishop(1995) example

Underlying function $f(x)=0.5+0.4\text{sine}(2\pi x)$
sampled randomly for 30 points

added Gaussian noise to each data point

30 data points 30 hidden RBF units

fits all data points but creates oscillations due added noise and unconstrained between data points

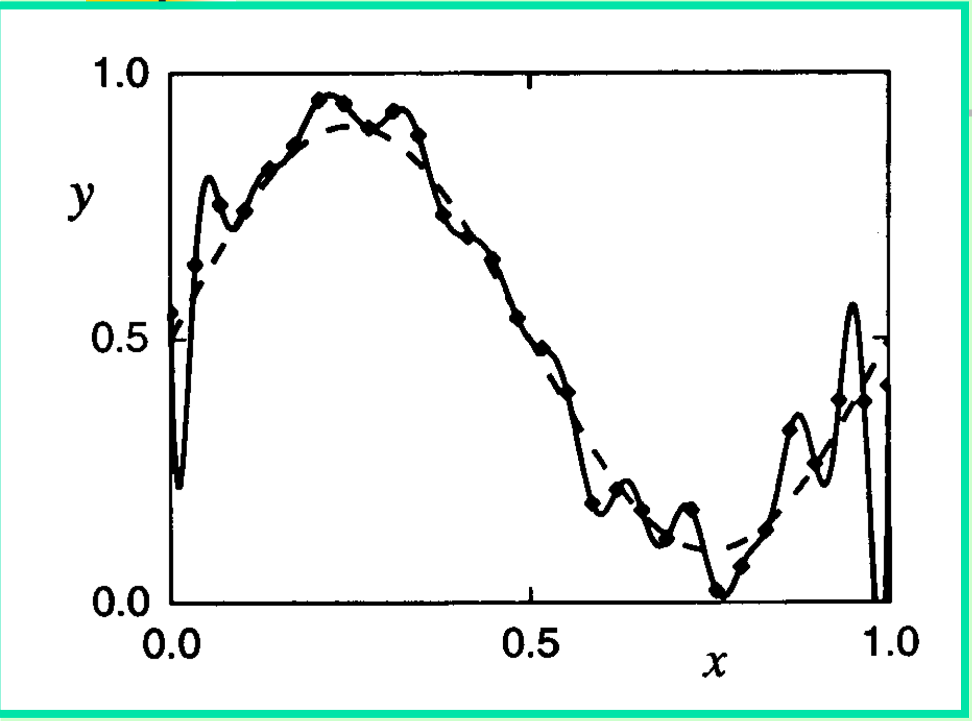

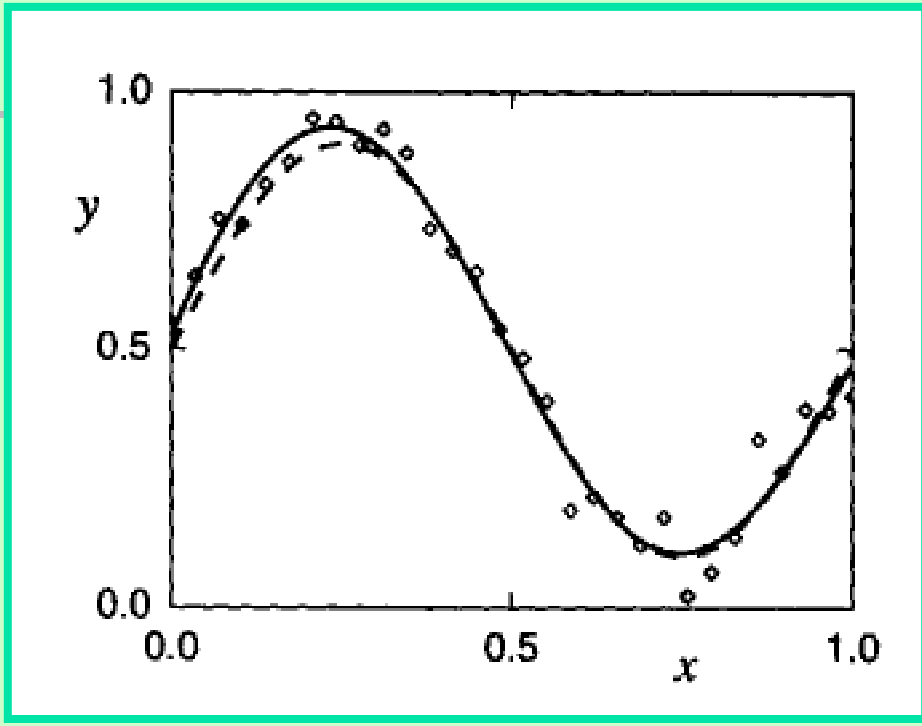

All Data Points

5 Basis functions

To fit an RBF to every data point is very inefficient due to the computational cost of matrix inversion and is very bad for generalization so:

- ✓ Use less RBF's than data points I.e. M<N
- ✓ Therefore don't necessarily have RBFs centred at data points
- ✓ Can include bias terms
- ✓ Can have Gaussian with general covariance matrices but there is a trade-off between complexity and the number of parameters to be found eg for *d* rbfs we have:

Fuzzy Modelling and Identification

Fuzzy Clustering with Application to Data-Driven Modelling

Introduction

- The ability to cluster data (concepts, perceptions, etc.)
  - essential feature of human intelligence.
- A cluster is a set of objects that are more similar to each other than to objects from other clusters.
- Applications of clustering techniques in pattern recognition and image processing.
- Some machine-learning techniques are based on the notion of similarity (decision trees, case-based reasoning)
- Non-linear regression and black-box modelling can be based on the partitioning data into clusters.

Section Outline

- **Basic concepts in clustering**
  - data set
  - partition matrix
  - distance measures
- **Clustering algorithms**
  - fuzzy c-means
  - Gustafson–Kessel
- **Application examples**
  - system identification and modelling
  - diagnosis

Examples of Clusters

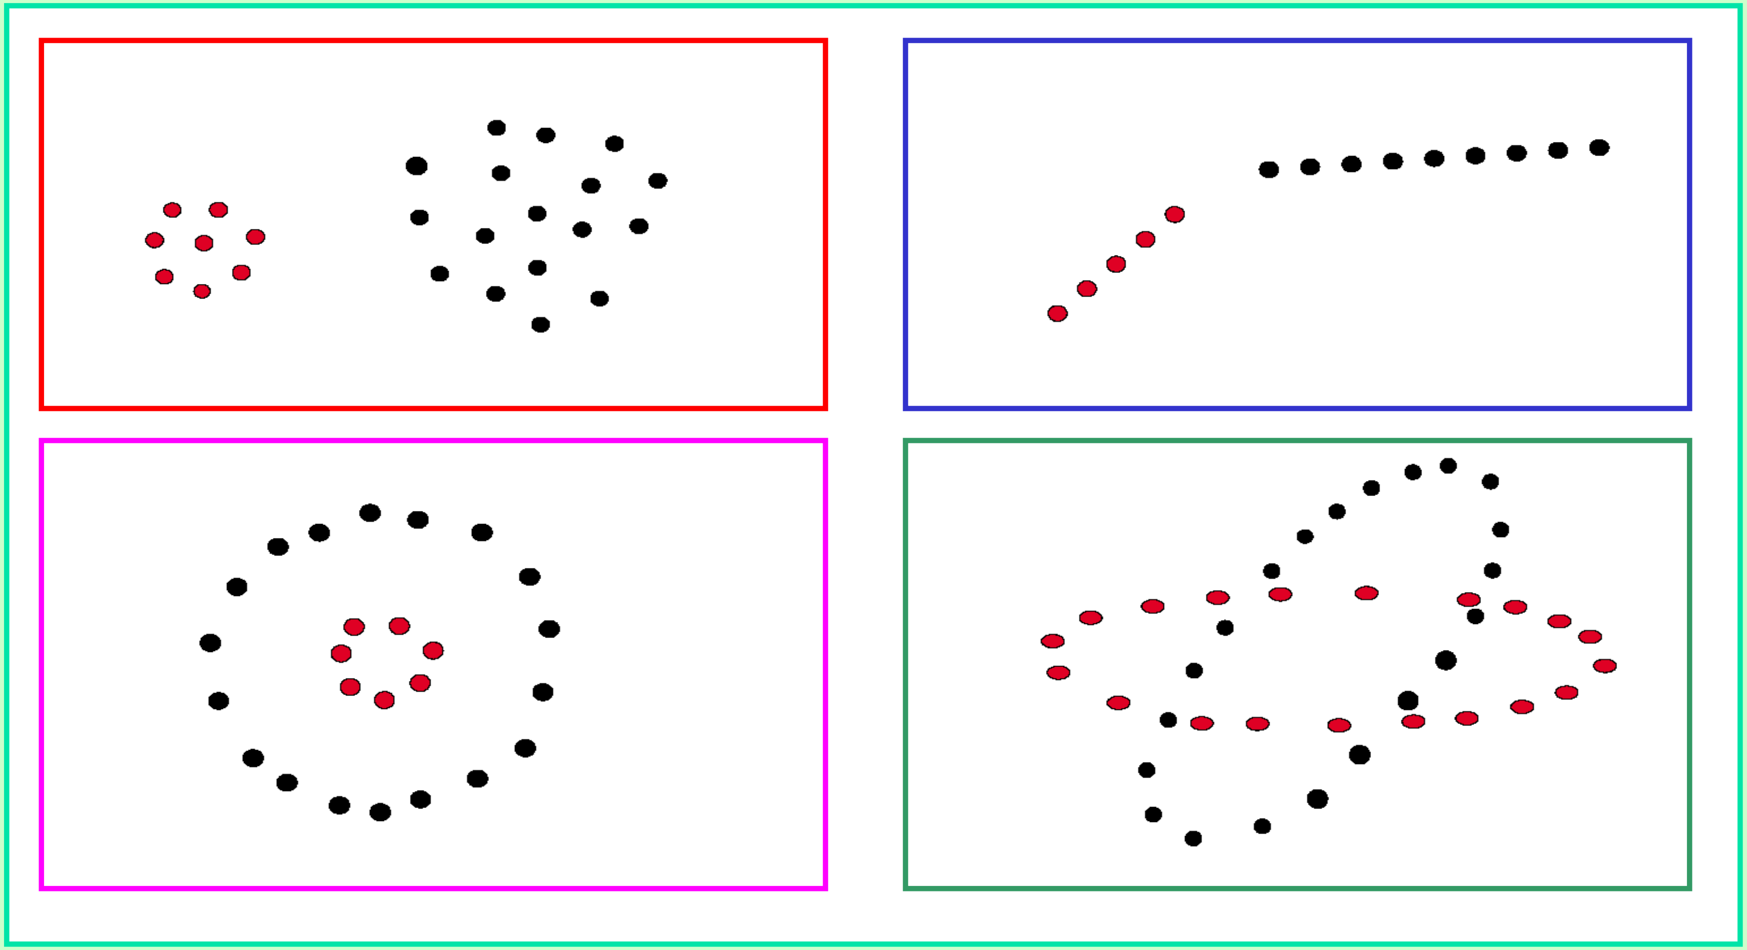

Problem Formulation

- Given is a set of data in $\boldsymbol{R}^n$ and the (estimated) number of clusters to look for (a difficult problem, more on this later).
- Find the partitioning of the data into subsets (clusters), such that samples within a subset are more similar to each other than to samples from other subsets.
- Similarity is mathematically formulated by using a distance measure (i.e., a dissimilarity function).
- Usually, each cluster will have a prototype and the distance is measured from this prototype.

Distance Measure

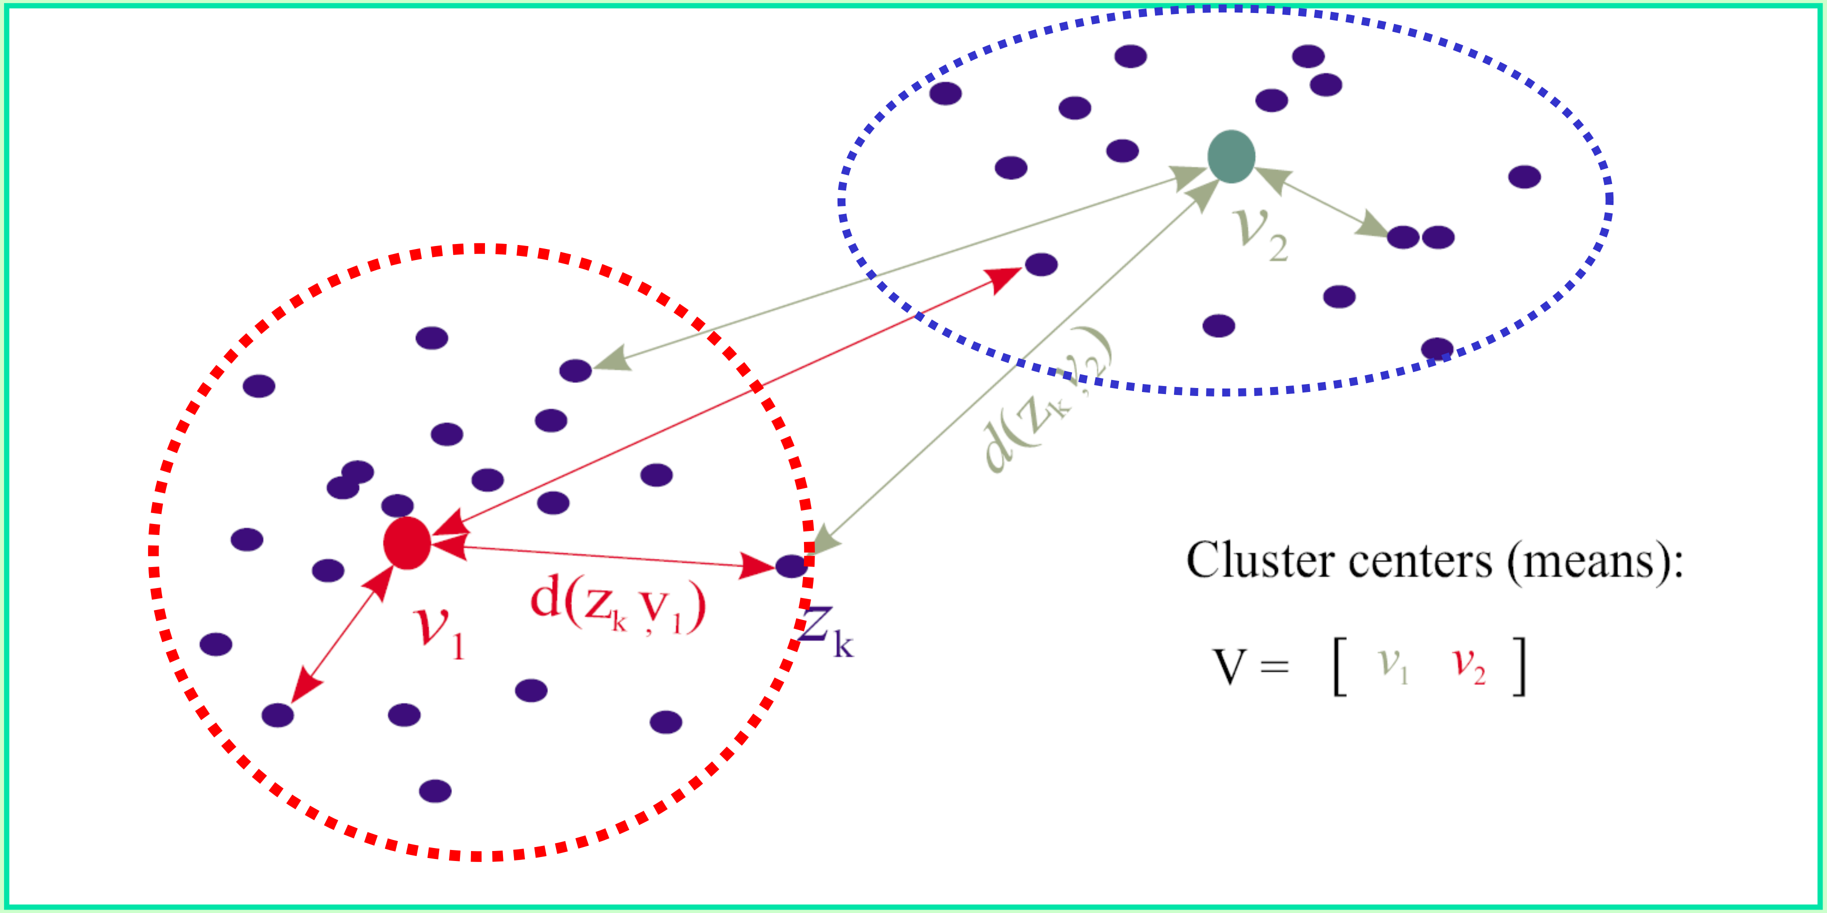

Distance Measures

- **Euclidean norm:**
  - $d^2(\mathbf{z}_j, \mathbf{v}_i) = (\mathbf{z}_j - \mathbf{v}_i)^T (\mathbf{z}_j - \mathbf{v}_i)$
- **Inner-product norm:**
  - $d^2{}_{\mathbf{A}_i}(\mathbf{z}_j, \mathbf{v}_i) = (\mathbf{z}_j - \mathbf{v}_i)^T \mathbf{A}_i (\mathbf{z}_j - \mathbf{v}_i)$
- **Many other possibilities . . .**

Generalized Prototypes (Varieties)

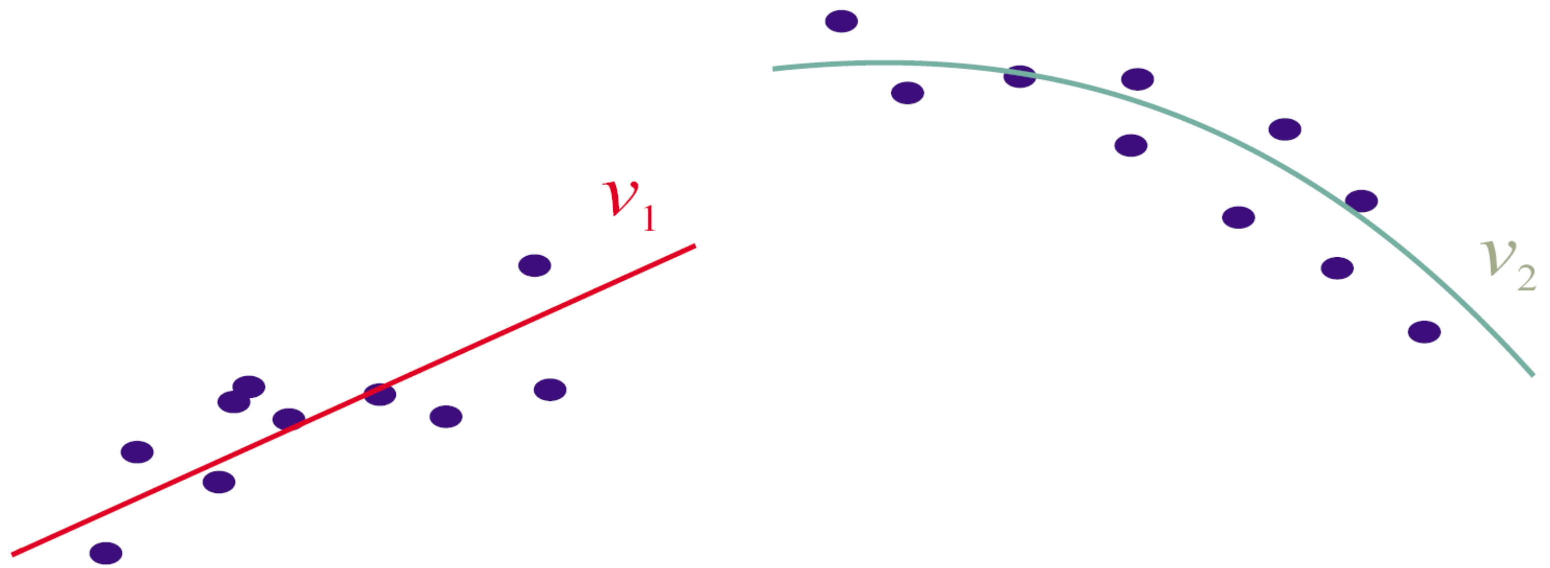

lines, circles, ellipses, regression functions, etc.

Corresponding Distance Measures

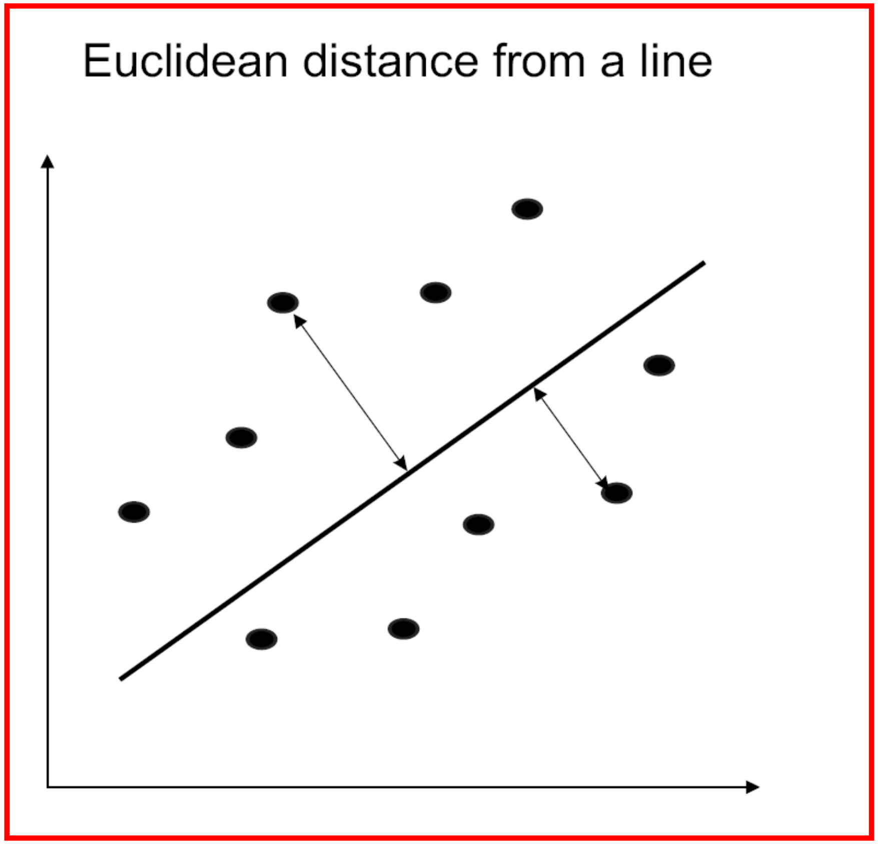

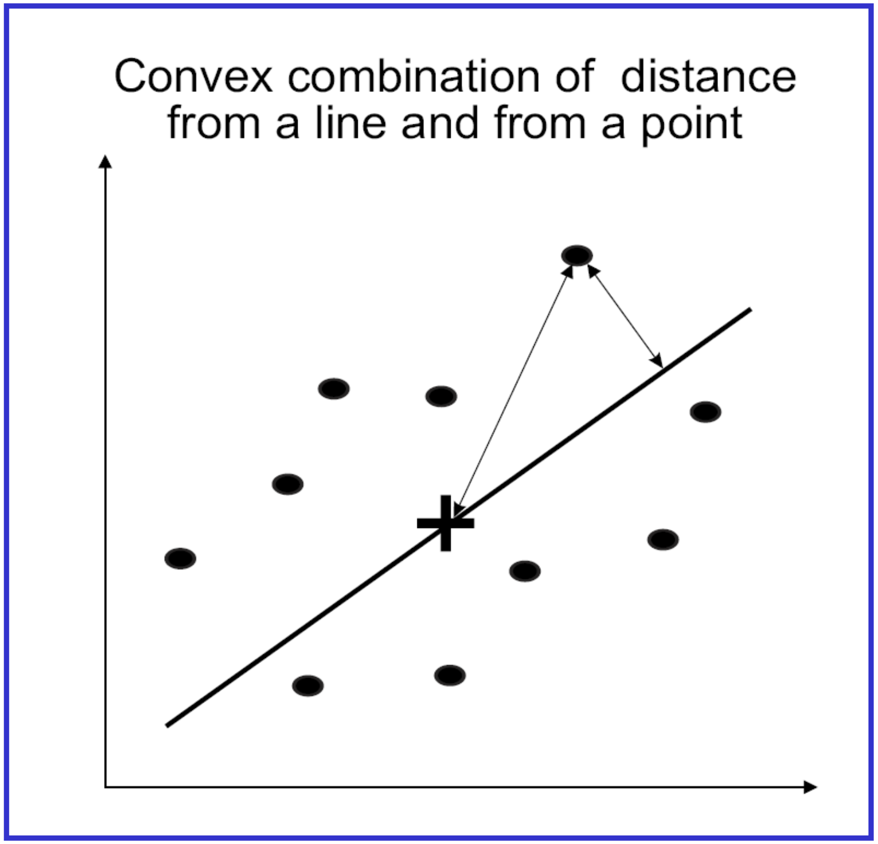

Mathematical Formulation of Clustering

- **Given the data:**

$$\mathbf{z}_k = \left[z_{1k},\ z_{2k},\ \ldots,\ z_{nk}\right]^T \in \Re^n,\ \ k = 1, \cdots, N$$

Find:

- **the partition matrix:**

$$\mathbf{U} = \begin{bmatrix} \mu_{11} & \cdots & \mu_{1k} & \cdots & \mu_{1N} \\ \vdots & \cdots & \vdots & \cdots & \vdots \\ \mu_{c1} & \cdots & \mu_{ck} & \cdots & \mu_{cN} \end{bmatrix}$$

- **and the cluster prototype (centres):**

$$\mathbf{V} = \left\{\mathbf{v}_1,\ \mathbf{v}_2,\ \ldots,\ \mathbf{v}_c\right\},\ \ \mathbf{v}_i \in \Re^n$$

Fuzzy Clustering: an Optimisation Approach

- **Objective function (least-squares criterion):**

$$J(\mathbf{Z};\mathbf{V},\mathbf{U},\mathbf{A}) = \sum_{i=1}^{c}\sum_{j=1}^{N} \mu_{i,j}^{m} d_{\mathbf{A}_i}^{2}(\mathbf{z}_j,\mathbf{v}_i)$$

- **subject to constraints:**

$0 \le \mu_{i,j} \le 1, \quad i = 1,\dots,c,\ j = 1,\dots,N$		membership degree
$0 < \sum_{j=1}^{N} \mu_{i,j} < 1, \quad i = 1,\dots,c$		no cluster empty
$\sum_{i=1}^{c} \mu_{i,j} = 1, \quad j = 1,\dots,N$		total membership

Fuzzy c-Means Algorithm

Repeat:

1. **Compute cluster prototypes (means):**

$$v_i = \frac{\sum_{k=1}^{N} \mu_{i,k}^m \mathbf{z}_k}{\sum_{k=1}^{N} \mu_{i,k}^m}$$

2. **Calculate distances:**

$$d_{ik} = (\mathbf{z}_k - \mathbf{v}_i)^T (\mathbf{z}_k - \mathbf{v}_i)$$

3. **Update partition matrix:**

$$\mu_{ik} = \frac{1}{\sum_{j=1}^{c} (d_{ik}/d_{jk})^{1/(m-1)}}$$

until $\|\Delta \mathbf{U}\| < \epsilon$

$$(i = 1, \cdots, c. \; k = 1, \cdots, N)$$

Failure to Discover Non-Spherical Clusters

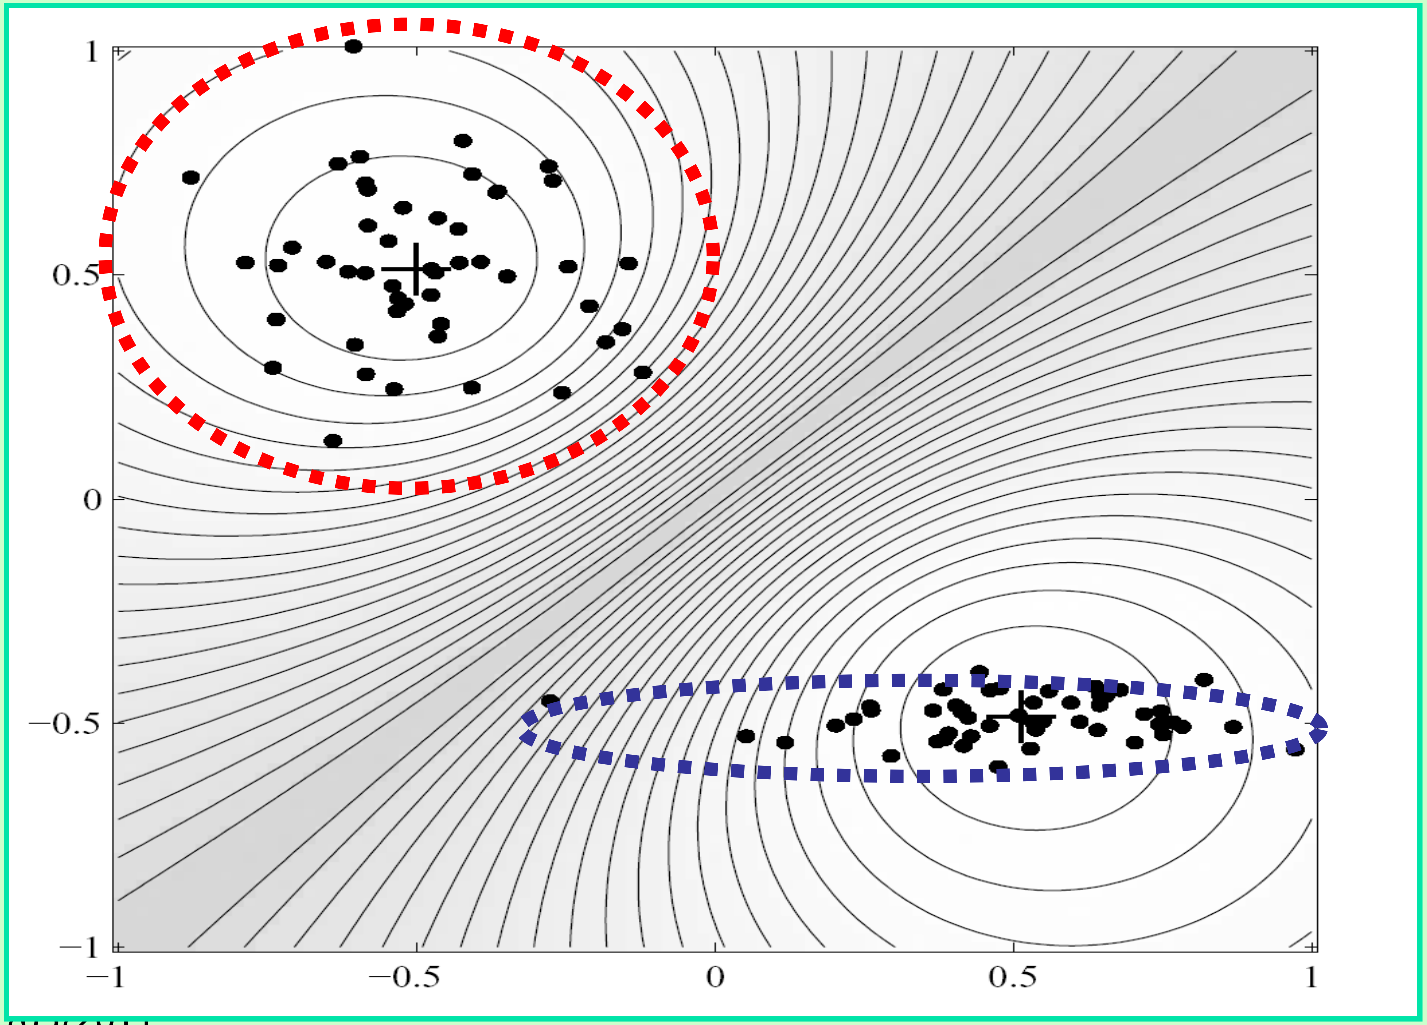

Adaptive Distance Measure

- **Inner-product norm:**

$$d^2_{\mathbf{A}_i}(\mathbf{z}_j, \mathbf{v}_i) = (\mathbf{z}_j - \mathbf{v}_i)^T \mathbf{A}_i (\mathbf{z}_j - \mathbf{v}_i)$$

- **norm-inducing matrix**

$$\mathbf{A}_i = \rho_i \det(\mathbf{F}_i)^{1/n} \mathbf{F}_i^{-1}$$

- **covariance matrix**

$$\mathbf{F}_i = \frac{\sum_{k=1}^{N} \mu_{ik}^m (\mathbf{z}_k - \mathbf{v}_i)(\mathbf{z}_k - \mathbf{v}_i)^T}{\sum_{k=1}^{N} \mu_{ik}^m}$$

Inner-Product Norm

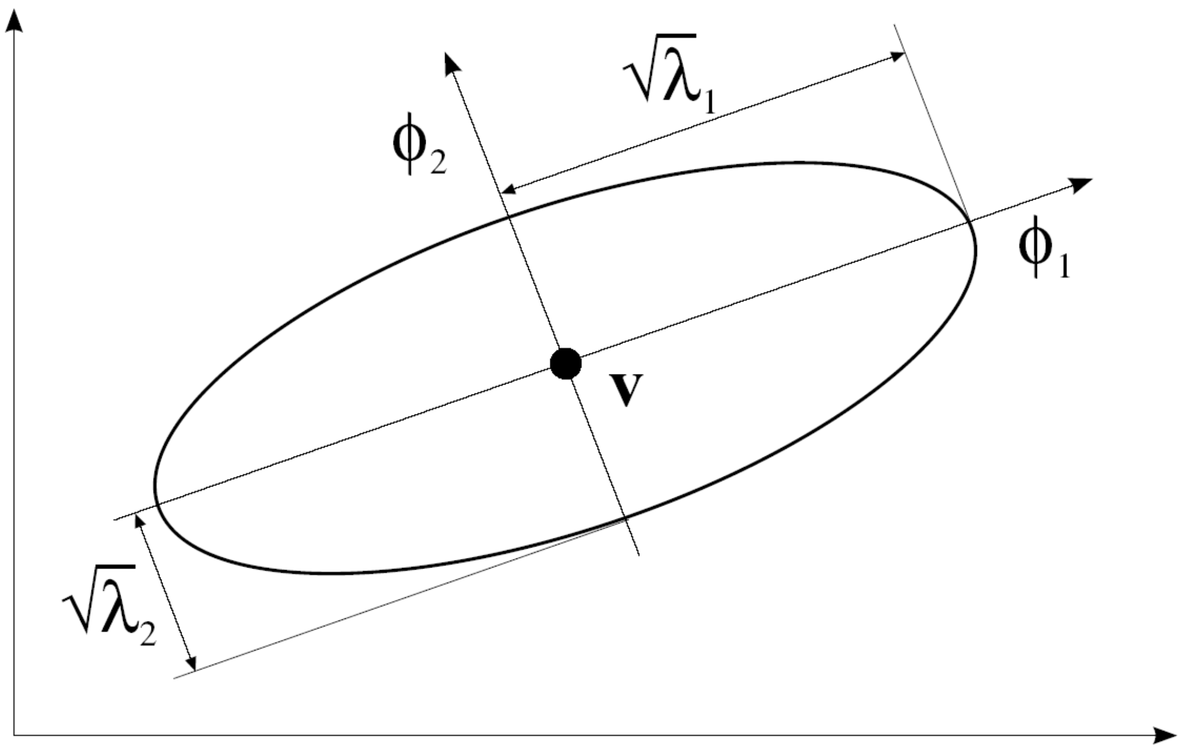

ellipsoid: $(\mathbf{z}-\mathbf{v})^T \mathbf{F}^{-1} (\mathbf{z}-\mathbf{v}) = \mathrm{const}$

Gustafson–Kessel Algorithm

Repeat:

1. Compute cluster prototypes (means):

$$v_i = \frac{\sum_{k=1}^{N} \mu_{i,k}^m \mathbf{z}_k}{\sum_{k=1}^{N} \mu_{i,k}^m}$$

2. Compute covariance matrices:

$$\mathbf{F}_i = \frac{\sum_{k=1}^{N} \mu_{ik}^m (\mathbf{z}_k - \mathbf{v}_i)(\mathbf{z}_k - \mathbf{v}_i)^T}{\sum_{k=1}^{N} \mu_{ik}^m}$$

3. Compute distances:

$$d_{ik} = (\mathbf{z}_k - \mathbf{v}_i)^T \rho_i \det(\mathbf{F}_i)^{1/n} \mathbf{F}_i^{-1} (\mathbf{z}_k - \mathbf{v}_i)$$

4. Compute partition matrix:

$$\mu_{ik} = \frac{1}{\sum_{j=1}^{c} (d_{ik}/d_{jk})^{1/(m-1)}}$$

until $\|\Delta \mathbf{U}\| < \epsilon$

Clusters of Different Shape and Orientation

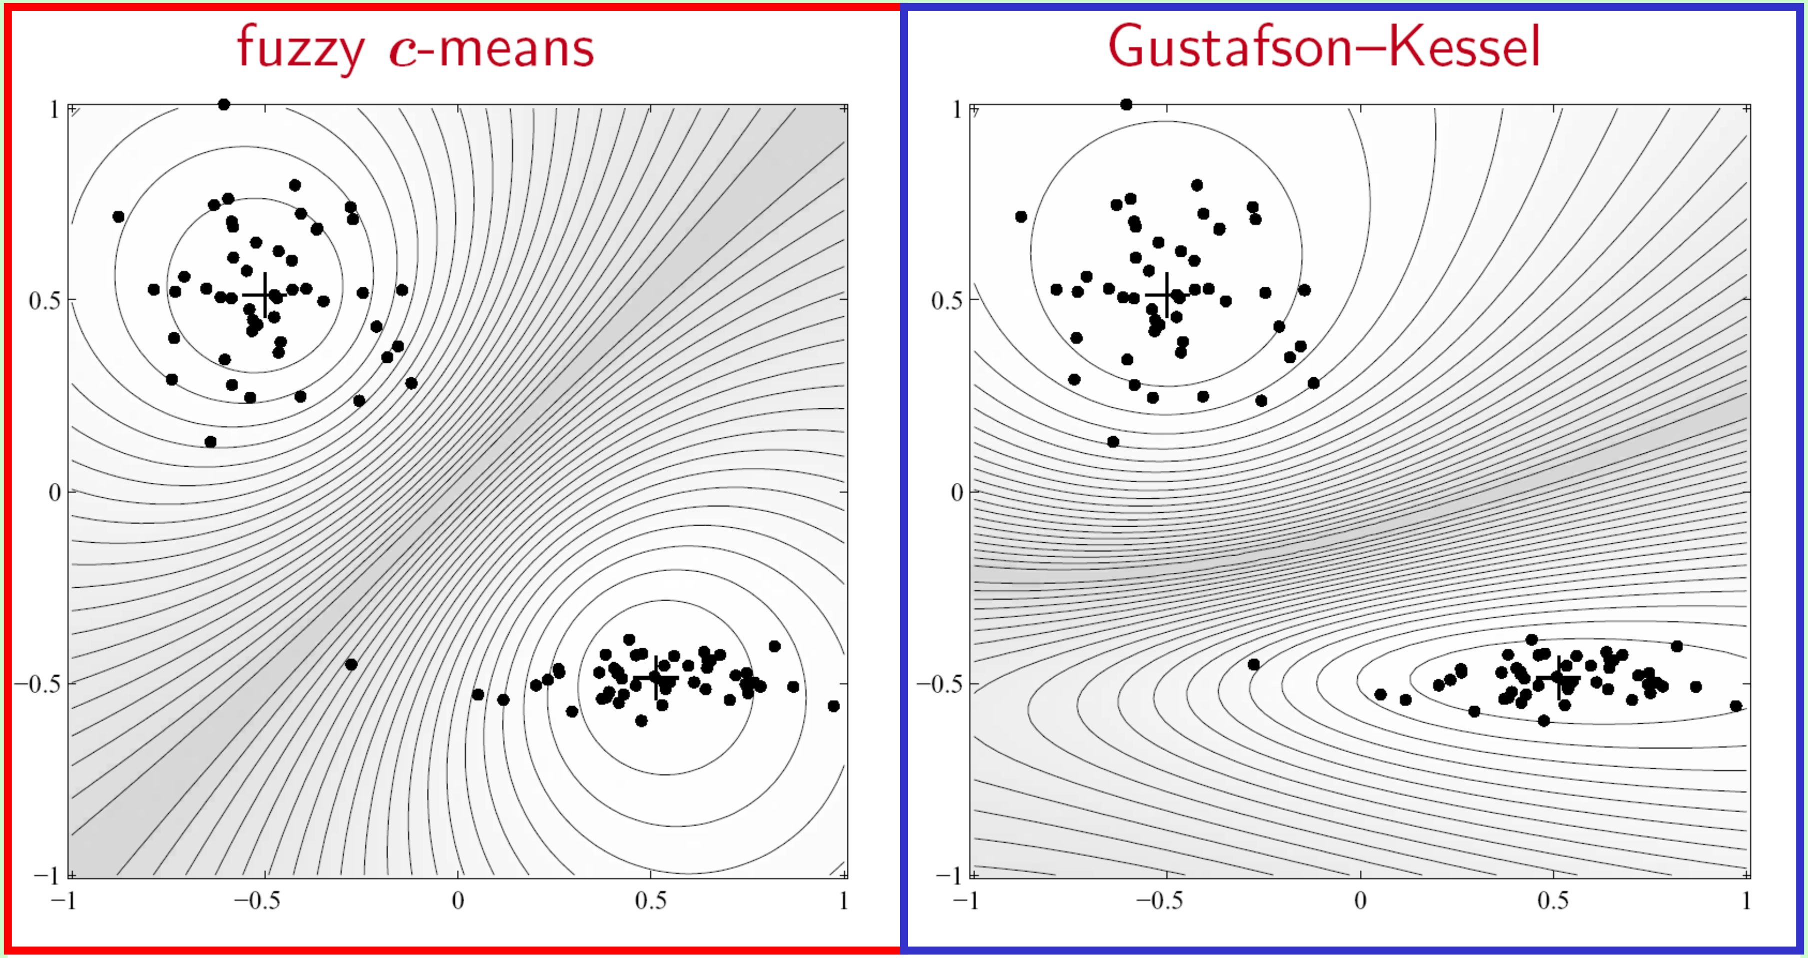

Number of Clusters

Validity measures

- Fuzzy hypervolume:

$$V_h = \sum_{i=1}^{c} [\det(\mathbf{F}_i)]^{1/2}$$

- Average within-cluster distance:

$$D_w = \frac{1}{c} \sum_{i=1}^{c} \frac{\sum_{k=1}^{N} \mu_{ik}^m D_{ik}^2}{\sum_{k=1}^{N} \mu_{ik}^m}$$

- Xie Beni index . . .

Validity Measures: Example

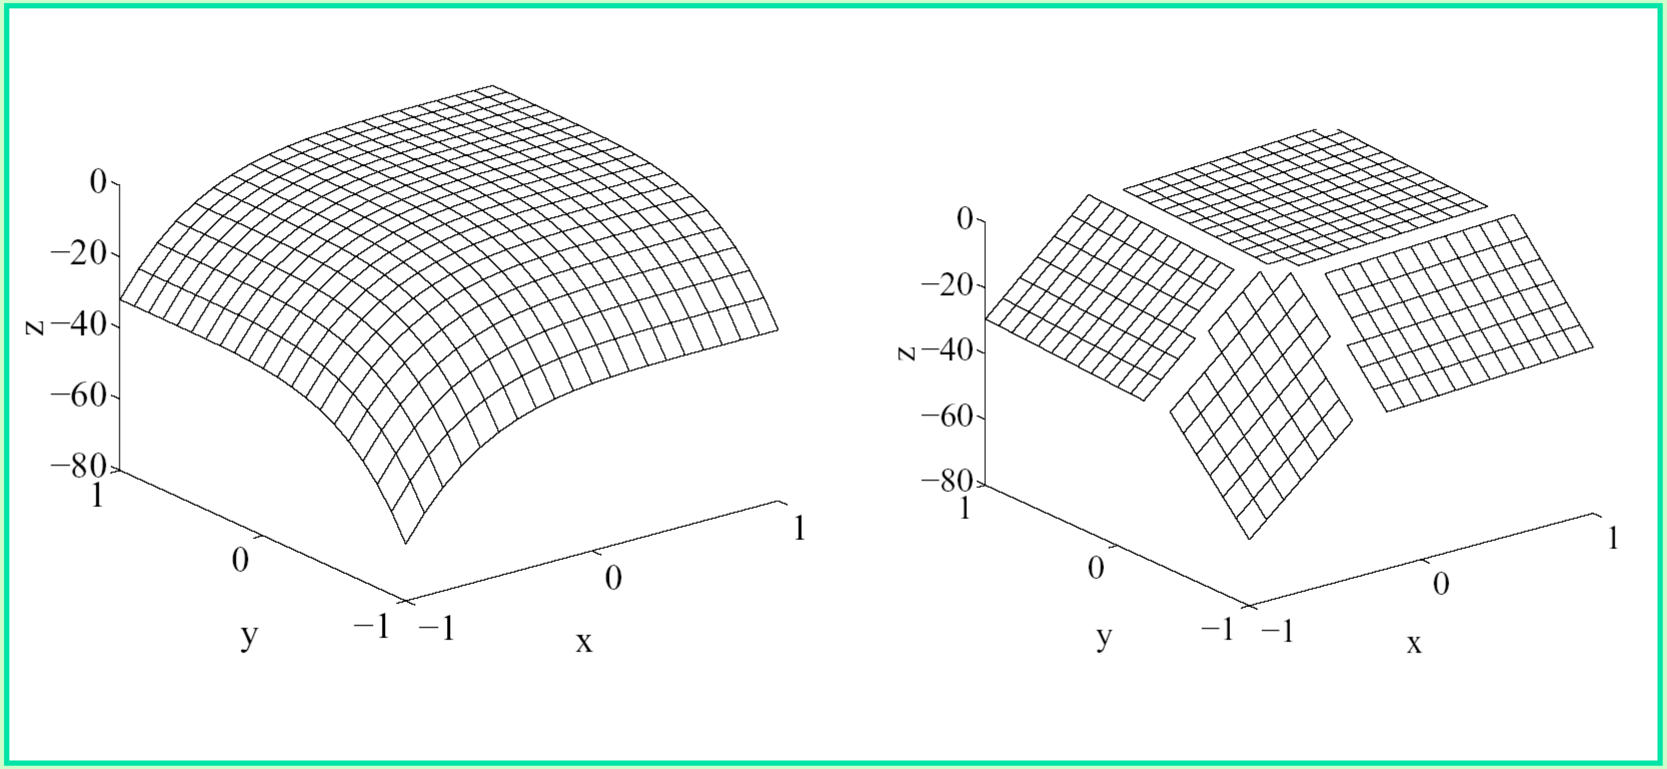

Data over 4 clusters

Validity Measures

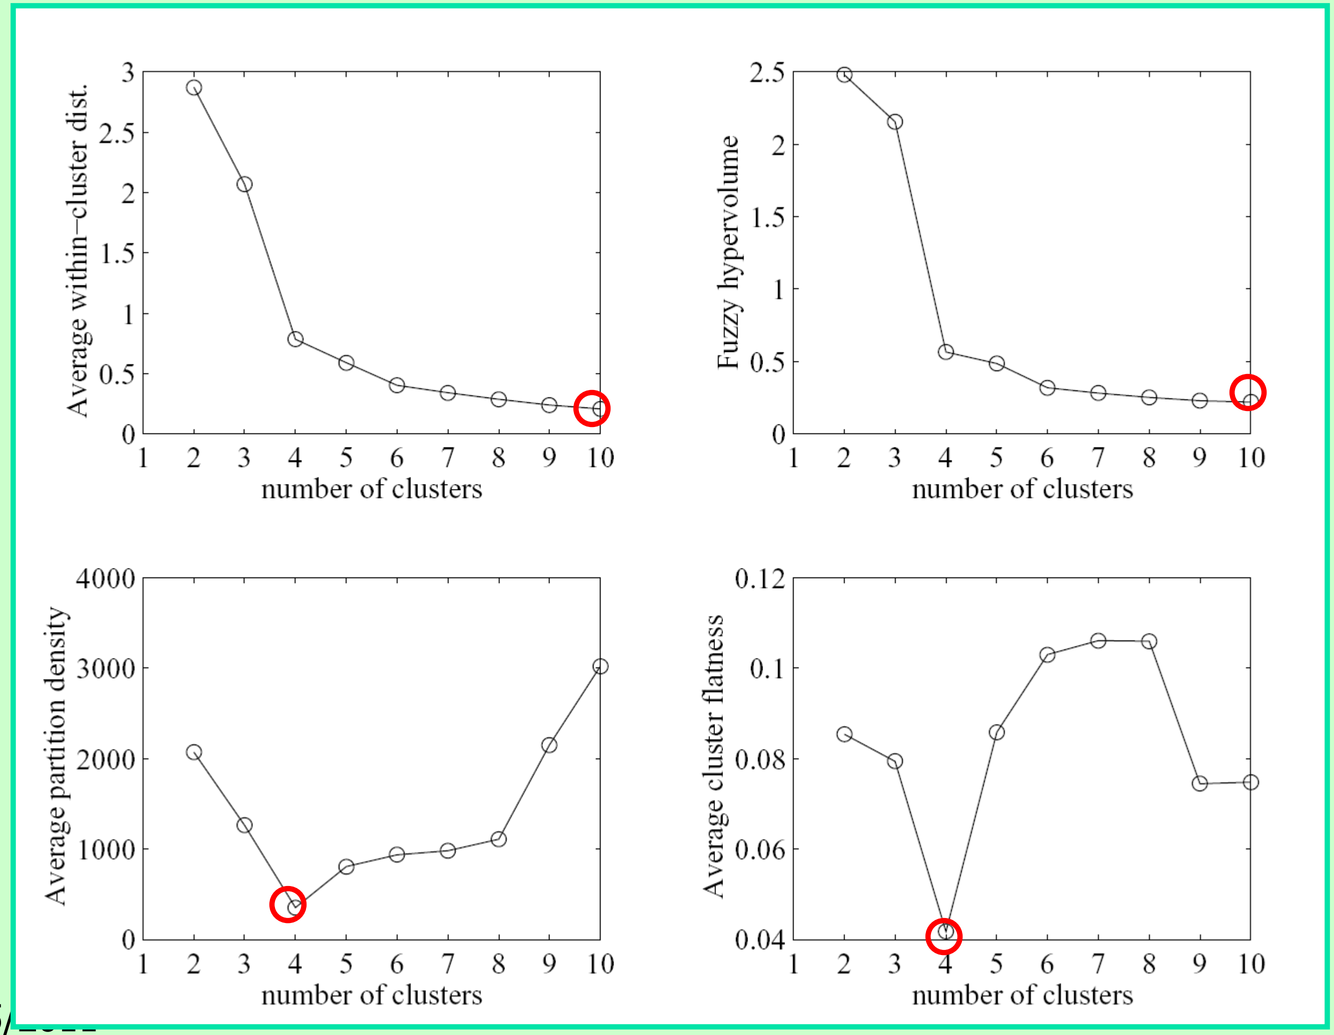

Number of Clusters

Compatible cluster merging

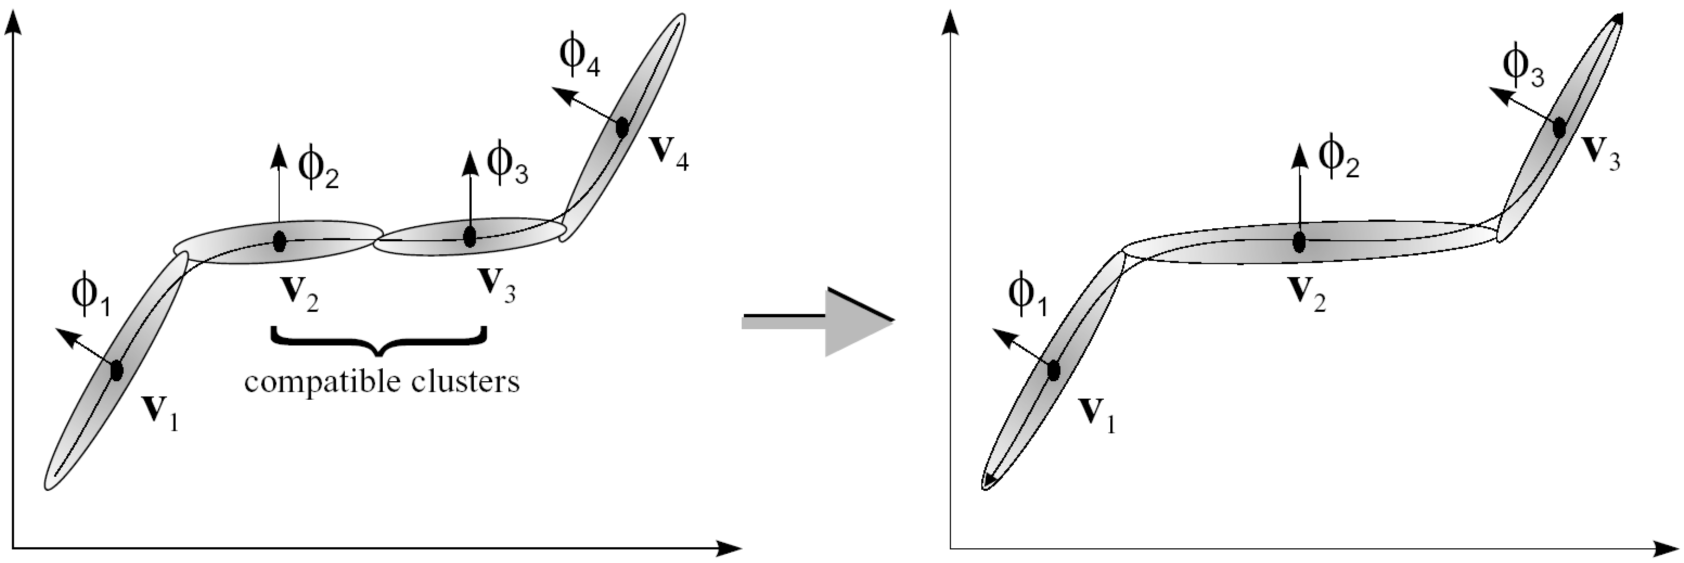

$$|\ \Phi_i \cdot \Phi_j\ | \geq k_1, \quad k_1 \rightarrow 1 \quad \text{and} \quad \|v_i - v_j\| \leq k_2, \quad k_2 \rightarrow 0$$

Data-Driven (Black-Box) Modelling

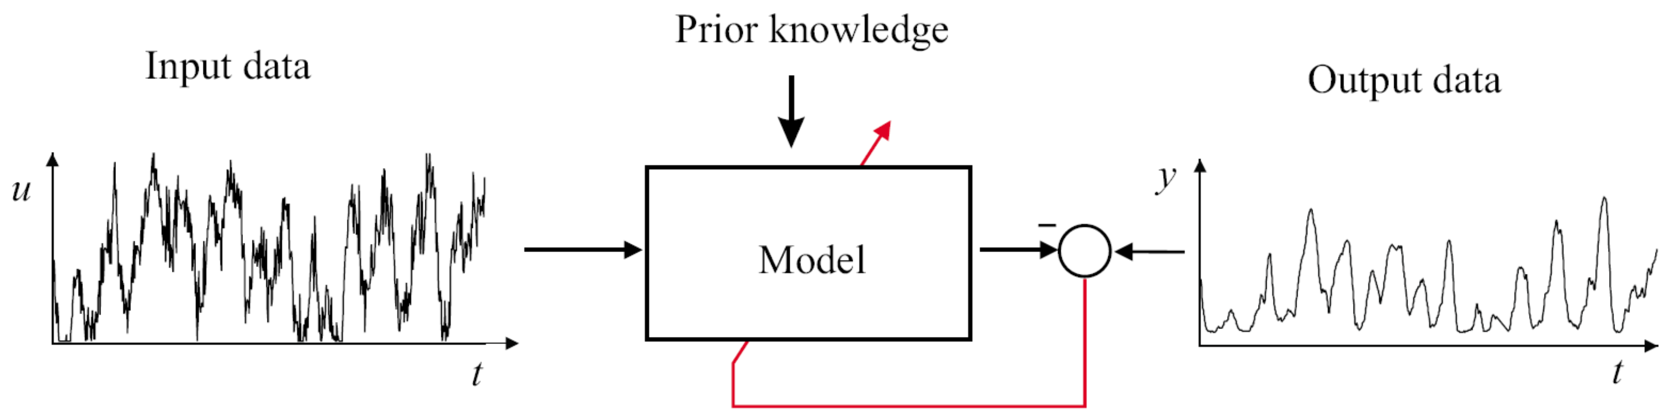

- Linear model (for linear systems only, limited in use)
- Neural network (black box, unreliable extrapolation)
- Rule-based model (more transparent, 'grey-box')

Extraction of Rules by Fuzzy Clustering

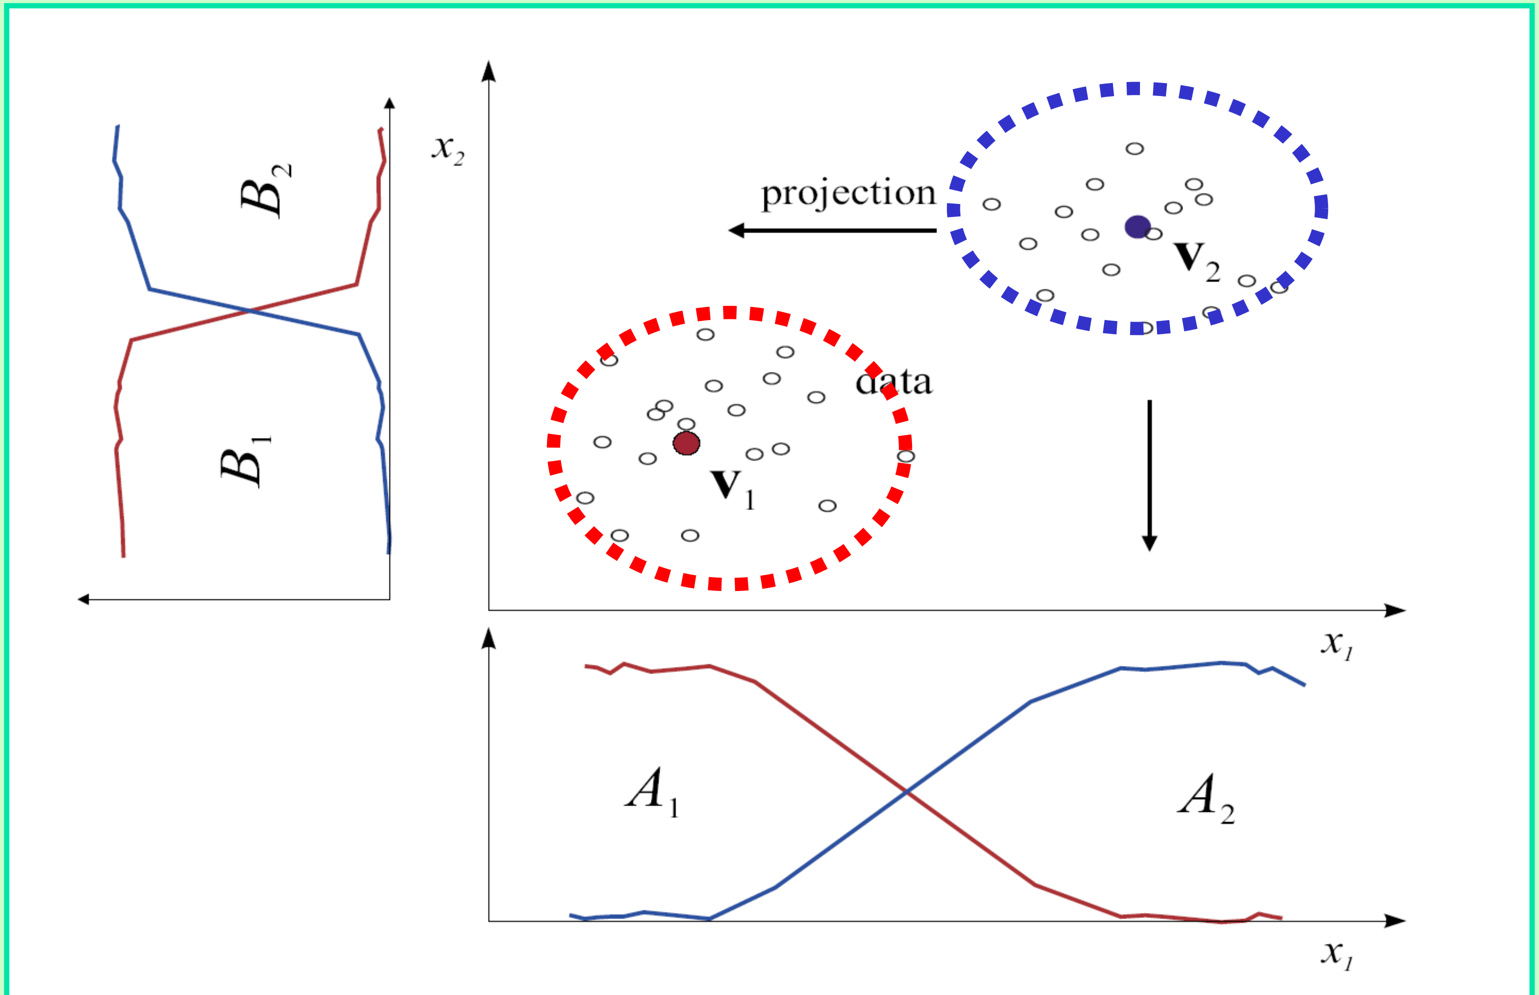

Extraction of Rules by Fuzzy Clustering

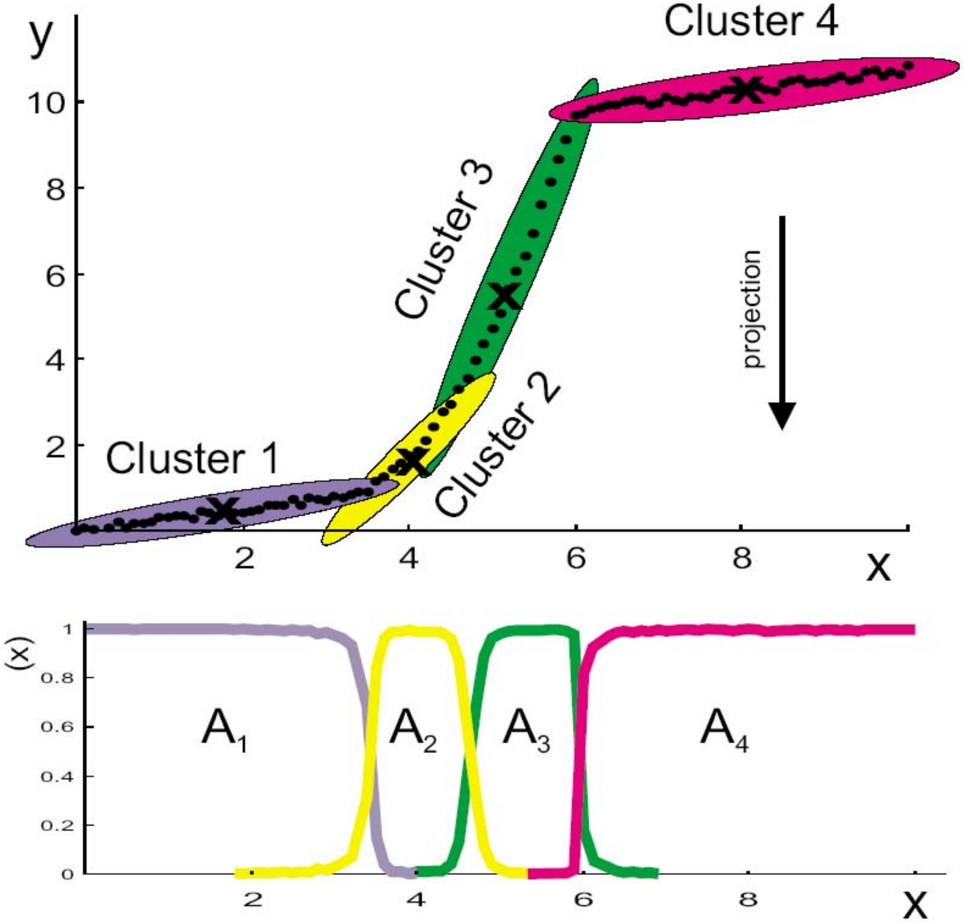

Takagi-Sugeno model

Rule-based description:

If x is A_1 then $y = a_1x + b_1$
If x is A_2 then $y = a_2x + b_2$

etc...

Example: Non-linear Autoregressive System (NARX)

$$x(k+1) = f(x(k)) + \epsilon(k)$$

$$f(x) = \begin{cases} 2x - 2, & 0.5 < x \\ -2x, & -0.5 \le x < 0.5 \\ 2x + 2, & x \le -0.5 \end{cases}$$

Structure Selection and Data Preparation

1. Choose model order p

$$x(k+1) = f(\underbrace{x(k), x(k-1), \ldots, x(k-p+1)}_{\mathbf{x}(k)})$$

2. Form pattern matrix **Z** to be clustered

$$\mathbf{Z}^T = \begin{bmatrix} x(1) & x(2) & \ldots & x(p) & x(p+1) \\ x(2) & x(3) & \ldots & x(p+1) & x(p+2) \\ \vdots & \vdots & & \vdots & \vdots \\ x(N-p) & x(N-p+1) & \ldots & x(N-1) & x(N) \end{bmatrix}$$

Clustering Results

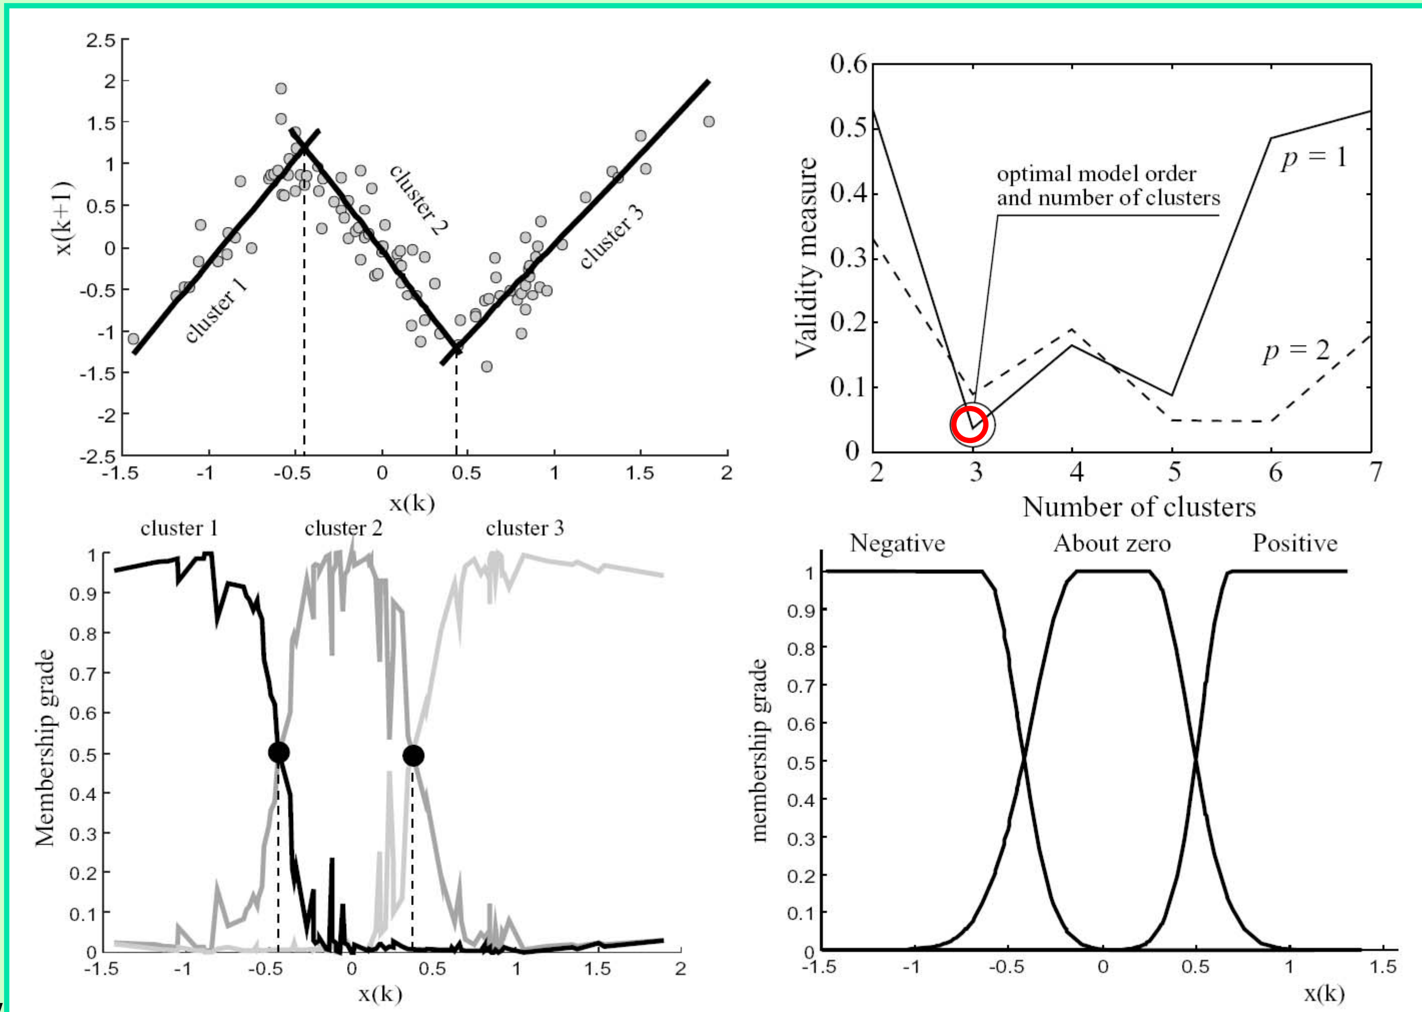

Rules Obtained

1) If $x(k)$ is $Positive$ then $x(k+1) = 2.0244x(k) - 2.0289$

2) If $x(k)$ is $About\ zero$ then $x(k+1) = -1.8852x(k) + 0.0005$

3) If $x(k)$ is $Negative$ then $x(k+1) = 1.9050x(k) + 1.9399$

original function: $$f(x) = \begin{cases} 2x-2, & 0.5 < x \\ -2x, & -0.5 \leq x < 0.5 \\ 2x+2, & x \leq -0.5 \end{cases}$$

Identification of Pressure Dynamics

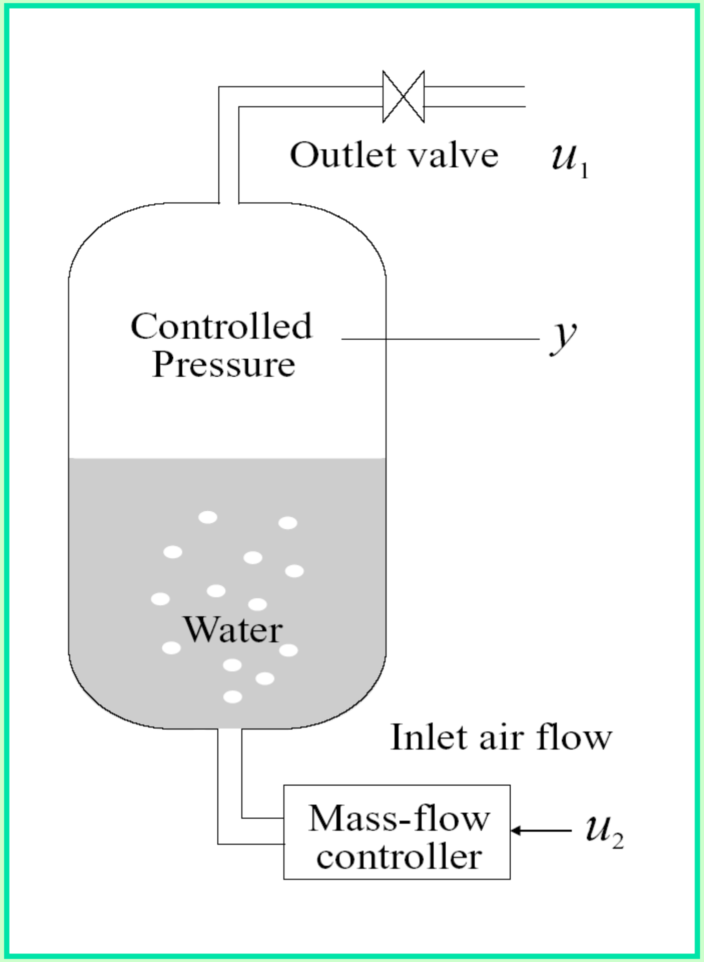

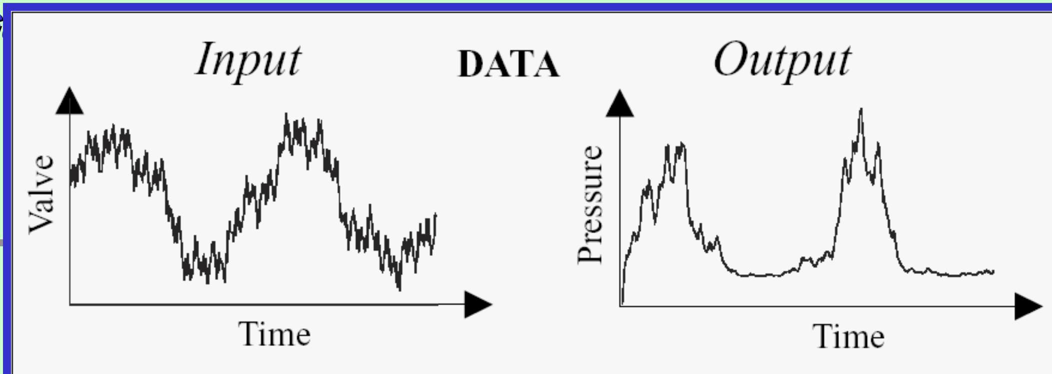

Rules **FUZZY MODEL** *Membership f.*

1) If Valve is Open and
Pressure is Low then ....

2) If Valve is Closed and
Pressure is High then ....

3) ...

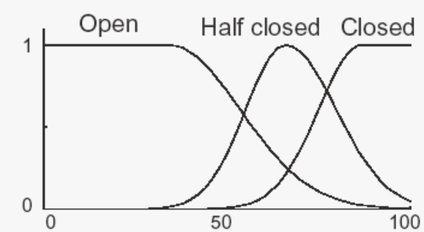

PREDICTIVE CONTROLLER

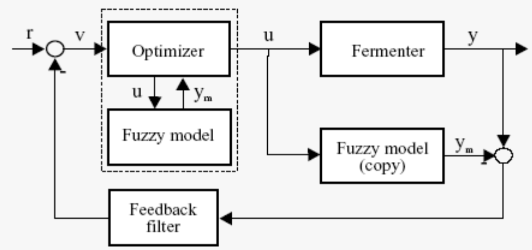

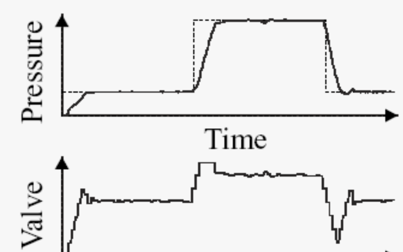

Concluding Remarks

- **Optimisation approach to clustering**
  - ✓ effective for metric (e.g., real-valued) data
  - ✓ accurate results for small to medium complexity problems
  - ✓ for large problems, convergence to local optima, slow
- **Many other techniques**
  - ✓ agglomerative methods
  - ✓ hierarchical splitting methods
  - ✓ graph-theoretic methods
- **Variety of applications**

Application Examples

Neural Networks for

Non-linear Identification

Nonlinear System Identification

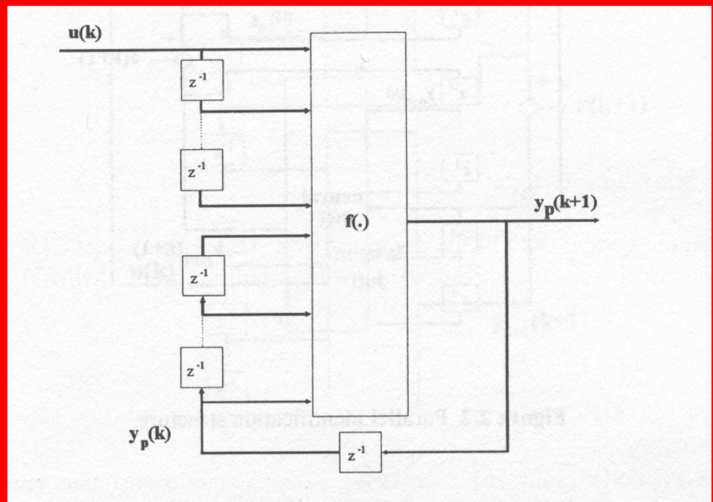

Figure 2.1 Input-output model

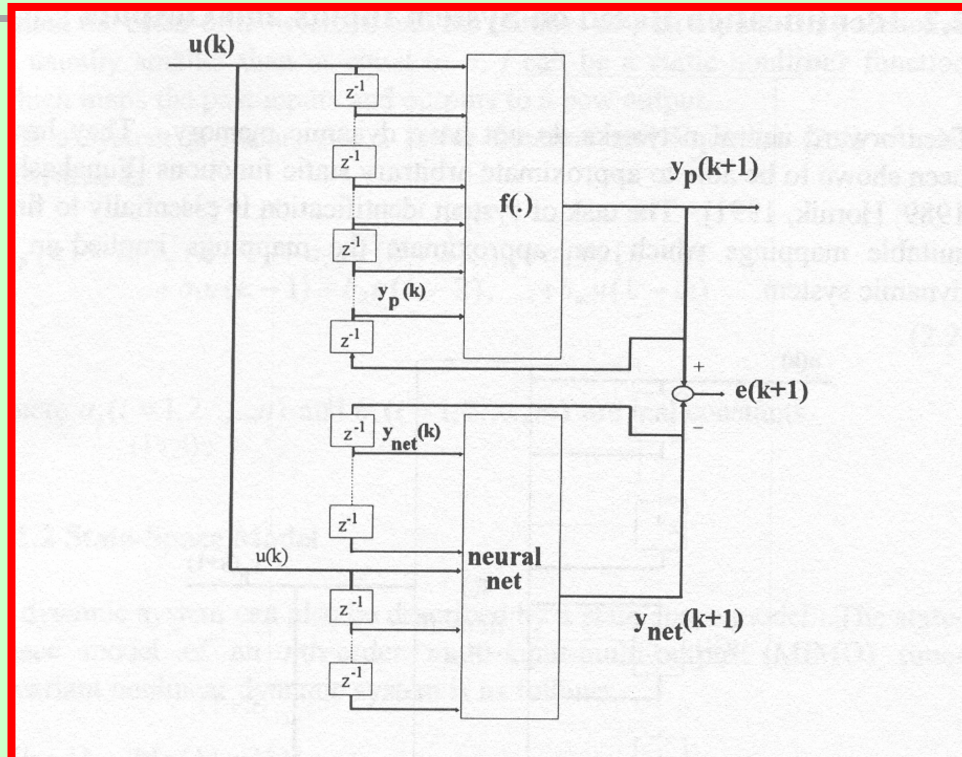

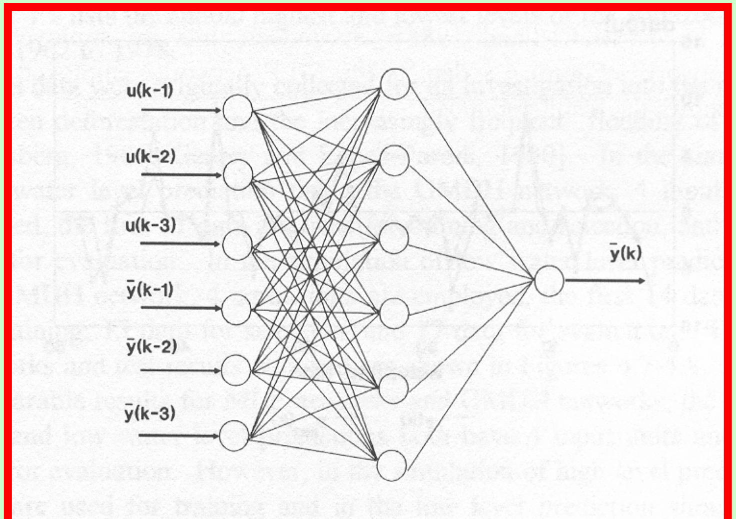

Target function: $y_p(k+1) = f(.)$

Identified function: $y_{NET}(k+1) = F(.)$

Estimation error: $e(k+1)$

Nonlinear System Identification

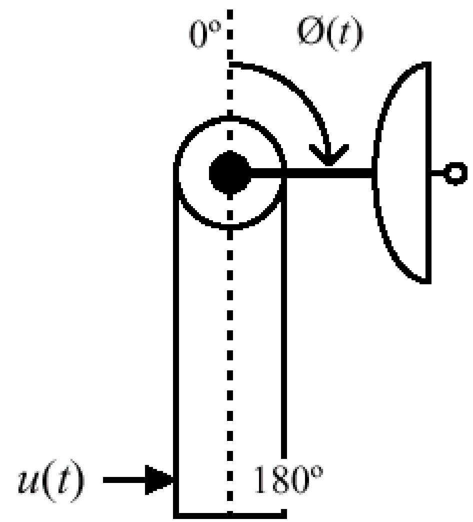

$$\frac{d}{dt}\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ 9.81 \sin x_1 - 2 x_2 + u \end{bmatrix}$$

$$x_1 = \varnothing$$

$$x_2 = \frac{d\varnothing}{dt}$$

```
deg2rad = pi/180;
angle = [-20:40:200]*deg2rad;
vel = [-90:36:90]*deg2rad;
force = -30:6:30;
```

```
angle2 = [-20:10:200]*deg2rad;
Pm = [combvec(angle,vel,force);
   [angle2; zeros(2,length(angle2))]];
```

Neural network input generation Pm

Nonlinear System Identification

```
S1 = 8;
[S2,Q] = size(Tm);
mnet = newff(minmax(Pm),[S1 S2],{'tansig' 'purelin'},'trainlm');
```

```
mnet.trainParam.goal = (0.0037^2);
mnet = train(mnet,Pm,Tm);
```

Neural network target Tm

Neural network response (angle & velocity)

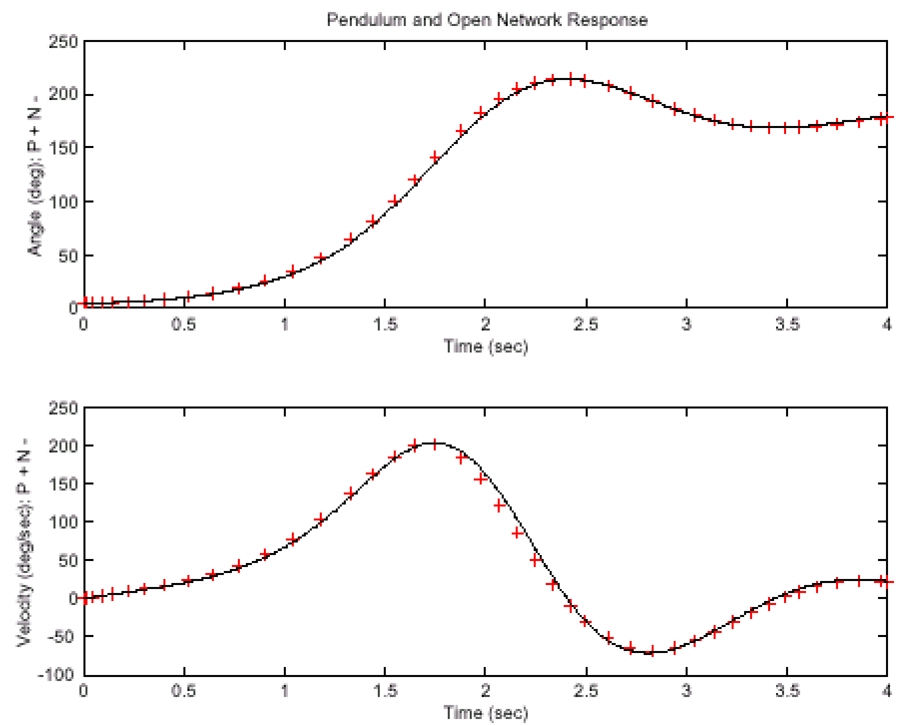

Matlab NNtool GUI (Graphical User Interface)