

Introduction to Artificial Intelligence and its applications

Dr. Mohammed Abu Mallouh
Mechatronics Eng.
Hashemite University
Zarqa, Jordan



- Boston Dog
- Use machine learning.



Play :VI Boston Dog

What is Artificial Intelligence (AI)

- **Intelligence** is related with human ability to store and recall fact, solve a given problem based on known fact and relevant theorem.
- Artificial Intelligence (**AI**) is the ability of an electronic device (**computer**) to accomplish any tasks that ordinary would have been handled by human.

Another definition for AI

-**Intelligence** is the ability to understand and learn things.

-**Intelligence** is the ability to think and understand instead of doing things by instinct or automatically.

(Essential English Dictionary, Collins, London, 2008)


Intelligence as **the ability to learn and understand, to solve problems and to make decisions.**

Another definition for AI

- Intelligence:
“the capacity to learn and solve problems” (Webster’s dictionary)
- In particular, the ability to solve novel problems the ability to act rationally the ability to act like humans.

Another definition for AI


- The term artificial intelligence was first coined by John McCarthy in 1956 when he held the first academic conference on the subject. • Isn't there a solid definition of intelligence that doesn't depend on relating it to human intelligence? Not yet. The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent. We understand some of the mechanisms of intelligence and not others.



-One of the most significant papers on machine intelligence, “**Computing Machinery and Intelligence**”, was written by the British mathematician **Alan Turing** over fifty years ago . However, it still stands up well under the test of time, and the Turing’s approach remains universal.


-Turing did not provide definitions of machines and thinking, he just avoided semantic arguments by inventing a game , the **Turing Imitation Game**.

■ Turing Imitation Game: Phase 1



-The imitation game originally included two phases. In the first phase, the interrogator, a man and a woman are each placed in separate rooms. The interrogator's objective is to work out who is the man and who is the woman by questioning them. The man should attempt to deceive the interrogator that he is the woman, while the woman has to convince the interrogator that she is the woman.

Turing Imitation Game: Phase 2




-In the second phase of the game, the man is replaced by a computer programmed to deceive the interrogator as the man did. It would even be programmed to make mistakes and provide fuzzy answers in the way a human would . If the computer can fool the interrogator as often as the man did, we may say this computer has passed the intelligent behavior test.

- The history of artificial intelligence
- The birth of artificial intelligence(1943-1956)
- The rise of artificial intelligence, or the era of great expectations (1956-late 1960s)
- In the sixties, AI researchers attempted to simulate the thinking process by inventing general methods for solving broad classes of problems. They used the general-purpose search mechanism to find a solution to the problem. Such approaches, now referred to as **weak methods**, applied weak information about the problem domain.

Unfulfilled promises, or the impact of reality
(late 1960s-early 1970s)


The main difficulties for AI in the late 1960s were :

- Because AI researchers were developing general methods for broad classes of problems, early programs contained little or even no knowledge about a problem domain.



-The technology of expert systems, or the key to success(early 1970s -mid -1980s)

-Probably the most important development in the seventies was the realization that the domain for intelligent machines had to be sufficiently restricted. Previously, AI researchers had believed that clever search algorithms and reasoning techniques could be invented to emulate general, human-like, problem-solving methods. A general-purpose search mechanism could rely on elementary reasoning steps to find complete solutions and could use weak knowledge about domain.




-When weak methods failed, researchers finally realized that the only way to deliver practical results was to solve typical cases in narrow areas of expertise, making large reasoning steps.

-A 1986 survey reported a remarkable number successful expert system applications in different areas: chemistry, electronics, engineering, geology, management, medicine, process control and military science (Waterman, 1986). Although Waterman found nearly 200 expert systems , most of the applications were in the field of medical diagnosis. Seven years later a similar survey reported over 2500 developed expert systems (Durkin, 1994). The new growing area was business and manufacturing, which accounted for about 60% of the applications. Expert system technology had clearly matured.

What is Artificial Intelligence (AI)

- Computational models of human behavior?
- Programs that behave (externally) like humans
 - Computational models of human “thought”
- Programs that operate (internally) the way humans do

Artificial Intelligence Techniques

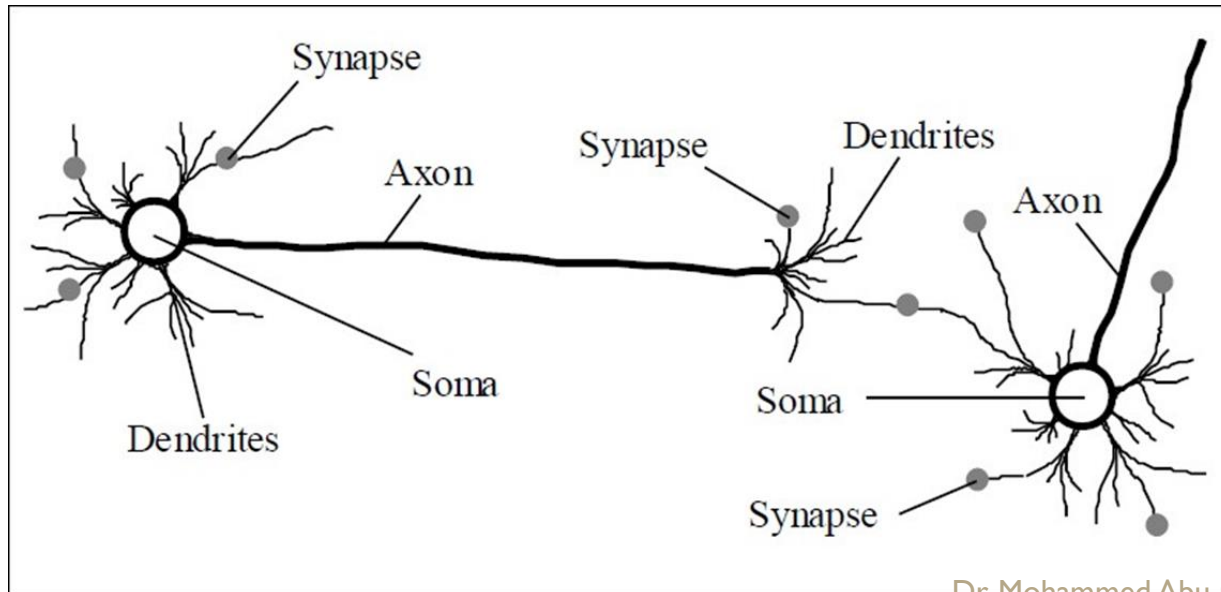


Applications

- Control
- Estimation
- System Identification
- Optimization


Biological Neural network

- Computational models of human “thought”
- Programs that operate (internally) the way humans do
- The brain consists of a densely interconnected set of nerve cells, or basic information-processing units, called neurons.
- The human brain incorporates nearly 10 billion neurons and 60 trillion connections, synapses, between them.




Artificial Neural network

Architecture of a typical artificial neural network



Artificial Neural network

The neuron as a simple computing element
Diagram of a neuron



Artificial Neural network

Advantages

- Learning capabilities
- Generalization
- No Mathematical model
- Fault tolerance
- Parallel processing

Drawbacks

- Lack of design techniques
- Computational effort

Applications

- Control
- Estimation
- System Identification
- Optimization

Genetic algorithms

Evolutionary computation, or learning by doing
(early 1970s-onwards)

-Natural intelligence is a product of evolution.

Therefore, by simulating biological evolution, we might expect to discover how living systems are propelled towards high-level intelligence.

-Nature learns by doing; biological systems are not told how to adapt to a specific environment – they simply compete for survival.

Genetic algorithms

- The evolutionary approach AI is based on the computational models of natural selection and genetics.
- Evolutionary computation works by simulating a population of individuals, evaluating their performance, generating a new population, and repeating this process a number of times.
- Evolutionary computation combines three main techniques: **genetic algorithms**, **evolutionary strategies** and **genetic programming**.

-Advantages

- Derivative free
- Avoid local minimal

-Application

- Optimization
- Parameter tuning and estimation

Fuzzy logic

- Computational models of human behavior?
- Programs that behave (externally) like humans

The new era of knowledge engineering, or computing with words(late 1980s-onwards)

- Neural network technology offers more natural interaction with the real word than do systems based on symbolic reasoning. Neural network s can learn, adept to changes in problem's environment, establish patterns in situations where rules are not known, and deal with fuzzy or incomplete information.
- However, they lack explanation facilities and usually act as a black box. The process of training neural networks is slow, and frequent retraining can cause serious difficulties.

Fuzzy logic

- Very important technology dealing with vague, imprecise and uncertain knowledge and data is **fuzzy logic**.
- Human experts do not usually think in probability values, but in such terms as often, generally, sometimes, occasionally and rarely. Fuzzy logic is concerned with capturing the meaning of words, human reasoning and decision making. Fuzzy logic provides the way to break through the computational bottlenecks of traditional expert systems.
- At the heart of fuzzy logic lies the concept of a **linguistic variable**. The values of the linguistic variable are words rather than numbers.

Fuzzy logic

Benefits derived from the application of fuzzy logic models in knowledge-based and decision-support systems can be summarized as follow:

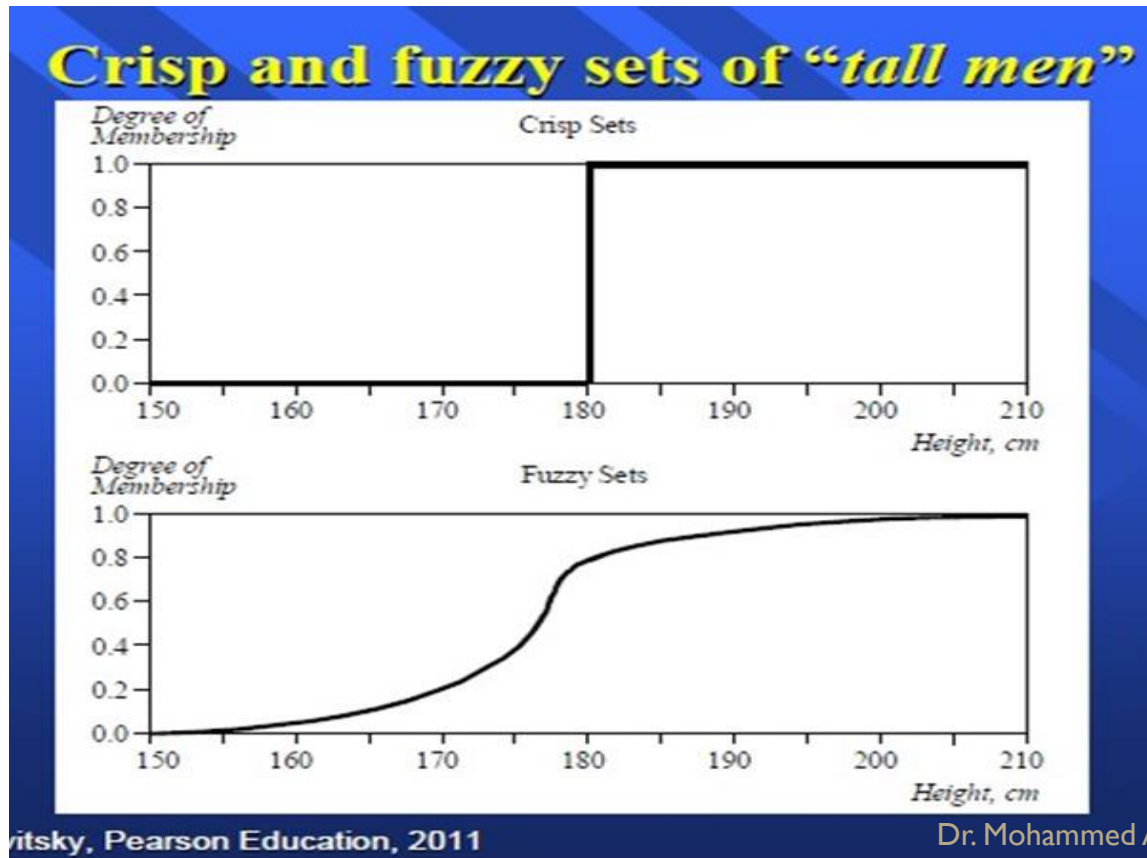
-Improved computational power: Fuzzy rule-based systems perform faster than conventional expert systems and require fewer rules. A fuzzy expert systems merges the rules, making them more powerful. Lotfi Zadeh believes that in a few years most expert systems will use fuzzy logic to solve highly nonlinear and computationally difficult problems.

Fuzzy logic

Computational models of human behavior?

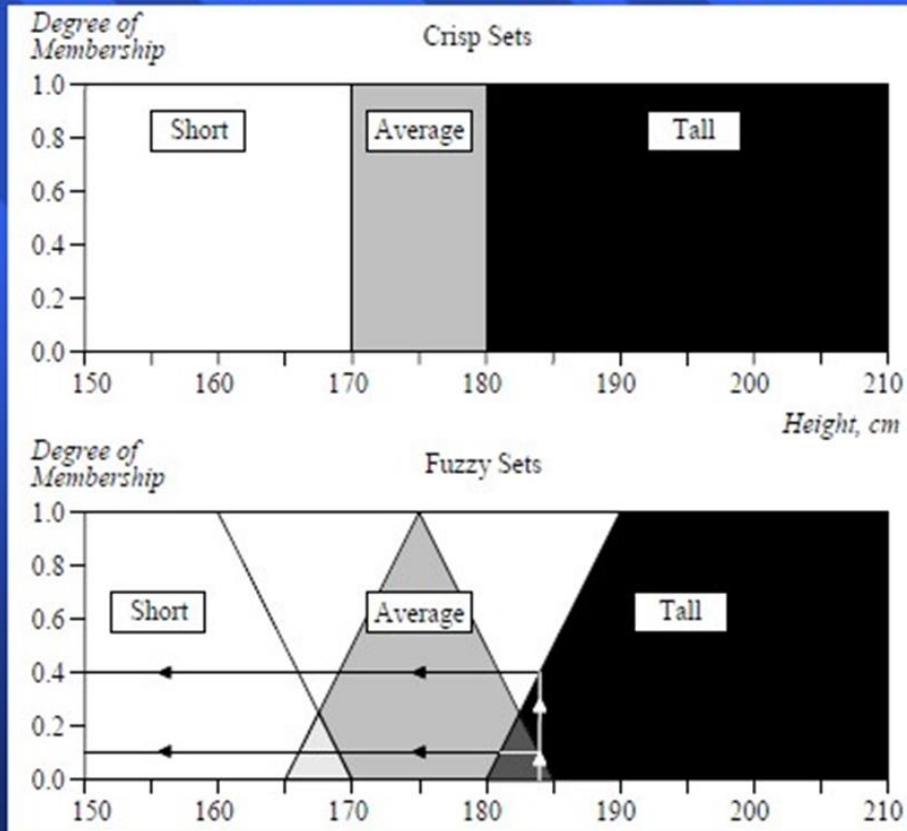
Programs that behave (externally) like humans

Unlike two-valued Boolean logic,, fuzzy logic is multi-valued.. It deals with degrees of membership and degrees of truth.. Fuzzy logic uses the continuum of logical values between 0 (completely false) and 1 (completely true).



Fuzzy logic

Crisp and fuzzy sets of short, average and tall men



Fuzzy logic

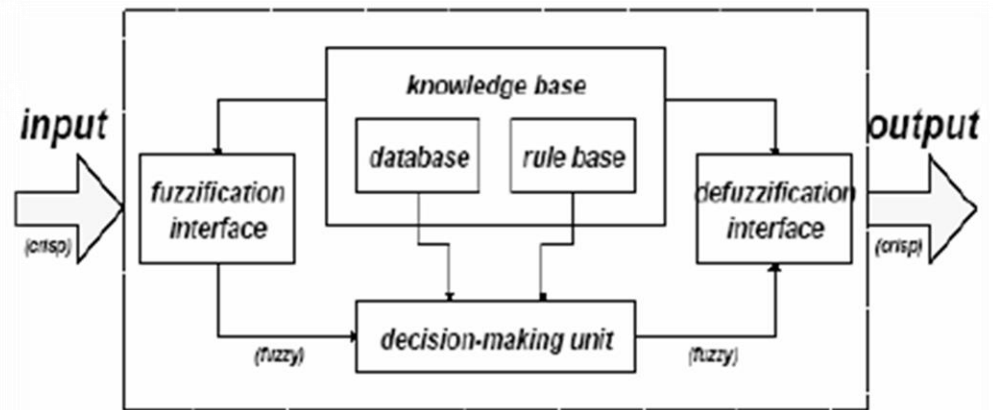
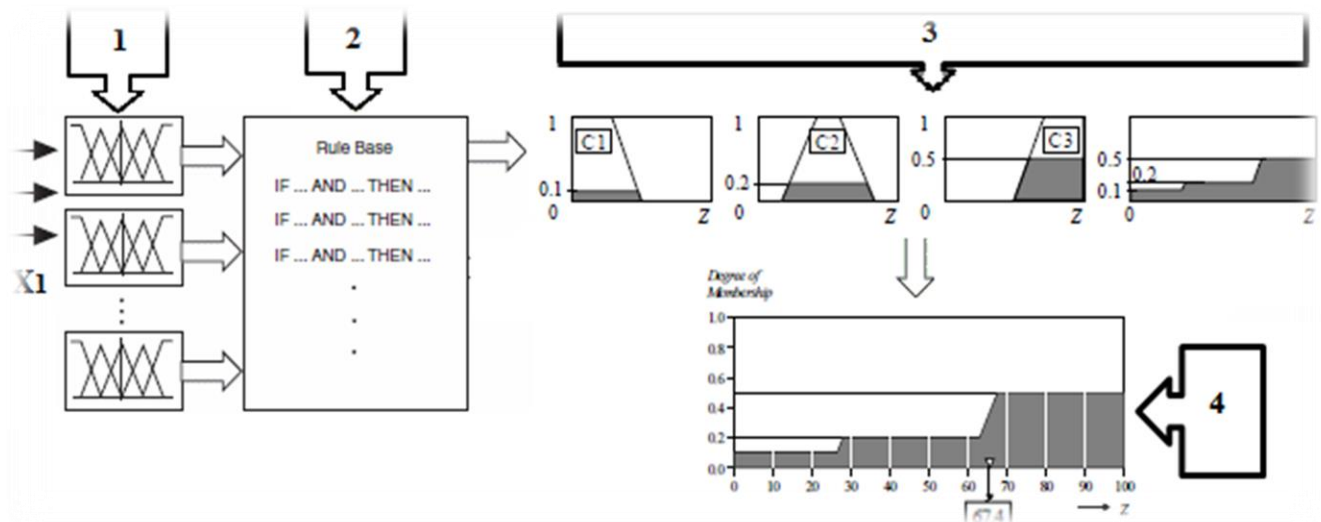


Fig. 4. Fuzzy Logic system structure.




Applications and examples of AI

- Deep Blue defeated the world chess champion Garry Kasparov in 1997.
- In 1997, the Deep Blue chess program created by IBM, beat the current world chess champion, Gary Kasparov.




What is Deep Blue ?



IBM Deep Blue

Vs



World chess champion
Garry Kasparov in 1997

A blue-themed graphic titled "What is Deep Blue ?" comparing the IBM Deep Blue computer system and the world chess champion Garry Kasparov. On the left is a photograph of the tall, black IBM Deep Blue server rack. On the right is a photograph of Garry Kasparov standing with his arms crossed. The word "Vs" is centered between the two images.

Deep Blue

Play :V2 Deep Blue

Autonomous Robotic ground vehicle

- Defense Advanced Research Projects Agency (DARPA) Grand Challenge
- Cash prizes (\$1 to \$2 million) offered to first robots to complete a long course completely unassisted.
- Stimulates research in vision, robotics, planning, machine learning.

The DARPA Grand Challenge

Autonomous Robotic Ground Vehicles
Los Angeles – Las Vegas
March 13, 2004
www.darpa.mil/grandchallenge



The Challenge

- Navigate 300 miles of rugged terrain between Los Angeles and Las Vegas
- Winner of \$1 million cash prize is first to complete course in prescribed time
- No drivers allowed – unmanned vehicles only

DARPA

- 2004 Grand Challenge:
 - 150 mile route in Nevada desert
 - Furthest any robot went was about 7 miles

- 2005 Grand Challenge:
 - 132 mile race
 - Narrow tunnels, winding mountain passes, etc
 - Stanford 1st, CMU 2nd, both finished in about 6 hours

- 2007 Urban Grand Challenge

Stanley Autonomous Robotic ground vehicle winner of DARPA 2005



Stanley Robot

Stanford Racing Team www.stanfordracing.org

Play :V3 Stanley Autonomous Robotic ground
vehicle winner of DARPA

Face recognition

Searching.....

Found

You (and 15 other people) say this is
Henrik Andersen (8)
Also known as... (4 alternatives)
Possible matches (opens in a new window)

Record Navigation Wizard

Google Search results

Google

Web Images Groups News

Hours Search

Search

the web pages from the I

Web Results 1 - 10 of about 446,000

Personalization Engine

News results for **spurs** - View today's top stories

Duncan returns as **Spurs** top pick - Arizona Republic - 9 11

Duncan O-moralizing to **Spurs** opponents - USA Today - 7

Tottenham **interview**

The official site with news, team, players, fixtures, match reports, travel, members club, history, stadium, and links.

www.spurs.co.uk/ - 5k - 7 Mar 2006 - Cached - Similar pages

OK Cancel

Google

Dr. Mohammed Abu-mallouh - AI 39


Aviation

DOT/FAA/CT-94/41

FAA Technical Center
Atlantic City International Airport,
N.J. 08405

Artificial Intelligence With Applications for Aircraft

19941129 049



- Optimizing the use of airspace.
- Reducing the cost of flying.
- Meeting Air Traffic Control(ATC) requirements.
- Aiding the decision making process of the flight crew.
- Aiding maintenance activity.
- Assisting data management.

Fuzzy Logic in Automotive Engineering

- Antilock Braking System (ABS) -Nissan and Mitsubishi.
- Engine Control-Nok and Nissan.
- Automatic transmission systems- Nissan, Honda, GM.
- Cruise control – Peugeot, Citroën.

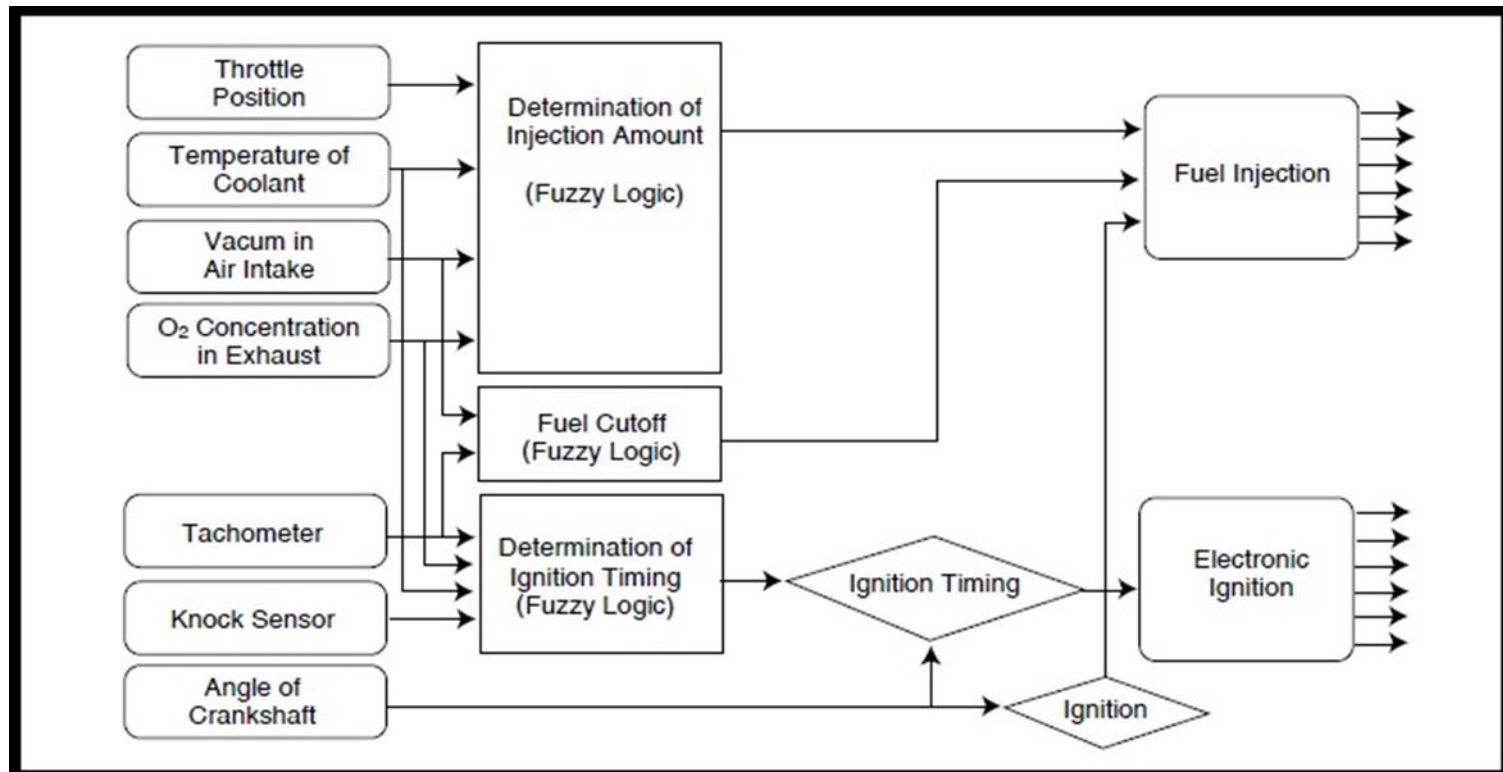
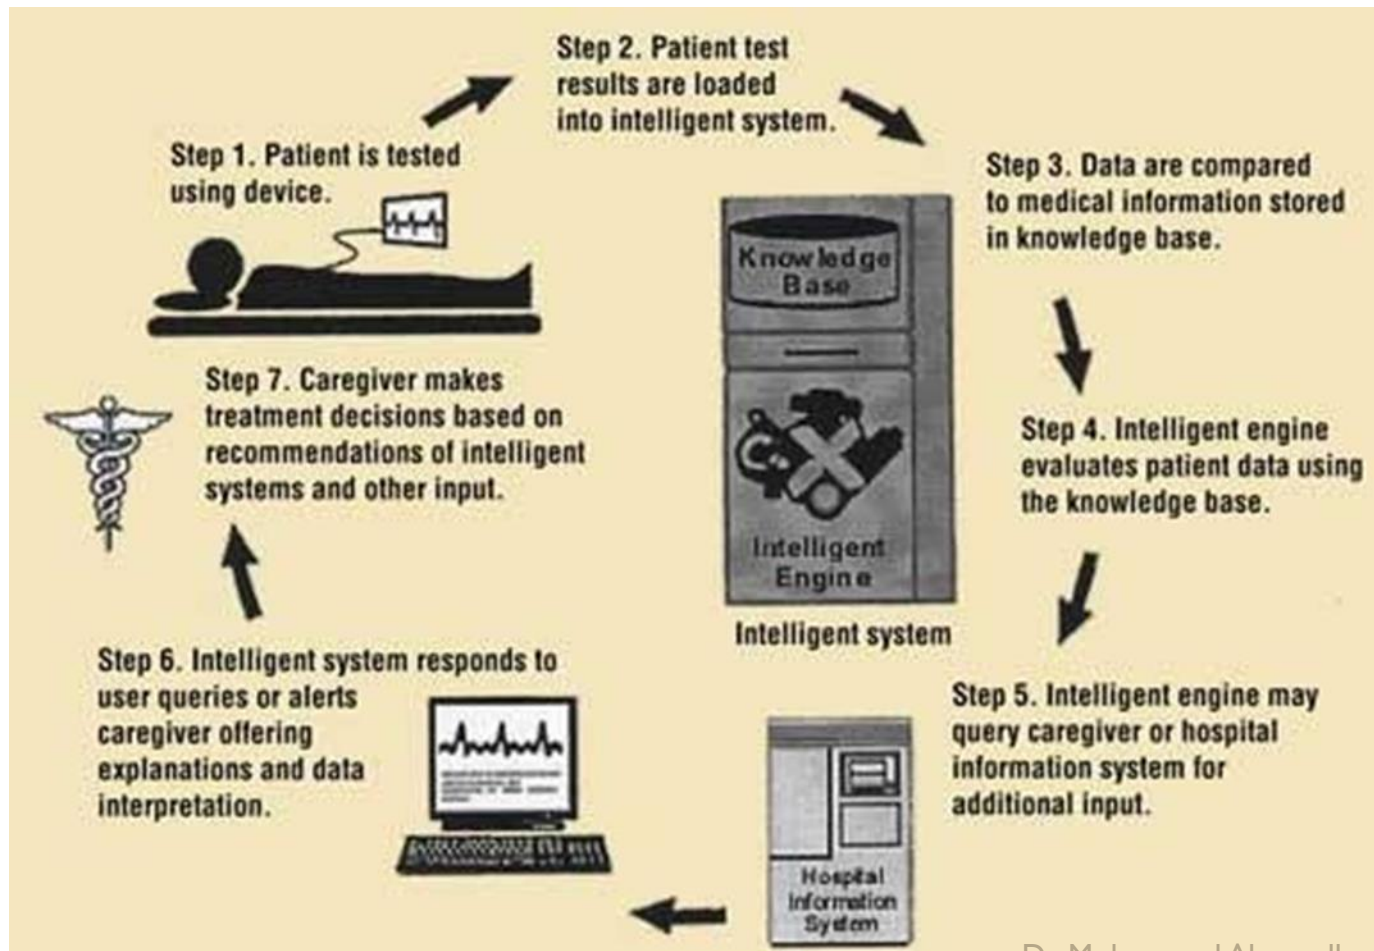


Figure 2—As you can see, the engine controller of NOK Corporation contains three fuzzy-logic modules.


AI In Medicine

- MYCIN: early expert system that used artificial intelligence to identify bacteria causing severe infections.



AI In Robotics

- Tennis playing robot.



AI In IBM Watson

- IBM Watson is a technology platform that uses natural language processing and machine learning to reveal insights from large amounts of unstructured data.



AI In IBM Watson

- Question Answering.
- Jeopardy! game player.

Play : V4 IBM Watson

Contemporary Issue In AI

Summary: you should be . . .
either shocked or be assured that
There's no magic in AI.
It's all about optimization,
probability and statistics,
logic, algorithms.
have a rough idea of the state-of-the-art of AI

References

1. Introduction to Artificial Intelligence, CS 271, Instructor: Professor Padhraic Smyth
2. <http://www.ibm.com/smarterplanet/us/en/ibmwatson/what-is-watson.html>
3. <https://www.youtube.com/watch?v=qOli7-Qx00k>
4. Introduction to Artificial Intelligence, CS540-2, Bryan R. Gibson
5. DR. A. F. ADEKOYA, Department of Computer Science, University of Agriculture, Abeokuta, Nigeria. www.unaab.edu.ng
6. <https://courses.csail.mit.edu/6.825/>
7. Negnevitsky, Pearson Education, 2011
8. <http://www.ncl.ac.uk/eece/staff/profile/shady.gadoue>

Lecture 4

- **Fuzzy expert systems:
Fuzzy logic**

- Introduction, or what is fuzzy thinking?

- Fuzzy sets

- Linguistic variables and hedges


- Operations of fuzzy sets


- Fuzzy rules

- Summary

Introduction, or what is fuzzy thinking?

- Experts rely on **common sense** when they solve problems.
- **How can we represent expert knowledge that uses vague and ambiguous terms in a computer?**
- Fuzzy logic is based on the idea that all things admit of degrees. Temperature, height, speed, distance- all come on a sliding scale.

- 
- Boolean logic uses sharp distinctions. It forces us to draw lines between members of a class and non-members. For instance, we may say, Tom is tall because his height is 181 cm. If we drew a line at 180 cm, we would find that David, who is 179 cm, is small. Is David really a small man or we have just drawn an arbitrary line in the sand?
 - Fuzzy logic reflects how people think. It attempts to model our sense of words, our decision making and our common sense. As a result, it is leading to new, more human, intelligent systems.



- In 1965 **Lotfi Zadeh**, published his famous paper “Fuzzy set Zadeh extended the work on possibility theory into a formal system of mathematical logic, and introduced a new concept for applying natural language terms.


This new logic for representing and manipulating fuzzy terms was called **fuzzy logic**, and Zadeh became the Master of fuzzy logic.

■ **Why fuzzy?**


As Zadeh said, the term is concrete, immediate and descriptive; we all know what it means. However, many people in the West were repelled by the word fuzzy, because it is usually used in a negative sense.

■ **Why logic?**

Fuzziness rests on fuzzy set theory, and fuzzy logic is just a small part of that theory.

- 
- ❖ Fuzzy logic is a set of mathematical principles for knowledge representation based on degrees of membership.
 - ❖ Unlike two-valued Boolean logic, fuzzy logic is **multi-valued**. It deals with **degrees of membership** and **degrees of truth**.
 - ❖ Fuzzy logic uses the continuum of logical values between 0 (completely false) and 1 (completely true). Instead of just black and white, it employs the spectrum of colors, accepting that things can be partly true and partly false at the same time.

Range of logical values in Boolean and fuzzy logic




Fuzzy sets

- The classical example in fuzzy sets is tall men. The elements of the fuzzy set “tall men” are all men, but their degrees of membership depend on their height.

Name	Height, cm	Degree of Membership	
		<i>Crisp</i>	<i>Fuzzy</i>
Chris	208	1	1.00
Mark	205	1	1.00
John	198	1	0.98
Tom	181	1	0.82
David	179	0	0.78
Mike	172	0	0.24
Bob	167	0	0.15
Steven	158	0	0.06
Bill	155	0	0.01
Peter	152	0	0.00


Degree of Membership

Crisp Sets




Degree of Membership

Fuzzy Sets



Crisp and fuzzy sets of “tall men”



■ The x-axis represents **the universe of discourse** - the range of all possible values applicable to a chosen variable. In our case, the variable is the man height. According to this representation, the universe of men's heights consists of all tall men.

■ The y-axis represents **the membership value of the fuzzy set**. In our case, the fuzzy set of “tall men” maps height values into corresponding membership values.

A fuzzy set is a set with fuzzy boundaries.

■ Let X be the universe of discourse and its elements be denoted as x . In the classical set theory, **crisp set A of X** is defined as function $F_A(x)$ called the characteristic function of A

$$F_A(x) : X \rightarrow \{0,1\}, \text{ where } F_A(x) = \begin{cases} 1, & \text{if } x \in A \\ 0, & \text{if } x \notin A \end{cases}$$

This set maps universe X to a set of two elements. For any element x of universe X , characteristic function $F_A(x)$ is equal to 1 if x is an element of set A , and is equal to 0 if x is not an element of A .

In the fuzzy theory, fuzzy set A of universe X is defined by function $\mu_A(x)$: called the membership function of set A .

$\mu_A(x) : X \rightarrow [0, 1]$, where $\mu_A(x) = 1$ if x is totally in A;
 $\mu_A(x) = 0$ if x is not in A;
 $0 < \mu_A(x) < 1$ if x is partly in A.

This set allows a continuum of possible choices. For any element x of universe X, membership function $\mu_A(x)$ equals the degree to which x is an element of fuzzy set A. This degree, a value between 0 and 1, represents the **degree of membership**, also called **membership** value, of element x in set A.

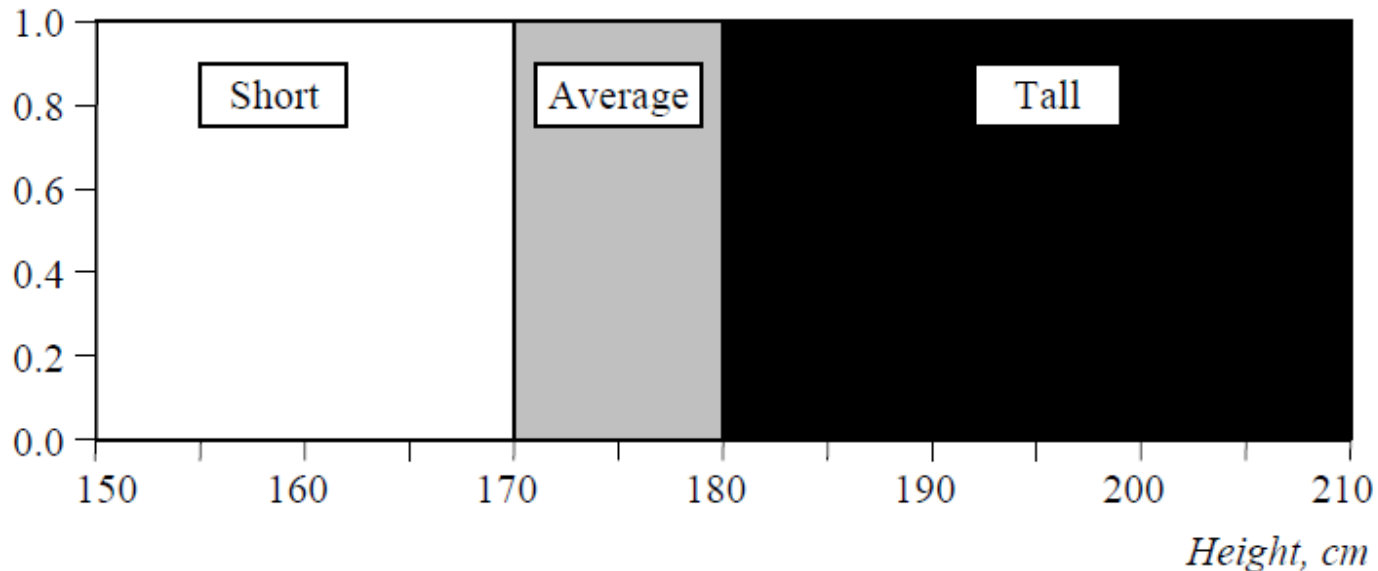
How to represent a fuzzy set in a computer ?

- First, we determine the membership functions. In our “tall men” example, we can obtain fuzzy sets of tall, short and average men.
- The universe of discourse - the men’s heights - consists of three sets: short, average and tall men. As you will see, a man who is 184 cm tall is a member of the average men set with a degree of membership of 0.1, and at the same time, he is also a member of the tall men set with a degree of 0.4.

Crisp and fuzzy sets of short, average and tall men


Degree of Membership

Crisp Sets



Degree of Membership

Fuzzy Sets



Linguistic variables and hedges


- At the root of fuzzy set theory lies the idea of linguistic variables.
- **A linguistic variable is a fuzzy variable.** For example, the statement “John is tall” implies that the linguistic variable John takes the linguistic value tall.

In fuzzy expert systems, linguistic variables are used in fuzzy rules. For example:

IF wind is strong
THEN sailing is good

IF project duration is long
THEN completion risk is high

IF speed is slow
THEN stopping distance is short



➤ The range of possible values of a linguistic variable represents the universe of discourse of that variable. For example, the universe of discourse of the linguistic variable speed might have the range between 0 and 220 km/h and may include such fuzzy subsets as very slow, slow, medium, fast, and very fast.

■ A linguistic variable carries with it the concept of fuzzy set qualifiers, called **hedges**.

■ **Hedges** are terms that modify the shape of fuzzy sets. They include adverbs such as very, somewhat, quite, more or less and slightly.

Fuzzy rules

- In 1973, Lotfi Zadeh published his second most influential paper. This paper outlined a new approach to analysis of complex systems, in which Zadeh suggested capturing human knowledge in fuzzy rules.

What is a fuzzy rule?

- A fuzzy rule can be defined as a conditional statement in the form:

IF x is A
THEN y is B

- where x and y are linguistic variables; and A and B are linguistic values determined by fuzzy sets on the universe of discourses X and Y , respectively.

What is the difference between classical and fuzzy rules?

A classical IF-THEN rule uses binary logic, for example,

Rule: 1

IF speed is > 100

THEN stopping distance is long

Rule: 2

IF speed is < 40

THEN stopping distance is short

The variable speed can have any numerical value between 0 and 220 km/h, but the linguistic variable stopping_distance can take either value long or short. In other words, classical rules are expressed in the black-and-white language of Boolean logic.

We can also represent the stopping distance rules in a fuzzy form:

Rule: 1

IF speed is fast


THEN stopping distance is long

Rule: 2

IF speed is slow

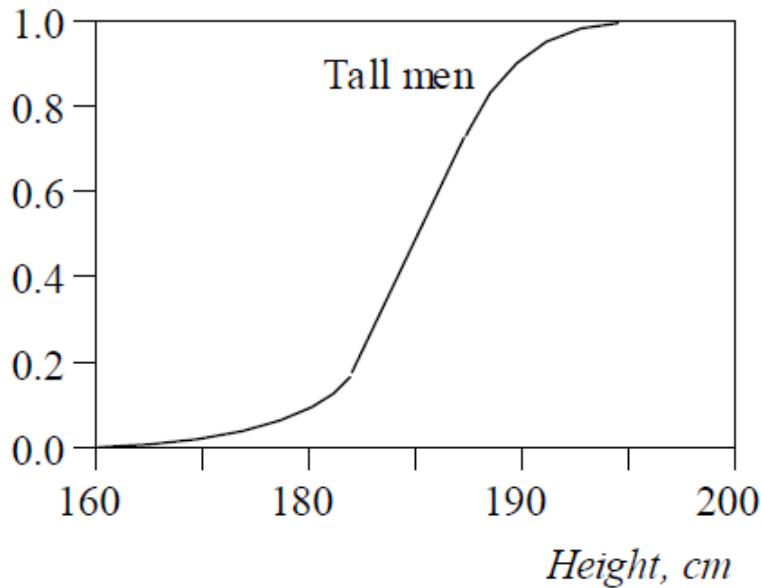
THEN stopping distance is short

- In fuzzy rules, the linguistic variable speed also has the range (the universe of discourse) between 0 and 220 km/h, but this range includes fuzzy sets, such as slow, medium and fast. The universe of discourse of the linguistic variable stopping_distance can be between 0 and 300 m and may include such fuzzy sets as short, medium and long.

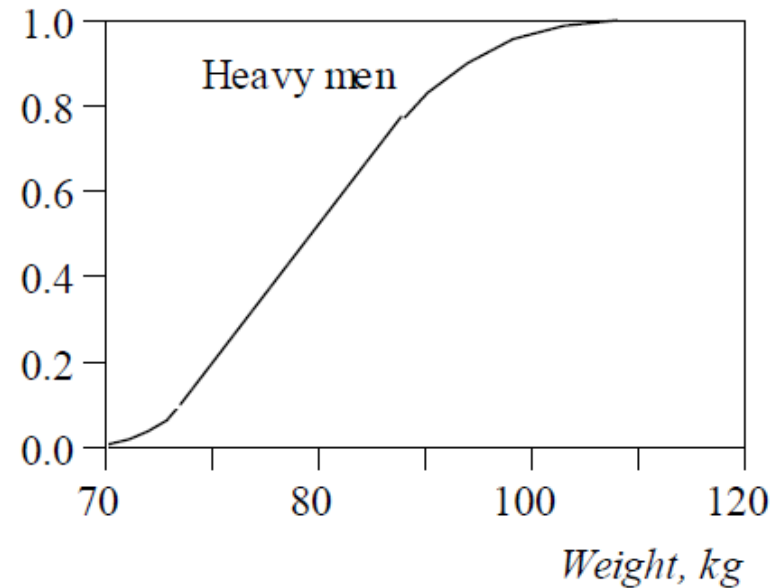
- 
- Fuzzy rules relate fuzzy sets.
 - In a fuzzy system, all rules fire to some extent, or in other words they fire partially. If the antecedent is true to some degree of membership, then the consequent is also true to that same degree.

Fuzzy sets of tall and heavy men

Degree of Membership





Degree of Membership



These fuzzy sets provide the basis for a weight estimation model. The model is based on a relationship between a man's height and his weight:

IF height is tall
THEN weight is heavy

- The value of the output or a truth membership grade of the rule consequent can be estimated directly from a corresponding truth membership grade in the antecedent. This form of fuzzy inference uses a method called **monotonic selection**.






A fuzzy rule can have multiple antecedents, for example:

IF project_duration is long
AND project_staffing is large
AND project_funding is inadequate
THEN risk is high

IF service is excellent
OR food is delicious
THEN tip is generous



The consequent of a fuzzy rule can also include multiple parts, for instance:

IF temperature is hot

THEN hot water is reduced;
 cold-water is increased

Lecture 5

Fuzzy expert systems:

Fuzzy inference

- Mamdani fuzzy inference
- Sugeno fuzzy inference
- Case study
- Summary

Fuzzy inference

The most commonly used fuzzy inference technique is the so-called Mamdani method. In 1975, Professor **Ebrahim Mamdani** of London University built one of the first fuzzy systems to control a steam engine and boiler combination. He applied a set of fuzzy rules supplied by experienced human operators.

Mamdani fuzzy Inference

- The Mamdani-style fuzzy inference process is performed in four steps:
 - fuzzification of the input variables,
 - rule evaluation;
 - aggregation of the rule outputs, and finally
 - defuzzification.

We examine a simple two-input one-output problem that includes three rules:

Rule: 1

IF x is A3
OR y is B1
THEN z is C1

Rule: 2

IF x is A2
AND y is B2
THEN z is C2

Rule: 3

IF x is A1
THEN z is C3

Rule: 1

IF *project_funding* is *adequate*
OR *project_staffing* is *small*
THEN *risk* is *low*

Rule: 2

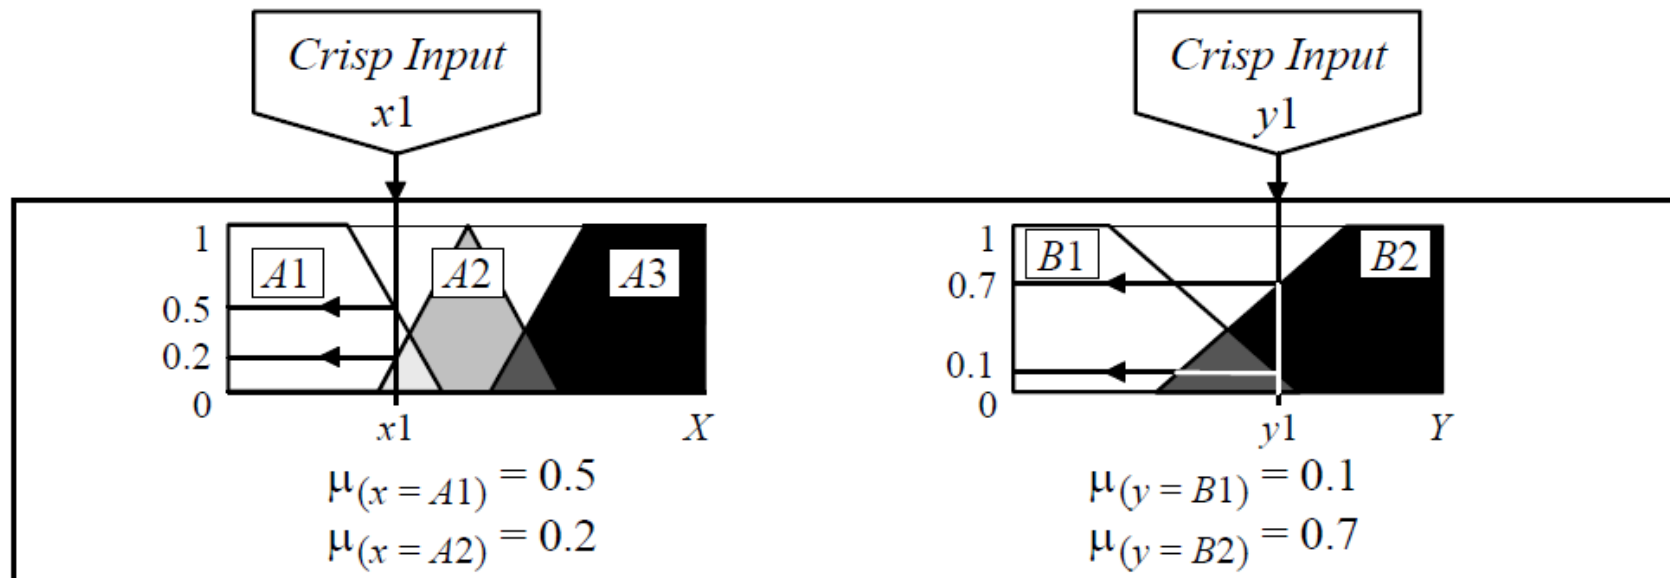
IF *project_funding* is *marginal*
AND *project_staffing* is *large*
THEN *risk* is *normal*

Rule: 3

IF *project_funding* is *inadequate*
THEN *risk* is *high*

Step I: Fuzzification

The first step is to take the crisp inputs, x_1 and y_1 (project funding and project staffing), and determine the degree to which these inputs belong to each of the appropriate fuzzy sets.



Step 2: Rule Evaluation

The second step is to take the fuzzified inputs, $\mu(x=A1) = 0.5$, $\mu(x=A2) = 0.2$, $\mu(y=B1) = 0.1$ and $\mu(y=B2) = 0.7$, and apply them to the antecedents of the fuzzy rules. If a given fuzzy rule has multiple antecedents, the fuzzy operator (AND or OR) is used to obtain a single number that represents the result of the antecedent evaluation. This number (the truth value) is then applied to the consequent membership function.

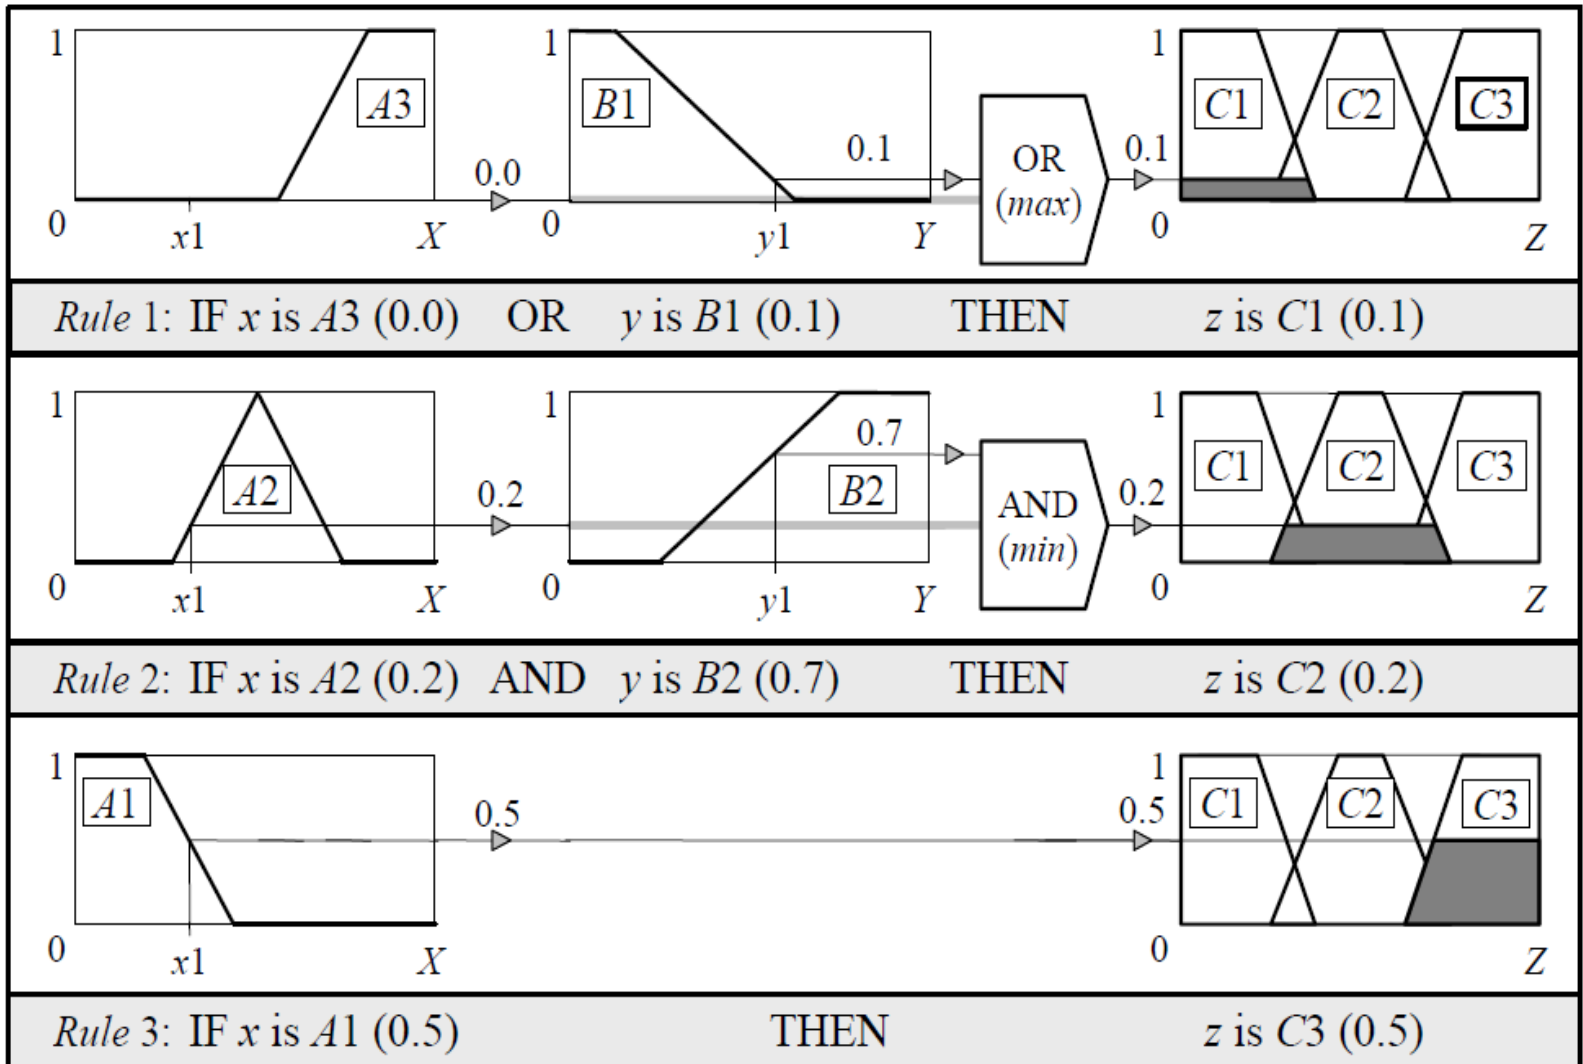
- To evaluate the disjunction of the rule antecedents, we use the **OR fuzzy operation**. Typically, fuzzy expert systems make use of the classical fuzzy operation **union**:

$$\mu_{A \cup B}(x) = \max[\mu_A(x), \mu_B(X)]$$

- Similarly, in order to evaluate the conjunction of the rule antecedents, we apply the **AND fuzzy operation intersection**:

$$\mu_{A \cap B}(x) = \min[\mu_A(x), \mu_B(X)]$$


Mamdani-style rule evaluation





Now the result of the antecedent evaluation can be applied to the membership function of the consequent.


■ The most common method of correlating the rule consequent with the truth value of the rule antecedent is to cut the consequent membership function at the level of the antecedent truth. This method is called clipping. Since the top of the membership function is sliced, the clipped fuzzy set loses some information. However, **clipping** is still often preferred because it involves less complex and faster mathematics, and generates an aggregated output surface that is easier to defuzzify.




While clipping is a frequently used method, **scaling** offers a better approach for preserving the original shape of the fuzzy set. The original membership function of the rule consequent is adjusted by multiplying all its membership degrees by the truth value of the rule antecedent. This method, which generally loses less information, can be very useful in fuzzy expert systems.

Clipped and scaled membership functions

Degree of Membership




Degree of Membership



Step 3: Aggregation of the rule outputs

Aggregation is the process of unification of the outputs of all rules. We take the membership functions of all rule consequents previously clipped or scaled and combine them into a single fuzzy set. The input of the aggregation process is the list of clipped or scaled consequent membership functions, and the output is one fuzzy set for each output variable.

Aggregation of the rule outputs




Step 4: Defuzzification

The last step in the fuzzy inference process is defuzzification. Fuzziness helps us to evaluate the rules, but the final output of a fuzzy system has to be a crisp number. The input for the defuzzification process is the aggregate output fuzzy set and the output is a single number.

- ❖ There are several defuzzification methods, but probably the most popular one is the **centroid technique**. It finds the point where a vertical line would slice the aggregate set into two equal masses. Mathematically this **centre of gravity (COG)** can be expressed as:


$$COG = \frac{\int_a^b \mu_A(x) x dx}{\int_a^b \mu_A(x) dx}$$

- Centroid defuzzification method finds a point representing the centre of gravity of the fuzzy set, A , on the interval, ab .
- A reasonable estimate can be obtained by calculating it over a sample of points.



Centre of gravity (COG):

$$COG = \frac{[(5+15+25) \times 0.1 + (35+45+55+65) \times 0.2 + (75+85+95) \times 0.5] (10)}{[0.1+0.1+0.1+0.2+0.2+0.2+0.2+0.5+0.5+0.5] (10)} = 71.7$$



Sugeno fuzzy inference

- Mamdani-style inference, as we have just seen, requires us to find the centroid of a two-dimensional shape by integrating across a continuously varying function. In general, this process is not computationally efficient.
- **Michio Sugeno** suggested to use a single spike, a singleton, as the membership function of the rule consequent. A **singleton**, or more precisely a **fuzzy singleton**, is a fuzzy set with a membership function that is unity at a single particular point on the universe of discourse and zero everywhere else.

Sugeno-style fuzzy inference is very similar to the Mamdani method. Sugeno changed only a rule consequent. Instead of a fuzzy set, he used a mathematical function of the input variable. The format of the **Sugeno-style fuzzy rule** is :

IF x is A
AND y is B
THEN z is $f(x, y)$

where x , y and z are linguistic variables; A and B are fuzzy sets on universe of discourses X and Y , respectively; and $f(x, y)$ is a mathematical function.


- The most commonly used **zero-order Sugeno fuzzy model** applies fuzzy rules in the following form:

IF x is A
AND y is B
THEN z is k


where k is a constant.

- In this case, the output of each fuzzy rule is constant. All consequent membership functions are represented by singleton spikes.

Sugeno-style rule evaluation




Sugeno-style aggregation of the rule outputs



Weighted average (WA):

$$WA = \frac{\mu(k1) \times k1 + \mu(k2) \times k2 + \mu(k3) \times k3}{\mu(k1) + \mu(k2) + \mu(k3)} = \frac{0.1 \times 20 + 0.2 \times 50 + 0.5 \times 80}{0.1 + 0.2 + 0.5} = 65$$

Sugeno-style defuzzification




How to make a decision on which method to apply - Mamdani or Sugeno?

- Mamdani method is widely accepted for capturing expert knowledge. It allows us to describe the expertise in more intuitive, more human-like manner. However, Mamdani-type fuzzy inference entails a substantial computational burden.
- On the other hand, Sugeno method is computationally effective and works well with optimization and adaptive techniques, which makes it very attractive in control problems, particularly for dynamic nonlinear systems.

Process of developing a fuzzy expert system

1. Specify the problem and define linguistic variables.
2. Determine fuzzy sets.
3. Elicit and construct fuzzy rules.
4. Encode the fuzzy sets, fuzzy rules and procedures to perform fuzzy inference into the expert system.
5. Evaluate and tune the system.




Artificial neural networks: Supervised learning


- Introduction, or how the brain works
- The neuron as a simple computing element
- The perceptron
- Multilayer neural networks
- Accelerated learning in multilayer neural networks
- Summary

Introduction, or how the brain works

- Machine learning involves adaptive mechanisms that enable computers to learn from experience, learn by example and learn by analogy.
- Learning capabilities can improve the performance of an intelligent system over time. The most popular approaches to machine learning are **artificial neural networks** and **genetic algorithms**.


- 
- A **neural network** can be defined as a model of reasoning based on the human brain.
 - The brain consists of a densely interconnected set of nerve cells, or basic information-processing units, called **neurons**.
 - The human brain incorporates nearly 10 billion neurons and 60 trillion connections, **synapses**, between them.
 - By using multiple neurons simultaneously, the brain can perform its functions much faster than the fastest computers in existence.
 - Each neuron has a very simple structure, but an army of such elements constitutes a tremendous processing power.
 - A neuron consists of a cell body, **soma**, a number of fibers called **dendrites**, and a single long fiber called the **axon**.

Biological neural network



- Our brain can be considered as a highly complex, non-linear and parallel information-processing system.
- ❑ Information is stored and processed in a neural network simultaneously throughout the whole network, rather than at specific locations. In other words, in neural networks, both data and its processing are **global** rather than local.
- ❑ Learning is a fundamental and essential characteristic of biological neural networks. The ease with which they can learn led to attempts to emulate a biological neural network in a computer.
- ❑ An artificial neural network consists of a number of very simple processors, also called neurons, which are analogous to the biological neurons in the brain.
- ❑ The neurons are connected by weighted links passing signals from one neuron to another.
- ❑ The output signal is transmitted through the neuron's outgoing connection. The outgoing connection splits into a number of branches that transmit the same signal. The outgoing branches terminate at the incoming connections of other neurons in the network.

Architecture of a typical artificial neural network

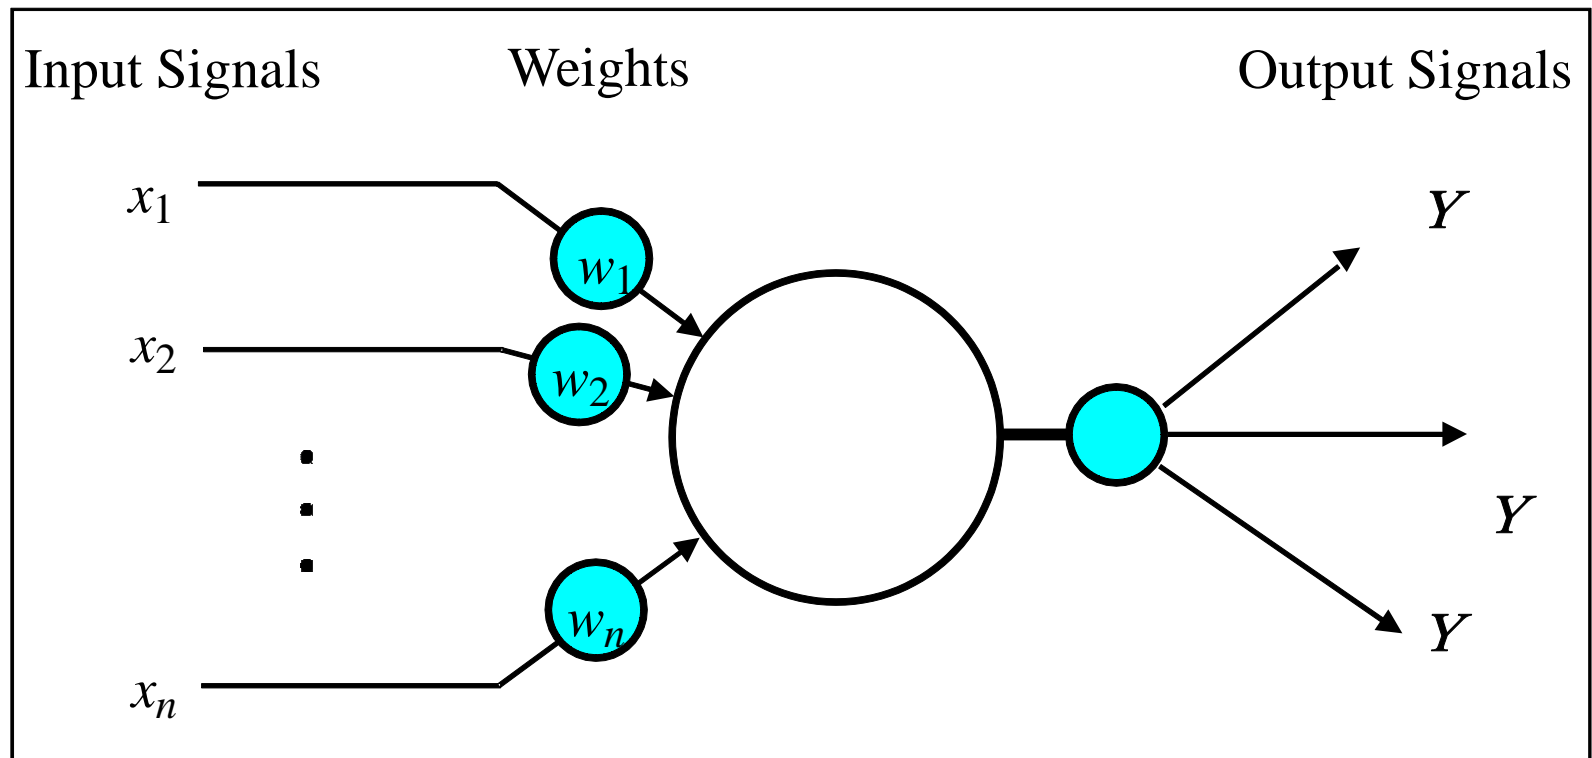


Analogy between biological and artificial neural networks

<i>Biological Neural Network</i>	<i>Artificial Neural Network</i>
<ul style="list-style-type: none">• Soma• Dendrite• Axon• Synapse	<ul style="list-style-type: none">• Neuron• Input• Output• Weight

The neuron as a simple computing element

Diagram of a neuron



□ The neuron computes the weighted sum of the input signals and compares the result with a **threshold value**, θ .

If the net input is less than the threshold, the neuron output is 0, But if the net input is greater than or equal to the threshold, the neuron becomes activated and its output attains a value +1.

The neuron uses the following transfer or **activation function**:


$$X = \sum_{i=1}^n x_i w_i$$

$$Y = \begin{cases} +1, & \text{if } X \geq \theta \\ 0, & \text{if } X < \theta \end{cases}$$

□ This type of activation function is called a **step function**.


Activation functions of a neuron

Step function




$$Y^{step} = \begin{cases} 1, & \text{if } X \geq 0 \\ 0, & \text{if } X < 0 \end{cases}$$

Sign function




$$Y^{sign} = \begin{cases} +1, & \text{if } X \geq 0 \\ -1, & \text{if } X < 0 \end{cases}$$

Sigmoid function



$$Y^{sigmoid} = \frac{1}{1 + e^{-X}}$$

Linear function



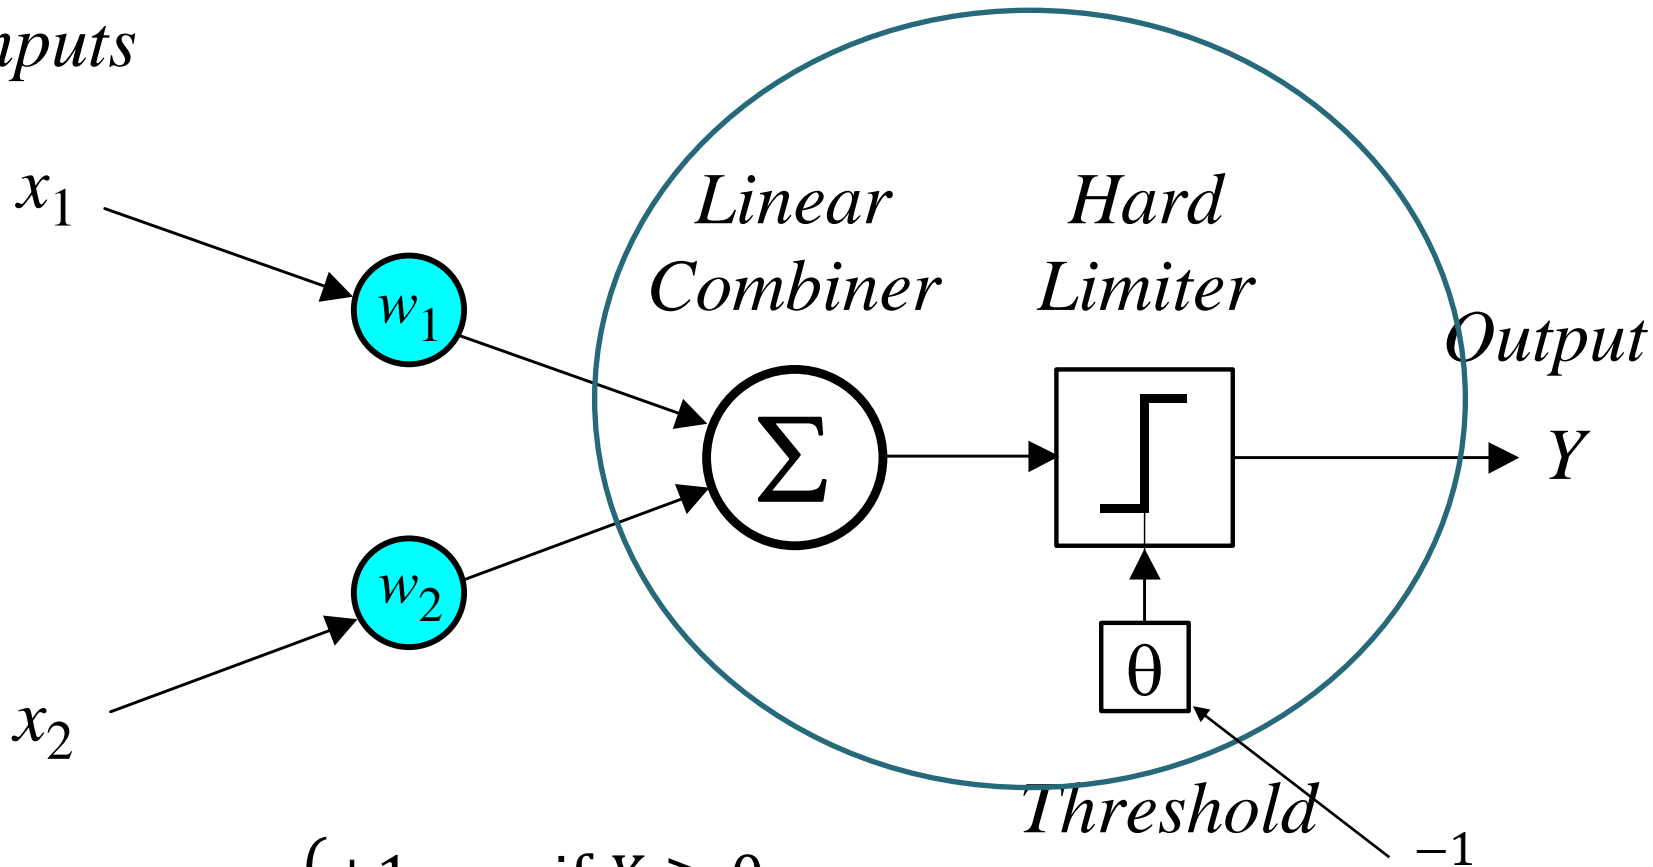
$$Y^{linear} = X$$

Can a single neuron learn a task?

- In 1958, **Frank Rosenblatt** introduced a training algorithm that provided the first procedure for training a simple ANN: a **perceptron**.
- The perceptron is the simplest form of a neural network. It consists of a single neuron with **adjustable** synaptic weights and a **hard limiter**.

Single-layer two-input perceptron

Inputs



$$Y = \begin{cases} +1, & \text{if } X \geq \theta \\ -1, & \text{if } X < \theta \end{cases}$$

$$X = \sum_{i=1}^n x_i w_i$$

The Perceptron


- The weighted sum of the inputs is applied to the hard limiter, which produces an output equal to +1 if its input is positive and -1 if it is negative.
- The aim of the perceptron is to classify inputs, x_1, x_2, \dots, x_n , into one of two classes, say A_1 and A_2 .
- In the case of an elementary perceptron, the n - dimensional space is divided by a hyper plane into two decision regions. The **hyper plane** is defined by the **linearly separable function**:

$$\sum_{i=1}^n x_i w_i - \theta = 0$$

Linear Separability

Consider a perceptron processor that only has 2 input . So ($\theta = w_0$)

$$a = -w_0 + w_1x_1 + w_2x_2$$



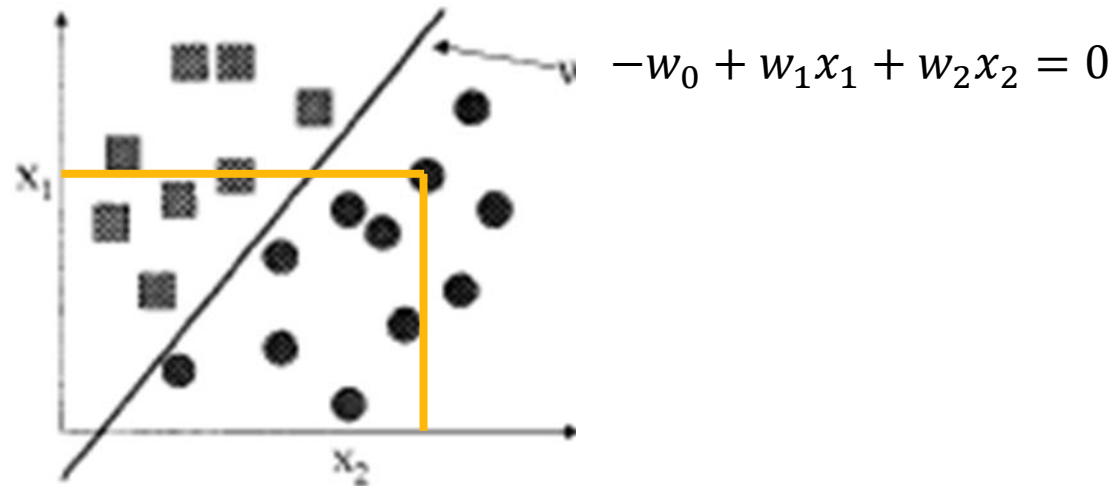
Which is just

Now let's look at the locus of points for which


$$w_0(-1) + w_1x_1 + w_2x_2 = 0$$

When plotted in the space of the values of x_1 and x_2 this is just the equation of a straight line .

$$(x_2 = ax_1 + b, \text{ where } a = w_1 / w_2 \text{ and } b = w_0 / w_2)$$




- For inputs x that fall on one side line , the activation (a) will be >0 , and thus the output (y) will be 1 . On the other side (and right on the line) the output will be -1 .
- So each set of values for w defines a straight line decision boundary and any possible line can be represented by some value of w .




$$-w_0 + w_1x_1 + w_2x_2 = 0$$

- So obviously , if we had a problem for which the examples did not separate nicely, we could not solve the problem completely using a perceptron processor.
- If the set of input/output pairs can be separated with a straight line , then it has the property of being linearly separable.
- (For input vectors of higher dimension, the boundary is a higher dimension hyper plane.Eg., a plane in a 3D input space)


Linear separability in the perceptron's



(a) Two-input perceptron.




(b) Three-input perceptron.



➤ It is important to distinguish between the ability to perform perfectly on the set of examples, and having a perceptron that is able to continue to perform for examples that were not used during the training of the perceptron's weights. That is an issue of how well the examples predict the subsequent inputs.


- Consider the XOR Problem : To form a decision boundary for the XOR function.




- This task is not linearly separable. It would require 2 decision boundaries to separate the two classes.

➤ This task can be solved with 2 perceptron processors (one for each decision boundary), and a third that resolves their outputs. There are many possible solutions in terms of the values of the weights and thresholds.


In general, networks of perceptron-like processors can solve most non-linearly separable tasks by using more than one layer of processors. Rosenblatt (and others) realized this in the 1960's, but did not have a learning rule that would work effectively with more than one level (it wasn't invented until the mid 1970's).




❖ What if we have more a single output ?



❖ Then it is really like each processor operating on its own with the same inputs.





❖ Can perceptrons be used for tasks that are not linearly separable ?

With modifications, yes.

- If the task is not linearly separable, we accept that the result will not be perfect, but we wish to minimize the number of errors.
- The number of errors is defined by the values of the weights.

How does the perceptron learn its classification tasks?

This is done by making small adjustments in the weights to reduce the difference between the actual and desired outputs of the perceptron. The initial weights are randomly assigned, usually in the range $[-0.5, 0.5]$, and then updated to obtain the output consistent with the training examples.

- If at iteration p , the actual output is $Y(p)$ and the desired output is $Y_d(p)$, then the error is given by:

$$e(p) = Y_d(p) - Y(p)$$

where $p = 1, 2, 3, \dots$

- Iteration p here refers to the p th training example presented to the perceptron.
- If the error, $e(p)$, is positive, we need to increase perceptron output $Y(p)$, but if it is negative, we need to decrease $Y(p)$.

The perceptron learning rule

$$w_i(p + 1) = w_i(p) + \alpha \cdot x_i(p) \cdot e(p)$$

where $p = 1, 2, 3, \dots$

α is the **learning rate**, a positive constant less than unity.

- A learning rule is a strategy by which example input/output pairs can be used to incrementally change the weights in away that gradually improves the performance of the network.

The perceptron learning rule involves the example input x , the computed output y , and the desired output d .

$$\text{If } y = 1, \text{ and } d = 0: w_i \leftarrow w_i - \alpha x_i (w_i = 1, \dots, n)$$

$$\text{If } y = 0, \text{ and } d = 1: w_i \leftarrow w_i + \alpha x_i (w_i = 1, \dots, n)$$

Where α is a small learning parameter

- Whenever the $(y=1, \text{ and } d=0)$ error occurs, it is because the activation a was too large, and decrementing the weights will reduce the activation for the same input. Similarly, the $(y=0, \text{ and } d=1)$ will result in appropriate increasing of activation for the same input .

Perceptron's training algorithm

Step 1: Initialisation

Set initial weights w_1, w_2, \dots, w_n and threshold θ to random numbers in the range $[-0.5, 0.5]$.

Perceptron's training algorithm (continued)

Step 2: Activation

Activate the perceptron by applying inputs $x_1(p), x_2(p), \dots, x_n(p)$ and desired output $Y_d(p)$. Calculate the actual output at iteration $p = 1$

$$Y(p) = \text{step} \left[\sum_{i=1}^n x_i(p) w_i(p) - \theta \right]$$

where n is the number of the perceptron inputs, and step is a step activation function.

Perceptron's training algorithm

If the error, $e(p)$, is positive, we need to increase perceptron output $Y(p)$, but if it is negative, we need to decrease $Y(p)$.

Step 3: Weight training

Update the weights of the perceptron

$$w_i(p + 1) = w_i(p) + \Delta w_i(p)$$

where $\Delta w_i(p)$ is the weight correction at iteration p .

$$\Delta w_i(p) = \alpha \cdot x_i(p) \cdot e(p)$$

The weight correction is computed by the **delta**

Step 4: Iteration

Increase iteration p by one, go back to *Step 2* and repeat the process until convergence.

Example of perceptron learning: the logical operation AND

Epoch	Inputs		Desired output Y_d	Initial weights		Actual output Y	Error e	Final weights	
	x_1	x_2		w_1	w_2			w_1	w_2
1	0	0	0	0.3	-0.1				
	0	1	0						
	1	0	0						
	1	1	1						
2	0	0	0						
	0	1	0						
	1	0	0						
	1	1	1						
3	0	0	0						
	0	1	0						
	1	0	0						
	1	1	1						
4	0	0	0						
	0	1	0						
	1	0	0						
	1	1	1						
5	0	0	0						
	0	1	0						
	1	0	0						
	1	1	1						

Threshold: $\theta = 0.2$; learning rate: $\alpha = 0.1$

Example of perceptron learning: the logical operation *AND*


Epoch	Inputs		Desired output Y_d	Initial weights		Actual output Y	Error e	Final weights	
	x_1	x_2		w_1	w_2			w_1	w_2
1	0	0	0	0.3	-0.1	0	0	0.3	-0.1
	0	1	0	0.3	-0.1	0	0	0.3	-0.1
	1	0	0	0.3	-0.1	1	-1	0.2	-0.1
	1	1	1	0.2	-0.1	0	1	0.3	0.0
2	0	0	0	0.3	0.0	0	0	0.3	0.0
	0	1	0	0.3	0.0	0	0	0.3	0.0
	1	0	0	0.3	0.0	1	-1	0.2	0.0
	1	1	1	0.2	0.0	1	0	0.2	0.0
3	0	0	0	0.2	0.0	0	0	0.2	0.0
	0	1	0	0.2	0.0	0	0	0.2	0.0
	1	0	0	0.2	0.0	1	-1	0.1	0.0
	1	1	1	0.1	0.0	0	1	0.2	0.1
4	0	0	0	0.2	0.1	0	0	0.2	0.1
	0	1	0	0.2	0.1	0	0	0.2	0.1
	1	0	0	0.2	0.1	1	-1	0.1	0.1
	1	1	1	0.1	0.1	1	0	0.1	0.1
5	0	0	0	0.1	0.1	0	0	0.1	0.1
	0	1	0	0.1	0.1	0	0	0.1	0.1
	1	0	0	0.1	0.1	0	0	0.1	0.1
	1	1	1	0.1	0.1	1	0	0.1	0.1

Threshold: $\theta = 0.2$; learning rate: $\alpha = 0.1$

Multilayer neural networks

- A multilayer perceptron is a feed forward neural network with one or more hidden layers.
- The network consists of an **input layer** of source neurons, at least one middle or **hidden layer** of computational neurons, and an **output layer** of computational neurons.
- The input signals are propagated in a forward direction on a layer-by-layer basis.

Multilayer perceptron with two hidden layers




What does the middle layer hide?

- A hidden layer “hides” its desired output. Neurons in the hidden layer cannot be observed through the input/output behaviour of the network. There is no obvious way to know what the desired output of the hidden layer should be.
- Commercial ANNs incorporate three and sometimes four layers, including one or two hidden layers. Each layer can contain from 10 to 1000 neurons. Experimental neural networks may have five or even six layers, including three or four hidden layers, and utilise millions of neurons.

BACKPROPAGATION NETWORKS

To overcome the limitations of perceptrons, networks require more than one processing layer. If we try to extend the error-correction learning (as used in Adeline) to more layers, we encounter the credit-assignment problem: How to determine which nodes are responsible for an outcome.


We consider a back propagation network with 2 processing layers, though they could have more. The layers are output, hidden, and input (which is just fan-in).




Each processing node uses

$$a_j = \sum_{i=0}^n X_i W_{ij}$$

- The activation of node j is the sum of each of its inputs x_0, x_1, \dots, x_n times each of its weights or the corresponding input W_{ij}
- Sometimes we refer to the inputs as y_0, y_1, \dots, y_n because they are often outputs of the nodes $1, \dots, n$.
- Each processing node has a bias (for threshold) input $x_0 = -1$.
- Each processing node creates an output signal, Which is the sigmoidal function.



$$y_j = f(a) = \frac{1}{1 + e^{-a}}$$



The training data consists of examples of inputs and desired outputs : $\{(x_1, d_1), \dots, (x_p, d_p)\}$


In general , we define a measure of error for any particular trial:

$$error = \sum_{i=1}^m e^2(y_k, d_k)$$

Where y_k are the outputs of each of the m output nodes, and d_k are the components of the desired result.

Usually the error function is the sum of square errors , or mean error.

$$e = d - a, \quad E = \frac{\sum e^2}{2}.$$



Consider the projection of a single weight ij in the error surface.

$\partial E / \partial W_{ij}$ is the direction of steepest ascent, and so $-\partial E / \partial W_{ij}$ is the direction of steepest descent.

We are interested in changing the weight to a value that decreases the error, so $-\partial E / \partial W_{ij}$ tells us in what direction to move the weight value.

$$a = \sum_{i=0}^n X_i W_i$$

we use gradient descent to determine the direction of change of the weights w

$$-\partial E / \partial W$$

We consider the adjustment to be made to individual weights w_{ij} , which is the weight to node j from its i th input. Regardless of what layer it is in, we want to make the adjustment:

$$\Delta W_{ij} = -\alpha \frac{\partial E}{\partial w_{ij}}$$

We have: (1) $E = (1/2) \sum e^2$ Total net work error

(2) $e_j = d_j - y_j$ Individual node error

(3) $y_j = f(a_j)$ Output is function of activation

(4) $a_j = \sum w_{ij} y_i$ Value of node's activation is sum of weight inputs

$$(1) E = \frac{1}{2} \sum e^2, \quad (2) e_j = d_j - y_j, \quad (3) y_j = f(a_j), \quad (4) a_j = \sum w_{ij} y_i$$

We calculate the change that we wish to make to each weight in the **output layer**.

$$\Delta w_{ij} = -\alpha \frac{\partial E}{\partial w_{ij}}$$

$$\Delta w_{ij} = -\alpha \frac{\partial E}{\partial e_j} \frac{\partial e_j}{\partial y_j} \frac{\partial y_j}{\partial a_j} \frac{\partial a_j}{\partial w_{ij}}$$

$$\Delta w_{ij} = -\alpha e_j \cdot (-1) \cdot f'(a_j) \cdot y_i$$

$$\Delta w_{ij} = \alpha \delta_j^0 y_i \quad \text{where } \delta_j^0 = e_j f'(a_j)$$

This is a lot like the rule that was used for the Perceptron. Constant \times error \times input.

The difference is that δ_j^0 , the error at output node j , is the usual error $(d_j - y_j)$

, scaled by a factor of $f'(a_j)$

$$f(x) = 1/(1 + e^{-x})$$

$$f'(x) = e^{-x}/(1 + e^{-x})^2$$

$$= 1/(1 + e^{-x}) - 1/(1 + e^{-x})^2$$

$$= f(x)(1 - f(x))$$

$$\text{so } \delta_j^0 = (d_j - y_j) y_j (1 - y_j)$$

•But bear in mind that this is only true if $f(x)$ is the sigmoidal function. Other output functions can also be used, so $f'(a_j)$ is the more general form.

❖Note that the output function for Back propagation must be differentiable.

$$(1) E = \frac{1}{2} \sum e^2, \quad (2) e_j = d_j - y_j, \quad (3) y_j = f(a_j), \quad (4) a_j = \sum w_{ij} y_i$$

Now, let's consider how to calculate Δw_{ij} for the **hidden nodes**.

$$\Delta w_{ij} = -\alpha \partial E / \partial w_{ij}$$


$$= -\alpha \partial E / \partial y_j \cdot \partial y_j / \partial a_j \cdot \partial a_j / \partial w_{ij} \quad \text{then using (3) and (4) as before}$$

$$= -\alpha \partial E / \partial y_j \cdot f'(a_j) \cdot y_i$$

Later we will see why these Outlined expressions are equal

$$= \alpha \sum (\delta_k^0 w_{jk}) \cdot f'(a_j) \cdot y_i$$

$$= \alpha \delta_j^h y_i \quad \text{where} \quad \delta_j^h = \sum (\delta_k^0 w_{jk}) \cdot f'(a_j)$$



The missing step in the derivation:

$$\boxed{\partial E / \partial y_j} = - \sum_k \delta_k^o w_{kj}$$

$$E = \frac{1}{2} \sum_k e_k^2 \quad \text{so}$$

$$\partial E / \partial y_j = \sum_k e_k \partial e_k / \partial y_j \quad (j \neq k, \text{ as it does for output nodes})$$

$$= \sum_k e_k \partial e_k / a_k \cdot \partial a_k / \partial y_j$$

$$e_k = d_k - y_k = d_k - f(a_k), \text{ so}$$

$$\partial e_k / a_k = -f'(a_k)$$

$$= - \sum_k e_k f'(a_k) \cdot \partial a_k / \partial y_j$$


$$a_k = \sum_j w_{kj} y_j, \text{ so}$$

$$\partial a_k / \partial y_j = w_{kj}$$

$$= - \sum_k e_k f'(a_k) \cdot w_{kj}$$

$$= - \sum_k \delta_k^o w_{kj}$$

$$\text{but } \delta_j^o = e_j f'(a_j)$$



1. Provide an example input x , allow the network to compute in feedforward mode, and produce output y .

2. Calculate the error at the output nodes by comparing the desired output d to the actual output y . $\delta_j^o = e_j f'(a_j)$

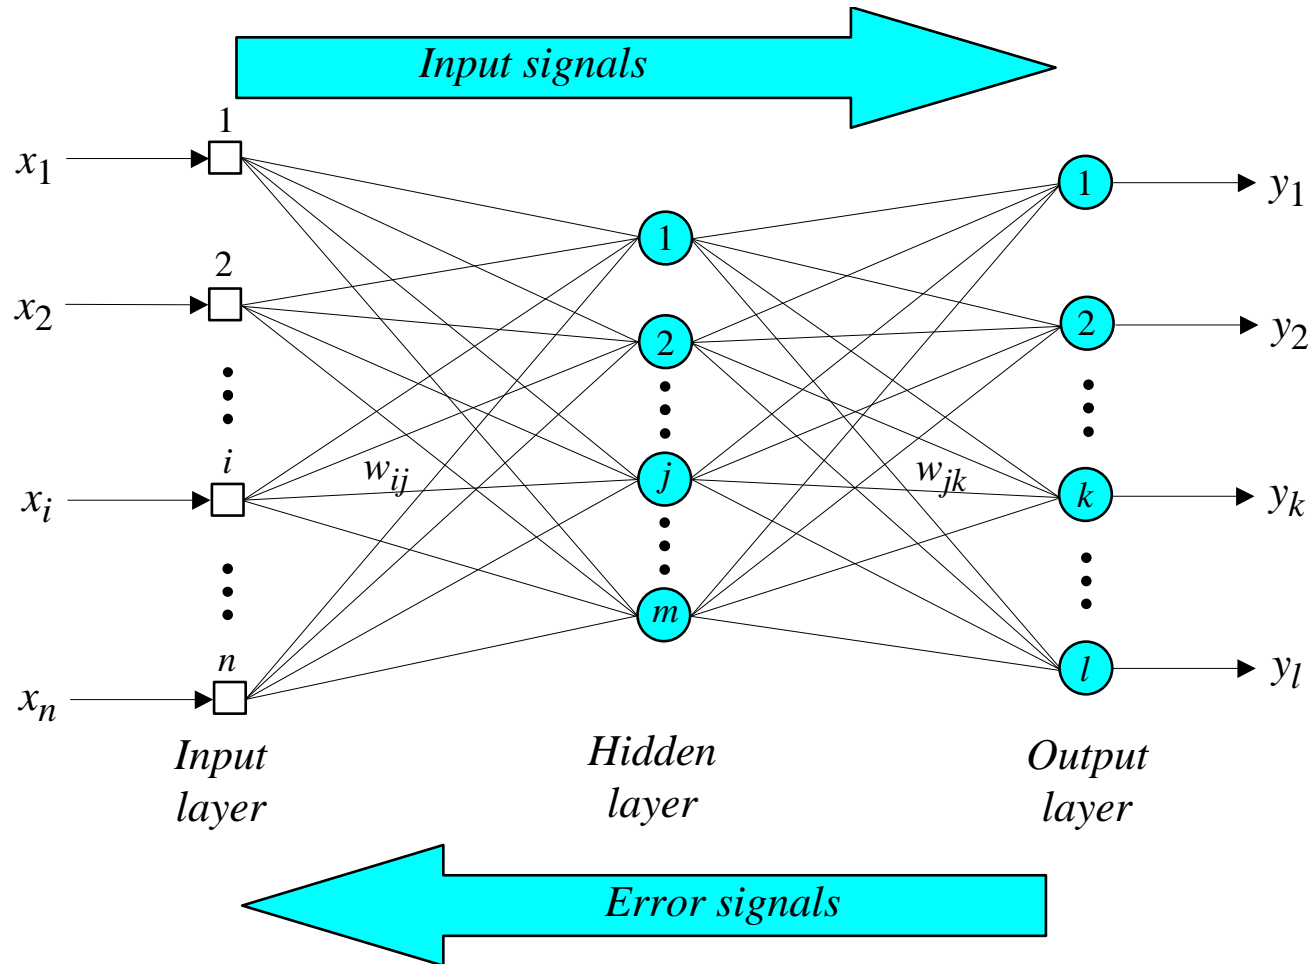
3. Calculate the error at the hidden nodes based on the output node errors. $\delta_j^h = \sum \delta_k^o w_{jk} \cdot f'(a_j)$ (error is propagated back to hidden layer)

4. Adjust all weights (from i to j) at output and hidden nodes on the basis of their calculated error. $\Delta w_{ij} = \alpha \delta_j y_i$

Back-propagation neural network

- Learning in a multilayer network proceeds the same way as for a perceptron.
- A training set of input patterns is presented to the network.
- The network computes its output pattern, and if there is an error - or in other words a difference between actual and desired output patterns - the weights are adjusted to reduce this error.
- In a back-propagation neural network, the learning algorithm has two phases.
- First, a training input pattern is presented to the network input layer. The network propagates the input pattern from layer to layer until the output pattern is generated by the output layer.
- If this pattern is different from the desired output, an error is calculated and then propagated backwards through the network from the output layer to the input layer. The weights are modified as the error is propagated.

Three-layer back-propagation neural network



The back-propagation training algorithm

Step 1: Initialisation

Set all the weights and threshold levels of the network to random numbers uniformly distributed inside a small range:

$$\left(-\frac{2.4}{F_i}, \quad +\frac{2.4}{F_i} \right)$$

where F_i is the total number of inputs of neuron i in the network. The weight initialisation is done on a neuron-by-neuron basis.

Step 2: Activation

Activate the back-propagation neural network by applying inputs $x_1(p), x_2(p), \dots, x_n(p)$ and desired outputs $y_{d,1}(p), y_{d,2}(p), \dots, y_{d,n}(p)$.

(A) Calculate the actual outputs of the neurons in the hidden layer:

$$y_j(p) = \textit{sigmoid} \left[\sum_{i=1}^n x_i(p) \cdot w_{ij}(p) - \theta_j \right]$$

where n is the number of inputs of neuron j in the hidden layer, and *sigmoid* is the *sigmoid* activation function.

Step 2: Activation (continued)

(B) Calculate the actual outputs of the neurons in the output layer:

$$y_k(p) = \textit{sigmoid} \left[\sum_{j=1}^m x_{jk}(p) \cdot w_{jk}(p) - \theta_k \right]$$

where m is the number of inputs of neuron k in the output layer.

Step 3: Weight training

Update the weights in the back-propagation network propagating backward the errors associated with output neurons.

(a) Calculate the error gradient for the neurons in the output layer:

$$\delta_k(p) = y_k(p) \cdot [1 - y_k(p)] \cdot e_k(p)$$

Where

$$e_k(p) = y_{d,k}(p) - y_k(p)$$

Calculate the weight corrections:

$$\Delta w_{jk}(p) = \alpha \cdot y_j(p) \cdot \delta_k(p)$$

Update the weights at the output neurons:

$$w_{jk}(p+1) = w_{jk}(p) + \Delta w_{jk}(p)$$

Step 3: Weight training (continued)

(b) Calculate the error gradient for the neurons in the hidden layer:

$$\delta_j(p) = y_j(p) \cdot [1 - y_j(p)] \cdot \sum_{k=1}^l \delta_k(p) w_{jk}(p)$$

Calculate the weight corrections:

$$\Delta w_{ij}(p) = \alpha \cdot y_i(p) \cdot \delta_j(p)$$

Update the weights at the hidden neurons:

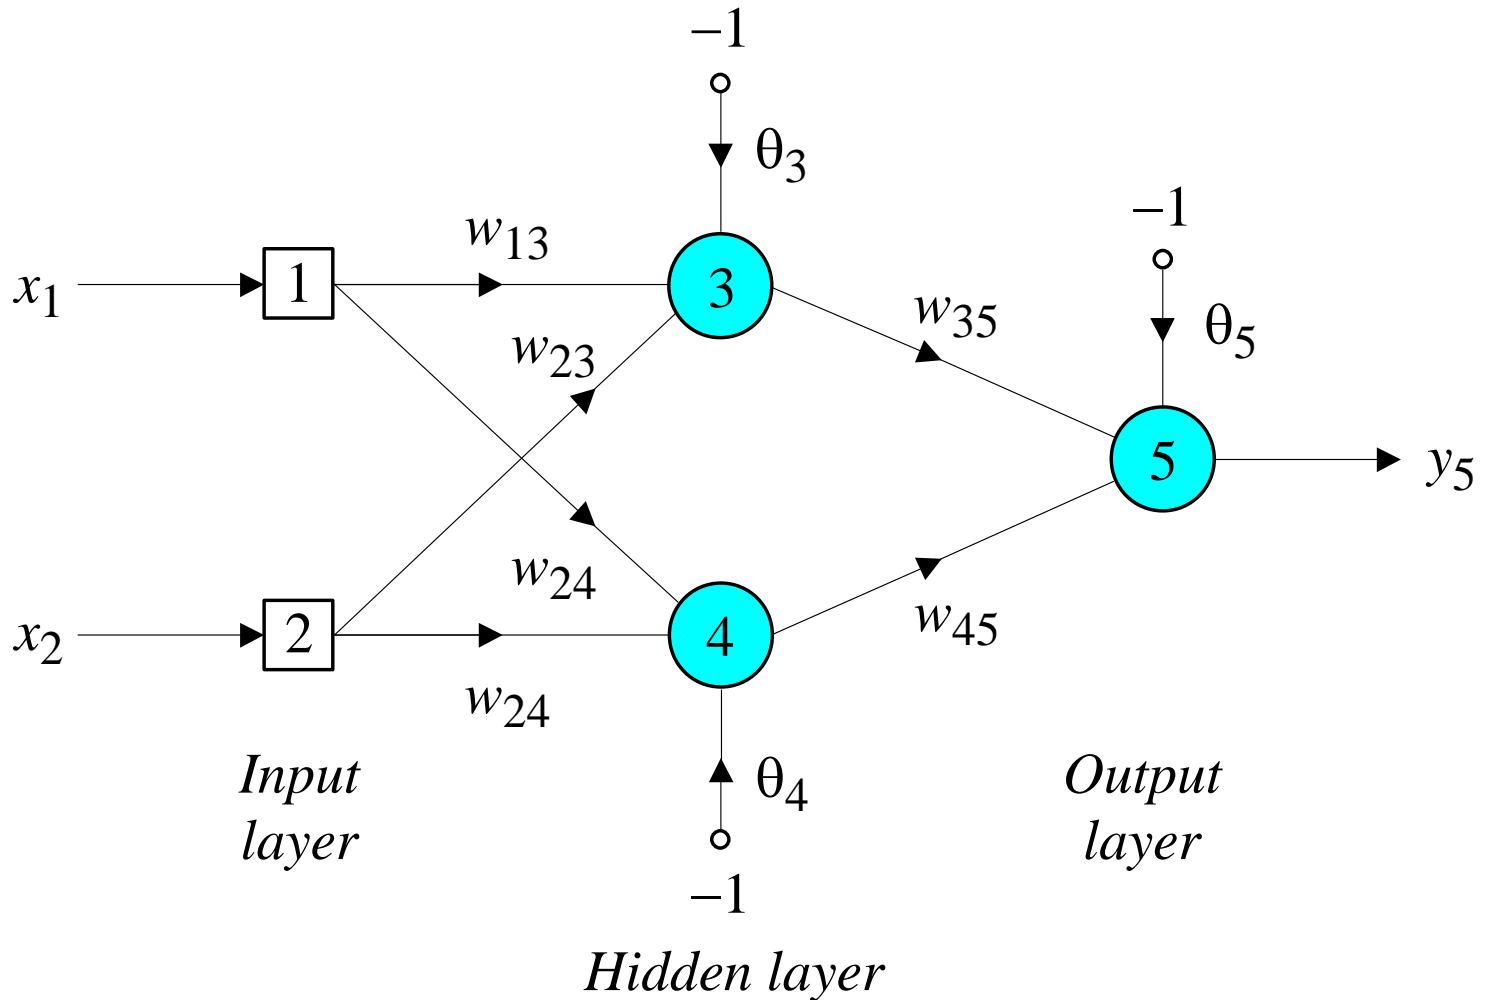
$$w_{ij}(p + 1) = w_{ij}(p) + \Delta w_{ij}(p)$$

Step 4: Iteration

Increase iteration p by one, go back to *Step 2* and repeat the process until the selected error criterion is satisfied.

As an example, we may consider the three-layer back-propagation network. Suppose that the network is required to perform logical operation *Exclusive-OR*. Recall that a single-layer perceptron could not do this operation. Now we will apply the three-layer net.

Three-layer network for solving the Exclusive-OR operation



- The effect of the threshold applied to a neuron in the hidden or output layer is represented by its weight, θ , connected to a fixed input equal to -1 .
- The initial weights and threshold levels are set randomly as follows:

$$w_{13} = 0.5, w_{14} = 0.9, w_{23} = 0.4, w_{24} = 1.0, w_{35} = -1.2,$$

$$w_{45} = 1.1, \theta_3 = 0.8, \theta_4 = -0.1 \text{ and } \theta_5 = 0.3.$$

➤ We consider a training set where inputs x_1 and x_2 are equal to 1 and desired output $y_{d,5}$ is 0. The actual outputs of neurons 3 and 4 in the hidden layer are calculated as

$$y_3 = \text{sigmoid}(x_1w_{13} + x_2w_{23} - \theta_3) = 1/[1 + e^{-(1 \times 0.5 + 1 \times 0.4 - 1 \times 0.8)}] = 0.5250$$

$$y_4 = \text{sigmoid}(x_1w_{14} + x_2w_{24} - \theta_4) = 1/[1 + e^{-(1 \times 0.9 + 1 \times 1.0 + 1 \times 0.1)}] = 0.8808$$

➤ Now the actual output of neuron 5 in the output layer is determined as:

$$y_5 = \text{sigmoid}(y_3w_{35} + y_4w_{45} - \theta_5) = 1/[1 + e^{-(-0.5250 \times 1.2 + 0.8808 \times 1.1 - 1 \times 0.3)}] = 0.5097$$

Thus, the following error is obtained:

$$e = y_{d,5} - y_5 = 0 - 0.5097 = -0.5097$$

- The next step is weight training. To update the weights and threshold levels in our network, we propagate the error, e , from the output layer backward to the input layer.
- First, we calculate the error gradient for neuron 5 in the output layer:

$$\delta_5 = y_5 (1 - y_5) e = 0.5097 \cdot (1 - 0.5097) \cdot (-0.5097) = -0.1274$$

- Then we determine the weight corrections assuming that the learning rate parameter, α , is equal to 0.1:

$$\Delta w_{35} = \alpha \cdot y_3 \cdot \delta_5 = 0.1 \cdot 0.5250 \cdot (-0.1274) = -0.0067$$

$$\Delta w_{45} = \alpha \cdot y_4 \cdot \delta_5 = 0.1 \cdot 0.8808 \cdot (-0.1274) = -0.0112$$

$$\Delta \theta_5 = \alpha \cdot (-1) \cdot \delta_5 = 0.1 \cdot (-1) \cdot (-0.1274) = -0.0127$$

- Next we calculate the error gradients for neurons 3 and 4 in the hidden layer:

$$\delta_3 = y_3(1 - y_3) \cdot \delta_5 \cdot w_{35} = 0.5250 \cdot (1 - 0.5250) \cdot (-0.1274) \cdot (-1.2) = 0.0381$$

$$\delta_4 = y_4(1 - y_4) \cdot \delta_5 \cdot w_{45} = 0.8808 \cdot (1 - 0.8808) \cdot (-0.1274) \cdot 1.1 = -0.0147$$

- We then determine the weight corrections:

$$\Delta w_{13} = \alpha \cdot x_1 \cdot \delta_3 = 0.1 \cdot 1 \cdot 0.0381 = 0.0038$$

$$\Delta w_{23} = \alpha \cdot x_2 \cdot \delta_3 = 0.1 \cdot 1 \cdot 0.0381 = 0.0038$$

$$\Delta \theta_3 = \alpha \cdot (-1) \cdot \delta_3 = 0.1 \cdot (-1) \cdot 0.0381 = -0.0038$$

$$\Delta w_{14} = \alpha \cdot x_1 \cdot \delta_4 = 0.1 \cdot 1 \cdot (-0.0147) = -0.0015$$

$$\Delta w_{24} = \alpha \cdot x_2 \cdot \delta_4 = 0.1 \cdot 1 \cdot (-0.0147) = -0.0015$$

$$\Delta \theta_4 = \alpha \cdot (-1) \cdot \delta_4 = 0.1 \cdot (-1) \cdot (-0.0147) = 0.0015$$

- At last, we update all weights and threshold:

$$w_{13} = w_{13} + \Delta w_{13} = 0.5 + 0.0038 = 0.5038$$

$$w_{14} = w_{14} + \Delta w_{14} = 0.9 - 0.0015 = 0.8985$$

$$w_{23} = w_{23} + \Delta w_{23} = 0.4 + 0.0038 = 0.4038$$

$$w_{24} = w_{24} + \Delta w_{24} = 1.0 - 0.0015 = 0.9985$$

$$w_{35} = w_{35} + \Delta w_{35} = -1.2 - 0.0067 = -1.2067$$

$$w_{45} = w_{45} + \Delta w_{45} = 1.1 - 0.0112 = 1.0888$$


$$\theta_3 = \theta_3 + \Delta \theta_3 = 0.8 - 0.0038 = 0.7962$$

$$\theta_4 = \theta_4 + \Delta \theta_4 = -0.1 + 0.0015 = -0.0985$$

$$\theta_5 = \theta_5 + \Delta \theta_5 = 0.3 + 0.0127 = 0.3127$$

- The training process is repeated until the sum of squared errors is less than 0.001.

Learning curve for operation *Exclusive-OR*

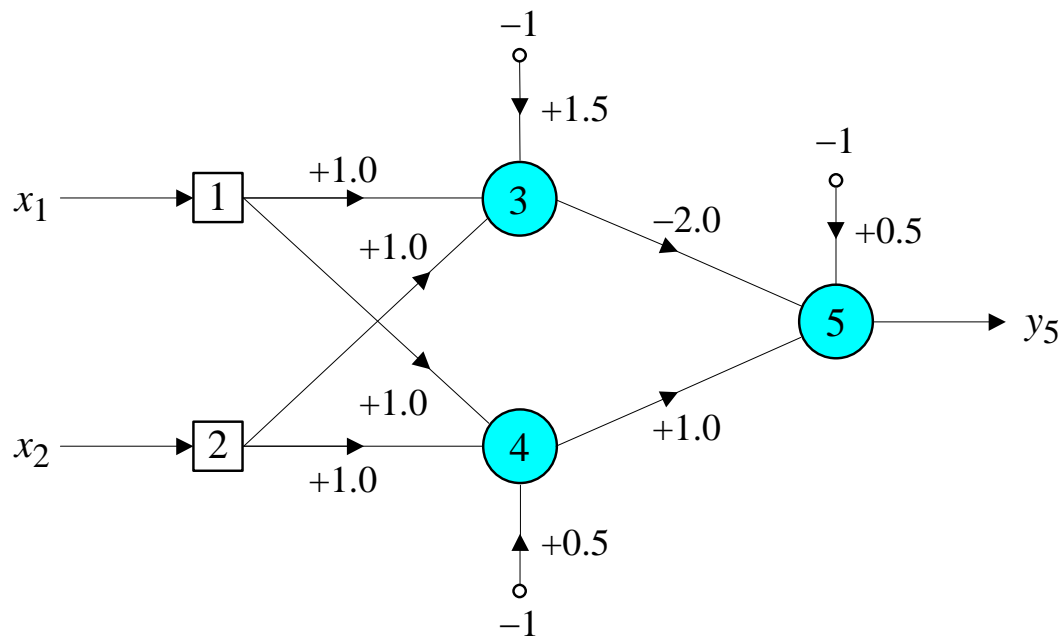


traingd Gradient Descent

Final results of three-layer network learning

Inputs		Desired output y_d	Actual output y_5	Error e	Sum of squared errors
x_1	x_2				
1	1	0	0.0155	-0.0155	0.0010
0	1	1	0.9849	0.0151	
1	0	1	0.9849	0.0151	
0	0	0	0.0175	-0.0175	

Network represented by McCulloch-Pitts model for solving the *Exclusive-OR* operation



Accelerated learning in multilayer neural networks

- We can accelerate training by including a **momentum term** in the delta rule:

$$\Delta w_{jk}(p) = \beta \cdot \Delta w_{jk}(p - 1) + \alpha \cdot y_j(p) \cdot \delta_k(p)$$


where β is a positive number ($0 \leq \beta < 1$) called the **momentum constant**. Typically, the momentum constant is set to 0.95.

<https://youtu.be/6iwvtzXZ4Mo?t=14>

Accelerated learning

- We can accelerate training by including a **momentum term** in the delta rule:



$$\Delta w_{jk}(p) = \beta \cdot \Delta w_{jk}(p - 1) + \alpha \cdot y_j(p) \cdot \delta_k(p)$$



<https://youtu.be/6iwvtzXZ4Mo?t=14>

Learning with momentum for operation *Exclusive-OR*

Training for 126 Epochs



Learning with adaptive learning rate

To accelerate the convergence and yet avoid the danger of instability, we can apply two heuristics:

Learning Rate Detail Explanation with Examples

LEARNING RATE



Play (k)

0:23 / 0:43

CC ⚙️ 🗉

Learning with adaptive learning rate

To accelerate the convergence and yet avoid the danger of instability, we can apply two heuristics:


Heuristic 1

If the change of the sum of squared errors has the same algebraic sign for several consequent epochs, then the learning rate parameter, α , should be increased.

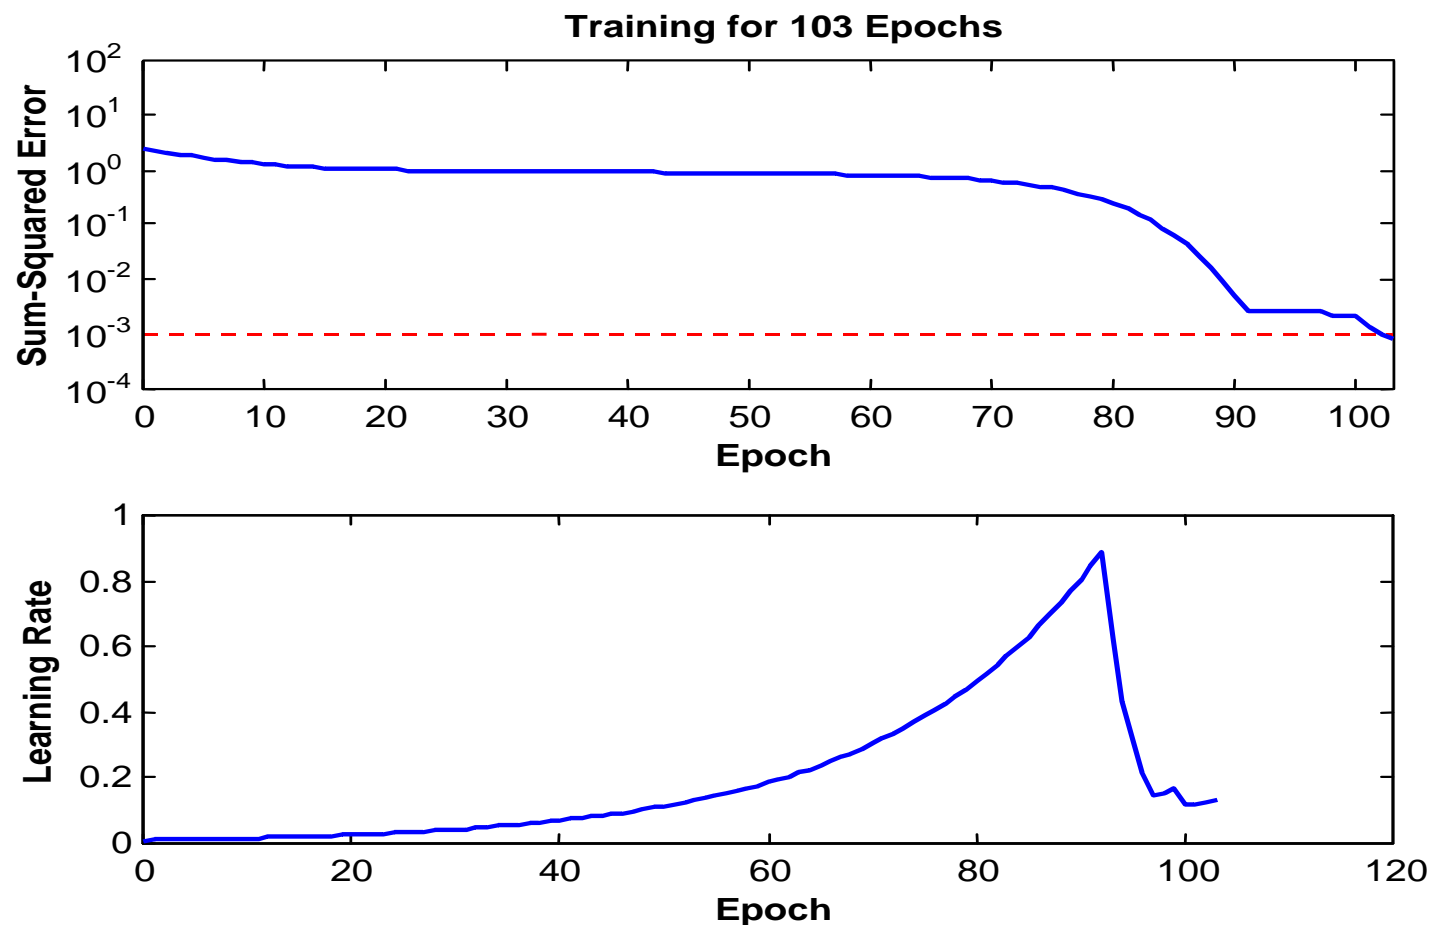
Heuristic 2

If the algebraic sign of the change of the sum of squared errors alternates for several consequent epochs, then the learning rate parameter, α , should be decreased.


<https://youtu.be/TOtKVUtpz-s>

- 
- If the error is less than the previous one, the learning rate is increased (typically by multiplying by 1.05).
 - If the sum of squared errors at the current epoch exceeds the previous value by more than a predefined ratio (typically 1.04), the learning rate parameter is decreased (typically by multiplying by 0.7) and new weights and thresholds are calculated.

Learning with adaptive learning rate



Learning with momentum and adaptive learning rate



Divide Data for Optimal Neural Network Training

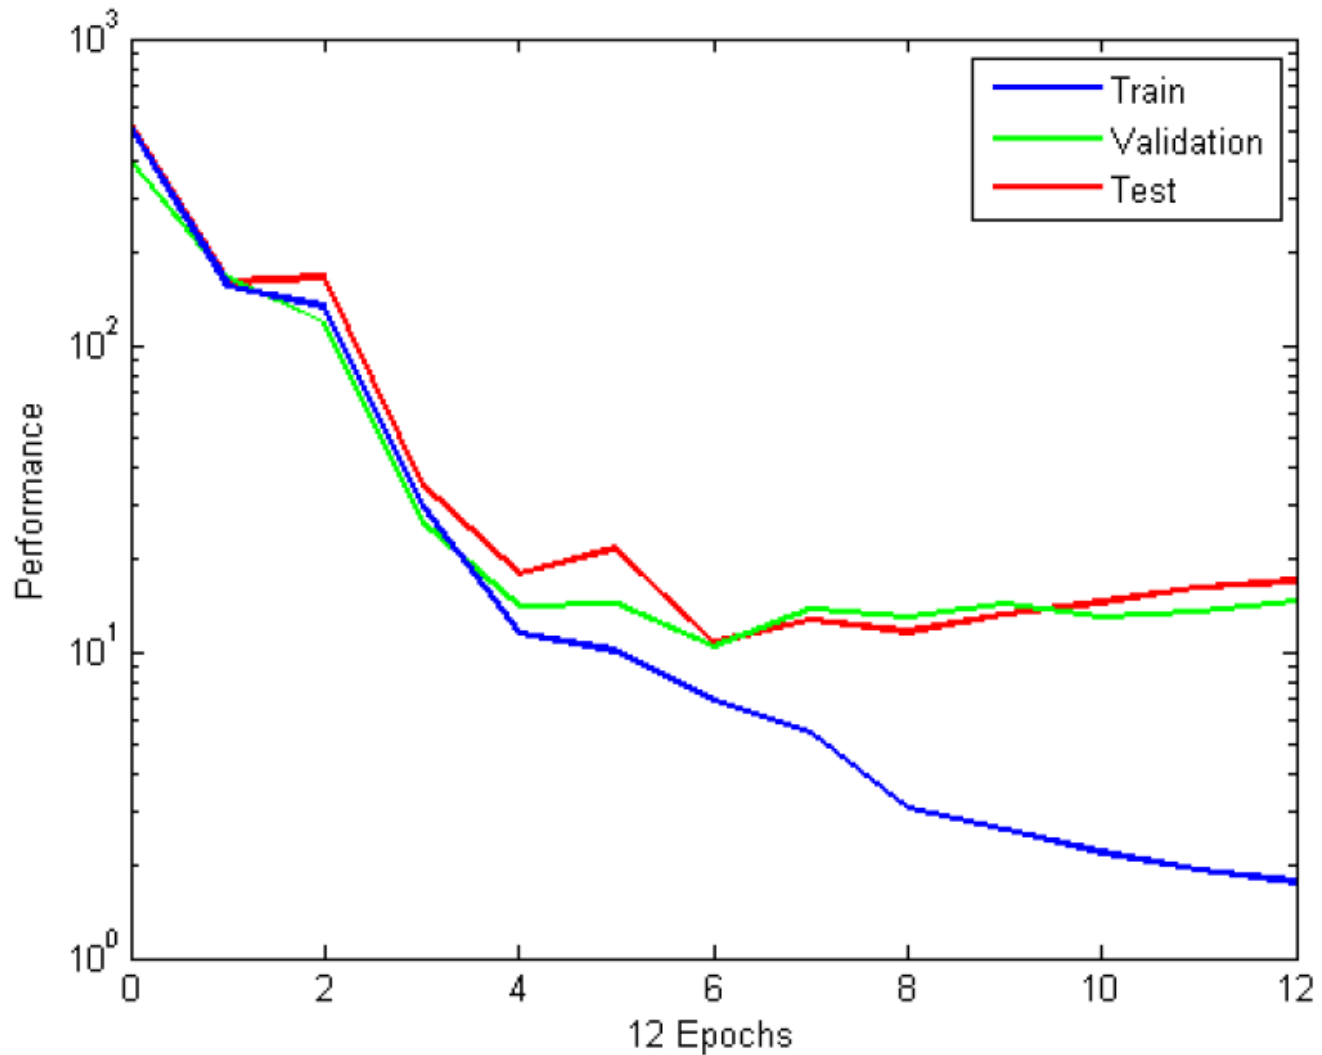
- One of the problems that occur during neural network training is called over fitting (overtraining). The error on the training set is driven to a very small value, but when new data is presented to the network the error is large. The network has memorized the training examples, but it has not learned to generalize to new situations.
- Note that if the number of parameters in the network is much smaller than the total number of points in the training set, then there is little or no chance of over fitting. If you can easily collect more data and increase the size of the training set, then there is no need to worry about over fitting.
- One of the methods for improving generalization of a neural network is early stopping.

Divide Data for Optimal Neural Network Training

When training multilayer networks, the general practice is to first divide the data into three subsets.

- 1) The first subset is the training set, which is used for computing the gradient and updating the network weights and biases.
- 2) The second subset is the validation set. The error on the validation set is monitored during the training process. The validation error normally decreases during the initial phase of training, as does the training set error. However, when the network begins to overfit the data, the error on the validation set typically begins to rise. The network weights and biases are saved at the minimum of the validation set error.
- 3) The test set represents a new data that were not used in training or validation to see how the system behaves for totally new data. It is also useful to plot the test set error during the training process. If the error on the test set reaches a minimum at a significantly different iteration number than the validation set error, this might indicate a poor division of the data set.

Divide Data for Optimal Neural Network Training



Applications

Industry	Business Applications
Aerospace	High-performance aircraft autopilot, flight path simulation, aircraft control systems, autopilot enhancements, aircraft component simulation, and aircraft component fault detection
Automotive	Automobile automatic guidance system, and warranty activity analysis
Banking	Check and other document reading and credit application evaluation
Defense	Weapon steering, target tracking, object discrimination, facial recognition, new kinds of sensors, sonar, radar and image signal processing including data compression, feature extraction and noise suppression, and signal/image identification
Electronics	Code sequence prediction, integrated circuit chip layout, process control, chip failure analysis, machine vision, voice synthesis, and nonlinear modeling
Entertainment	Animation, special effects, and market forecasting
Financial	Real estate appraisal, loan advising, mortgage screening, corporate bond rating, credit-line use analysis, credit card activity tracking, portfolio trading program, corporate financial analysis, and currency price prediction
Industrial	Prediction of industrial processes, such as the output gases of furnaces, replacing complex and costly equipment used for this purpose in the past
Insurance	Policy application evaluation and product optimization
Manufacturing	Manufacturing process control, product design and analysis, process and machine diagnosis, real-time particle identification, visual quality inspection systems, beer testing, welding quality analysis, paper quality prediction, computer-chip quality analysis, analysis of grinding operations, chemical product design analysis, machine maintenance analysis, project bidding, planning and management, and dynamic modeling of chemical process system
Medical	Breast cancer cell analysis, EEG and ECG analysis, prosthesis design, optimization of transplant times, hospital expense reduction, hospital quality improvement, and emergency-room test advisement
Oil and gas	Exploration
Robotics	Trajectory control, forklift robot, manipulator controllers, and vision systems
Speech	Speech recognition, speech compression, vowel classification, and text-to-speech synthesis
Securities	Market analysis, automatic bond rating, and stock trading advisory systems
Telecommunications	Image and data compression, automated information services, real-time translation of spoken language, and customer payment processing systems
Transportation	Truck brake diagnosis systems, vehicle scheduling, and routing systems

NN toolbox

As an example, the file housing.mat contains a predefined set of input and target vectors. The input vectors define data regarding real-estate properties and the target values define relative values of the properties. Load the data using the following command:

```
Load house_dataset
```

Loading this file creates two variables. The input matrix houseInputs consists of 506 column vectors of 13 real estate variables for 506 different houses. The target matrix houseTargets consists of the corresponding 506 relative valuations.

NN toolbox

Function	Algorithm
trainrp	Resilient Backpropagation
trainscg	Scaled Conjugate Gradient
traincgb	Conjugate Gradient with Powell/Beale Restarts
traincgf	Fletcher-Powell Conjugate Gradient
traincgp	Polak-Ribière Conjugate Gradient
trainoss	One Step Secant
traingdx	Variable Learning Rate Gradient Descent
traingdm	Gradient Descent with Momentum
traingd	Gradient Descent

Chapter 5:

Evolutionary Computation and Genetic algorithms

- Introduction, or can evolution be intelligent?
- Simulation of natural evolution
- Genetic algorithms
- Case study: maintenance scheduling with genetic algorithms
- Summary

Can evolution be intelligent ?

- Intelligence can be defined as the capability of a system to adapt its behavior to ever-changing environment.
- Evolutionary computation simulates evolution on a computer.
- The result of such a simulation is a series of optimization algorithms, usually based on a simple set of rules.
- Optimization iteratively improves the quality of solutions until an optimal, or at least feasible, solution is found.

Evolutionary Computation

- If, over successive generations, the organism survives, we can say that this organism is capable of learning to predict changes in its environment.
- The evolutionary approach is based on computational models of natural selection and genetics. We call them **evolutionary computation**, an umbrella term that combines **genetic algorithms**, **evolution strategies** and **genetic programming**.

Simulation of natural evolution

- Evolution can be seen as a process leading to the maintenance of a population's ability to survive and reproduce in a specific environment. This ability is called **evolutionary fitness**.
- Evolutionary fitness can also be viewed as a measure of the organism's ability to **anticipate** changes in its environment.
- The fitness, or the quantitative measure of the ability to predict environmental changes and respond adequately, can be considered as the quality that is optimized in natural life.

How is a population with increasing fitness generated?

- Let us consider a population of rabbits. Some rabbits are faster than others, and we may say that these rabbits possess superior fitness, because they have a greater chance of avoiding foxes, surviving and then breeding.
- If two parents have superior fitness, there is a good chance that a combination of their genes will produce an offspring with even higher fitness. Over time the entire population of rabbits becomes faster to meet their environmental challenges in the face of foxes.

Simulation of natural evolution

- All methods of evolutionary computation simulate natural evolution by creating a population of individuals, evaluating their fitness, generating a new population through genetic operations, and repeating this process a number of times.
- We will start with **Genetic Algorithms** (GAs) as most of the other evolutionary algorithms can be viewed as variations of genetic algorithms .

Genetic Algorithms

- In the early 1970s, John Holland introduced the concept of genetic algorithms.
- His aim was to make computers do what nature does.
- Holland was concerned with algorithms that manipulate strings of binary digits.
- Each artificial “chromosomes” consists of a number of “genes”, and each gene is represented by 0 or 1:

1	0	1	1	0	1	0	0	0	0	0	1	0	1	0	1
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

Genetic Algorithms

- Nature has an ability to adapt and learn without being told what to do.
- In other words, nature finds good chromosomes blindly. GAs do the same.
- Two mechanisms link a GA to the problem it is solving: **encoding** and **evaluation**.
- The GA uses a measure of fitness of individual chromosomes to carry out reproduction.
- As reproduction takes place, the crossover operator exchanges parts of two single chromosomes, and the mutation operator changes the gene value in some randomly chosen location of the chromosome.

Basic genetic algorithms

Step 1: Represent the problem variable domain as a chromosome of a fixed length, choose the size of a chromosome population N , the crossover probability p_c and the mutation probability p_m .

Step 2: Define a fitness function to measure the performance, or fitness, of an individual chromosome in the problem domain. The fitness function establishes the basis for selecting chromosomes that will be mated during reproduction

Step 3: Randomly generate an initial population of chromosomes of size N:

$$x_1, x_2, \dots, x_n$$

Step 4: Calculate the fitness of each individual chromosome:

$$f(x_1), f(x_2), \dots, f(x_n)$$

Step 5: Select a pair of chromosomes for mating from the current population. Parent chromosomes are selected with a probability related to their fitness



Step 6: Create a pair of offspring chromosomes by applying the genetic operators – crossover and mutation.

Step 7: Place the created offspring chromosomes in the new population

Step 8: Repeat **Step 5** until the size of the new chromosome population becomes equal to the size of the initial population, N .

Step 9: Replace the initial (parent) chromosome population with the new (offspring) population.

Step 10: Go to **Step 4**, and repeat the process until the termination criterion is satisfied

Genetic algorithms

- GA represents an iterative process. Each iteration is called a **generation**. A typical number of generations for a simple GA can range from 50 to over 500. The entire set of generations is called a **run**.
- Because GAs use a **stochastic** search method, the fitness of a population may remain stable for a number of generations before a superior chromosome appears.
- A common practice is to terminate a GA after a specified number of generations and then examine the best chromosomes in the population. If no satisfactory solution is found, the GA is restarted

Genetic algorithms: case study

- A simple example will help us to understand how a GA works.
- Let us find the maximum value of the function $(15x - x^2)$ where parameter x varies between 0 and 15.
- For simplicity, we may assume that x takes only integer values. Thus, chromosomes can be built with only four genes:

<i>Integer</i>	<i>Binary code</i>	<i>Integer</i>	<i>Binary code</i>	<i>Integer</i>	<i>Binary code</i>
1	0001	6	0110	11	1011
2	0010	7	0111	12	1100
3	0011	8	1000	13	1101
4	0100	9	1001	14	1110
5	0101	10	1010	15	1111


Genetic algorithms: case study

- Suppose that the size of the chromosome population N is 6
- The crossover probability p_c equals 0.7
- The mutation probability p_m equals 0.001
- The fitness function in our example is defined by:

$$f(x) = 15x - x^2$$

The fitness function and chromosome locations


<i>Chromosome label</i>	<i>Chromosome string</i>	<i>Decoded integer</i>	<i>Chromosome fitness</i>	<i>Fitness ratio, %</i>
X1	1 1 0 0	12	36	16.5
X2	0 1 0 0	4	44	20.2
X3	0 0 0 1	1	14	6.4
X4	1 1 1 0	14	14	6.4
X5	0 1 1 1	7	56	25.7
X6	1 0 0 1	9	54	24.8



Chromosome fitness




Sum of Chromosome fitness

- 
- In natural selection, only the fittest species can survive, breed, and thereby pass their genes on to the next generation.
 - GAs use a similar approach, but unlike nature, the size of the chromosome population remains unchanged from one generation to the next.
 - The last column in Table shows the ratio of the individual chromosome's fitness to the population's total fitness.
 - This ratio determines the chromosome's chance of being selected for mating. The chromosome's average fitness improves from one generation to the next.

Roulette wheel selection

The most commonly used chromosome selection techniques is the **roulette wheel selection**.




Crossover operator

- In our example, we have an initial population of 6 chromosomes. Thus, to establish the same population in the next generation, the roulette wheel would be spun six times.
- Once a pair of parent chromosomes is selected, the **crossover** operator is applied.

Crossover operator

- First, the crossover operator randomly chooses a crossover point where two parent chromosomes “break”, and then exchanges the chromosome parts after that point. As a result, two new offspring are created.
- If a pair of chromosomes does not cross over, then the chromosome cloning takes place, and the offspring are created as exact copies of each parent.

Crossover




- generate Random Number (RN1) [0-1]. if RN is less than P_c (0.7) then do crossover, otherwise no crossover
- crossover point is randomly selected (RN2) [1-3]

Mutation operator

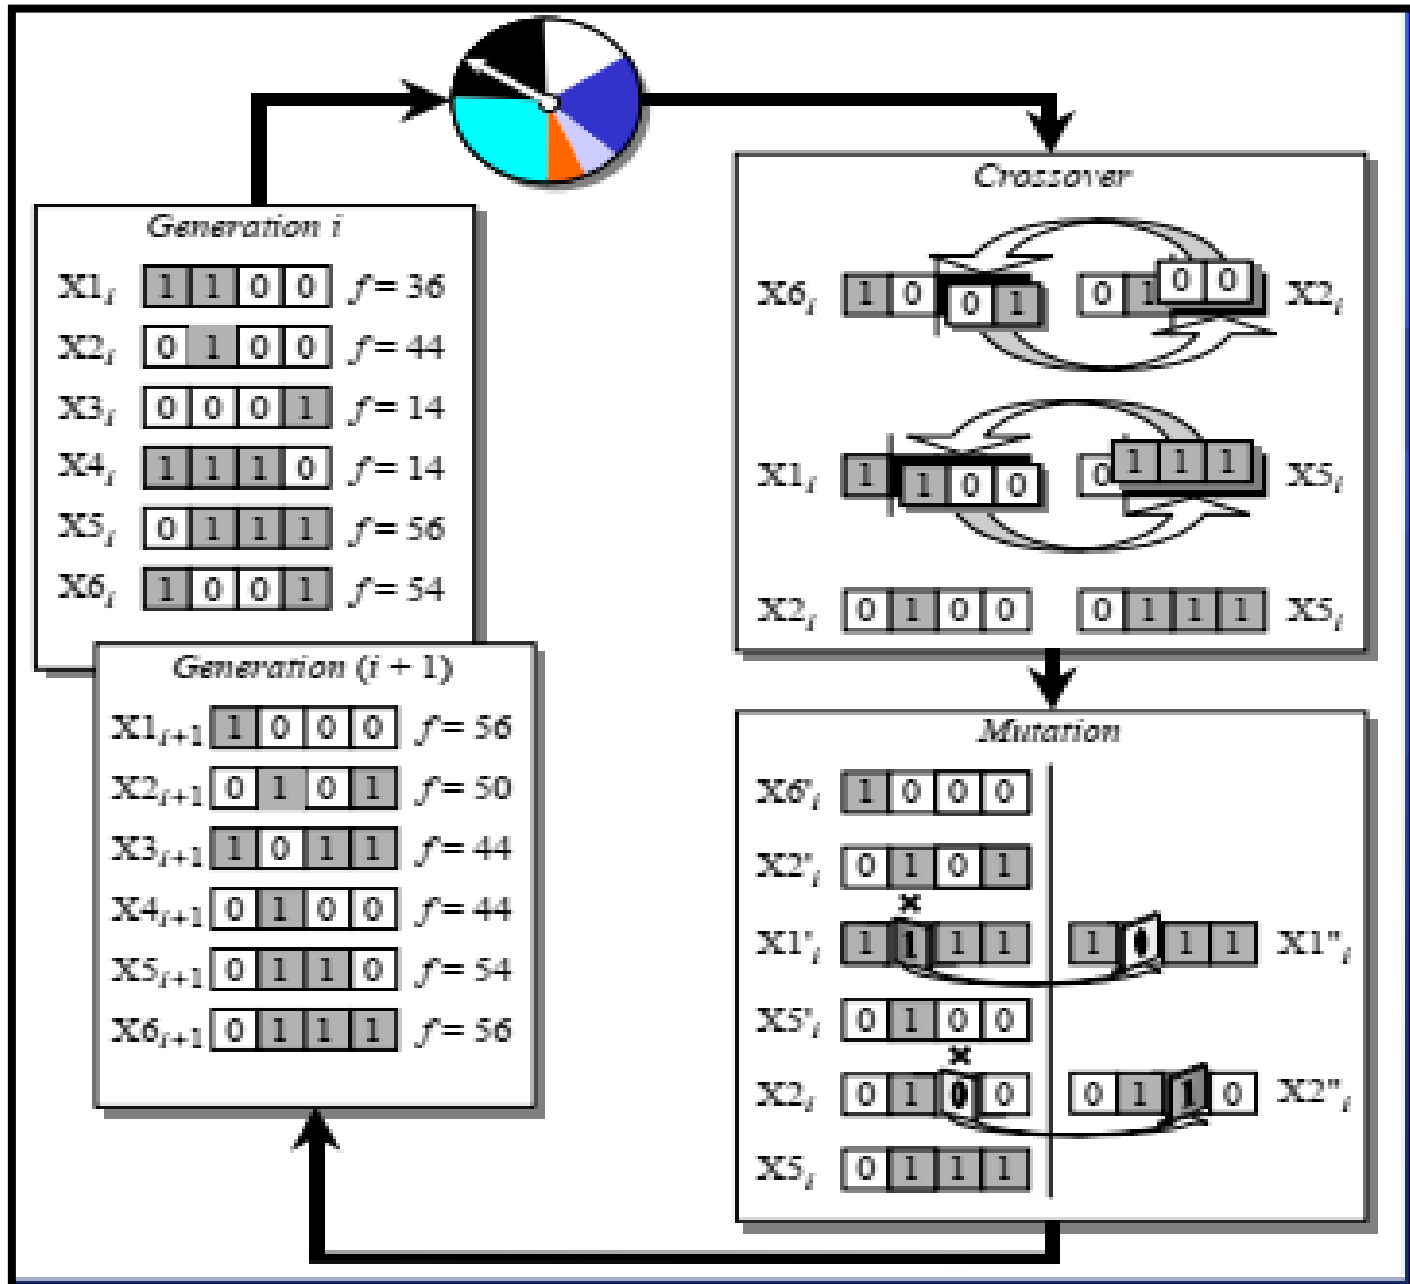
- Mutation represents a change in the gene.
- Mutation is a background operator. Its role is to provide a guarantee that the search algorithm is not trapped on a local optimum.
- The mutation operator flips a randomly selected gene in a chromosome.
- The mutation probability is quite small in nature, and is kept low for GAs, typically in the range between 0.001 and 0.01.

Mutation



- generate Random Number (RN1) [0-1]. if RN1 is less than P_m (0.001) then do mutation, otherwise no mutation.
- mutation point is randomly selected (RN2) [1-4]

The genetic algorithm cycle



Steps in the GA development

1. Specify the problem, define constraints and optimum criteria;
2. Represent the problem domain as a chromosome;
3. Define a fitness function to evaluate the chromosome performance;
4. Construct the genetic operators;
5. Run the GA and tune its parameters.

Case study: maintenance scheduling

- Maintenance scheduling problems are usually solved using a combination of search techniques and heuristics.
- These problems are complex and difficult to solve.

Case study

Scheduling of 7 units in 4 equal intervals

The problem constraints:

- The maximum loads expected during four intervals are 80, 90, 65 and 70 MW;
- Maintenance of any unit starts at the beginning of an interval and finishes at the end of the same or adjacent interval.

Case study

Scheduling of 7 units in 4 equal intervals

The problem constraints:

- The maintenance cannot be aborted or finished earlier than scheduled;
- The net reserve of the power system must be greater or equal to zero at any interval.
- $Net\ reserve = \sum capacity - \sum load$

The optimum criterion is the maximum of the net reserve at any maintenance period.

Case study

Unit data and maintenance requirements

<i>Unit number</i>	<i>Unit capacity, MW</i>	<i>Number of intervals required for unit maintenance</i>
1	20	2
2	15	2
3	35	1
4	40	1
5	15	1
6	15	1
7	10	1

Max capacity=150 MW

Case study

Unit gene pools


Unit 1:	1 1 0 0	0 1 1 0	0 0 1 1	
Unit 2:	1 1 0 0	0 1 1 0	0 0 1 1	
Unit 3:	1 0 0 0	0 1 0 0	0 0 1 0	0 0 0 1
Unit 4:	1 0 0 0	0 1 0 0	0 0 1 0	0 0 0 1
Unit 5:	1 0 0 0	0 1 0 0	0 0 1 0	0 0 0 1
Unit 6:	1 0 0 0	0 1 0 0	0 0 1 0	0 0 0 1
Unit 7:	1 0 0 0	0 1 0 0	0 0 1 0	0 0 0 1

Chromosome for the scheduling problem


Unit 1	Unit 2	Unit 3	Unit 4	Unit 5	Unit 6	Unit 7
0 1 1 0	0 0 1 1	0 0 0 1	1 0 0 0	0 1 0 0	0 0 1 0	1 0 0 0

Case study

The mutation operator



Performance graphs and the best maintenance schedules created in a population of 20 chromosomes




Unit number	Unit capacity, MW
1	20
2	15
3	35
4	40
5	15
6	15
7	10

(a) 50 generations

- The maximum loads during four intervals are 80, 90, 65 and 70 MW


$$fitness = Net\ reserve = \sum capacity - \sum load$$

Performance graphs and the best maintenance schedules created in a population of 20 chromosomes




(b) 100 generations

Performance graphs and the best maintenance schedules created in a population of 100 chromosomes




(a) Mutation rate is 0.001

Performance graphs and the best maintenance schedules created in a population of 100 chromosomes



(b) Mutation rate is 0.01

Performance graphs for 100 generations of 6 chromosomes: local maximum





Lecture 11

Hybrid intelligent systems: Neural expert systems and neuro-fuzzy systems

- Introduction
- Neural expert systems
- Neuro-fuzzy systems
- ANFIS: Adaptive Neuro-Fuzzy Inference System
- Summary



Introduction

- A **hybrid intelligent system** is one that combines at least two intelligent technologies. For example, combining a neural network with a fuzzy system results in a hybrid neuro-fuzzy system.



Comparison of Expert Systems, Fuzzy Systems, Neural Networks

	<i>FS</i>	<i>NN</i>
Knowledge representation	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Uncertainty tolerance	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Adaptability	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Learning ability	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Explanation ability	<input checked="" type="checkbox"/>	<input type="checkbox"/>
* The terms used for grading are:		
<input type="checkbox"/> - bad,	<input type="checkbox"/> - rather bad,	<input checked="" type="checkbox"/> - rather good and <input checked="" type="checkbox"/> - good



Neuro-fuzzy systems

- Fuzzy logic and neural networks are important tools in building intelligent systems.
- However, fuzzy systems lack the ability to learn and cannot adjust themselves to a new environment
- On the other hand, although neural networks can learn, they are opaque to the user.
- The merger of a neural network with a fuzzy system into one integrated system therefore offers a promising approach to building intelligent systems.



Neuro-fuzzy systems

- Integrated neuro-fuzzy systems can combine the parallel computation and learning abilities of neural networks with the humanlike knowledge representation and explanation abilities of fuzzy systems.
- As a result, neural networks become more transparent, while fuzzy systems become capable of learning.



Neuro-fuzzy systems

- A neuro-fuzzy system is, in fact, a neural network that is functionally equivalent to a fuzzy inference model.
- It can be trained to develop IF-THEN fuzzy rules and determine membership functions for input and output variables of the system.
- Expert knowledge can be easily incorporated into the structure of the neuro-fuzzy system.



How does a neuro-fuzzy system look?

- The structure of a neuro-fuzzy system is similar to a multi-layer neural network.
- In general, a neuro-fuzzy system has input and output layers, and three hidden layers that represent membership functions and fuzzy rules.

ANFIS: Adaptive Neuro-Fuzzy Inference System

- The Sugeno fuzzy model was proposed for a systematic approach to generating fuzzy rules from a given input-output data set. A typical Sugeno fuzzy rule can be expressed in the following form:

IF x_1 is A_1
AND x_2 is A_2
.....
AND x_m is A_m
THEN $y = f(x_1, x_2, \dots, x_m)$

where x_1, x_2, \dots, x_m are input variables; A_1, A_2, \dots, A_m are fuzzy sets; and y is either a constant or a linear function of the input variables. When y is a constant, we obtain a zero-order Sugeno fuzzy model in which the consequent of a rule is specified by a singleton. When y is a first-order polynomial, i.e.

$$y = k_0 + k_1x_1 + k_2x_2 + \dots + k_mx_m$$

we obtain a first-order Sugeno fuzzy model.

ANFIS: Adaptive Neuro-Fuzzy Inference System

- Jang's ANFIS is normally represented by a six-layer feedforward neural network.
- ANFIS architecture that corresponds to the first order Sugeno fuzzy model.
- For simplicity, we assume that the ANFIS has two inputs: x_1 and x_2 , and one output: y . Each input is represented by two fuzzy

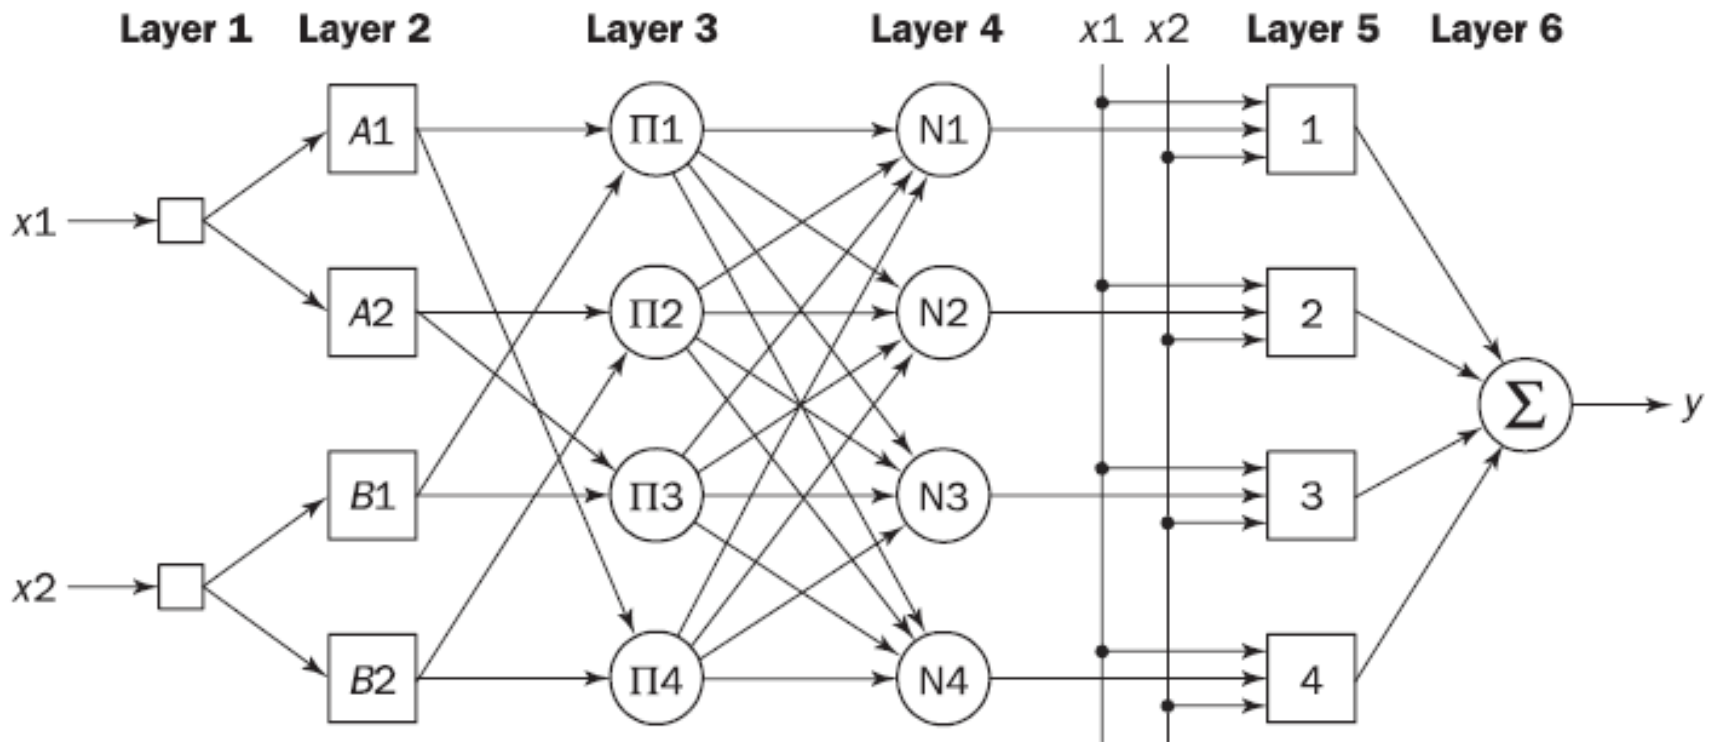


Figure 8.10 Adaptive Neuro-Fuzzy Inference System (ANFIS)

ANFIS: Adaptive Neuro-Fuzzy Inference System

- Each input is represented by two fuzzy sets, and the output by a first-order polynomial. The ANFIS implements four rules:

Rule 1:

IF x_1 is A_1

AND x_2 is B_1

THEN $y = f_1 = k_{10} + k_{11}x_1 + k_{12}x_2$

Rule 2:

IF x_1 is A_2

AND x_2 is B_2

THEN $y = f_2 = k_{20} + k_{21}x_1 + k_{22}x_2$

Rule 3:

IF x_1 is A_2

AND x_2 is B_1

THEN $y = f_3 = k_{30} + k_{31}x_1 + k_{32}x_2$

Rule 4:

IF x_1 is A_1

AND x_2 is B_2

THEN $y = f_4 = k_{40} + k_{41}x_1 + k_{42}x_2$


where x_1, x_2 are input variables; A_1 and A_2 are fuzzy sets on the universe of discourse X_1 ; B_1 and B_2 are fuzzy sets on the universe of discourse X_2 ; and k_{i0}, k_{i1} and k_{i2} is a set of parameters specified for rule i .

ANFIS: Adaptive Neuro-Fuzzy Inference System

- Let us now discuss the purpose of each layer in Jang's ANFIS.
- Layer 1 is the input layer. Neurons in this layer simply pass external crisp signals to Layer 2.

$$y_i^{(1)} = x_i^{(1)},$$

where $x_i^{(1)}$ is the input and $y_i^{(1)}$ is the output of input neuron i in Layer 1.




ANFIS: Adaptive Neuro-Fuzzy Inference System

Layer 2 is the fuzzification layer. Neurons in this layer perform fuzzification. In Jang's model, fuzzification neurons have a bell activation function.

A bell activation function, which has a regular bell shape, is specified as:

$$y_i^{(2)} = \frac{1}{1 + \left(\frac{x_i^{(2)} - a_i}{c_i} \right)^{2b_i}},$$

where $x_i^{(2)}$ is the input and $y_i^{(2)}$ is the output of neuron i in Layer 2; and a_i , b_i and c_i are parameters that control, respectively, the centre, width and slope of the bell activation function of neuron i .



ANFIS: Adaptive Neuro-Fuzzy Inference System

Layer 3 is the rule layer. Each neuron in this layer corresponds to a single Sugeno-type fuzzy rule. A rule neuron receives inputs from the respective fuzzification neurons and calculates the firing strength of the rule it represents. In an ANFIS, the conjunction of the rule antecedents is evaluated by the operator product. Thus, the output of neuron i in Layer 3 is obtained as,


$$y_i^{(3)} = \prod_{j=1}^k x_{ji}^{(3)},$$

where $x_{ji}^{(3)}$ are the inputs and $y_i^{(3)}$ is the output of rule neuron i in Layer 3.

For example,

$$y_{\Pi 1}^{(3)} = \mu_{A1} \times \mu_{B1} = \mu_1,$$

where the value of μ_1 represents the firing strength, or the truth value, of Rule 1.



ANFIS: Adaptive Neuro-Fuzzy Inference System

Layer 4 is the normalisation layer. Each neuron in this layer receives inputs from all neurons in the rule layer, and calculates the normalised firing strength of a given rule.


The normalised firing strength is the ratio of the firing strength of a given rule to the sum of firing strengths of all rules. It represents the contribution of a given rule to the final result.

Thus, the output of neuron i in Layer 4 is determined as,

$$y_i^{(4)} = \frac{x_{ii}^{(4)}}{\sum_{j=1}^n x_{ji}^{(4)}} = \frac{\mu_i}{\sum_{j=1}^n \mu_j} = \bar{\mu}_i,$$


where $x_{ji}^{(4)}$ is the input from neuron j located in Layer 3 to neuron i in Layer 4, and n is the total number of rule neurons. For example,

$$y_{N1}^{(4)} = \frac{\mu_1}{\mu_1 + \mu_2 + \mu_3 + \mu_4} = \bar{\mu}_1$$



ANFIS: Adaptive Neuro-Fuzzy Inference System

Layer 5 is the defuzzification layer. Each neuron in this layer is connected to the respective normalisation neuron, and also receives initial inputs, x_1 and x_2 . A defuzzification neuron calculates the weighted consequent value of a given



How does an ANFIS learn?

- An ANFIS uses a hybrid learning algorithm that combines the least-squares estimator and the gradient descent method
- First, initial activation functions are assigned to each membership neuron.
- In the ANFIS training algorithm, each epoch is composed from a forward pass and a backward pass.
- In the forward pass, a training set of input patterns (an input vector) is presented to the ANFIS, neuron outputs are calculated on the layer-by-layer basis, and rule consequent parameters are identified by the least squares estimator.

How does an ANFIS learn?

- In the Sugeno-style fuzzy inference, an output, y , is a linear function. Thus, given the values of the membership parameters and a training set of P input-output patterns, we can form P linear equations in terms of the consequent parameters as:

$$\left\{ \begin{array}{l} y_d(1) = \bar{\mu}_1(1)f_1(1) + \bar{\mu}_2(1)f_2(1) + \dots + \bar{\mu}_n(1)f_n(1) \\ y_d(2) = \bar{\mu}_1(2)f_1(2) + \bar{\mu}_2(2)f_2(2) + \dots + \bar{\mu}_n(2)f_n(2) \\ \vdots \\ y_d(p) = \bar{\mu}_1(p)f_1(p) + \bar{\mu}_2(p)f_2(p) + \dots + \bar{\mu}_n(p)f_n(p) \end{array} \right.$$

How does an ANFIS learn?

where y_d is a $P \times 1$ desired output vector,

$$y_d = \begin{bmatrix} y_d(1) \\ y_d(2) \\ \vdots \\ y_d(p) \\ \vdots \\ y_d(P) \end{bmatrix}$$

$$Y_d = AK$$

$$K = A^{-1}Y_d$$

A is a $P \times n(1 + m)$ matrix,

$$A = \begin{bmatrix} \bar{\mu}_1(1) & \bar{\mu}_1(1)x_1(1) & \cdots & \bar{\mu}_1(1)x_m(1) & \cdots & \bar{\mu}_n(1) & \bar{\mu}_n(1)x_1(1) & \cdots & \bar{\mu}_n(1)x_m(1) \\ \bar{\mu}_1(2) & \bar{\mu}_1(2)x_1(2) & \cdots & \bar{\mu}_1(2)x_m(2) & \cdots & \bar{\mu}_n(2) & \bar{\mu}_n(2)x_1(2) & \cdots & \bar{\mu}_n(2)x_m(2) \\ \vdots & \vdots & \cdots & \vdots & \cdots & \vdots & \vdots & \cdots & \vdots \\ \bar{\mu}_1(p) & \bar{\mu}_1(p)x_1(p) & \cdots & \bar{\mu}_1(p)x_m(p) & \cdots & \bar{\mu}_n(p) & \bar{\mu}_n(p)x_1(p) & \cdots & \bar{\mu}_n(p)x_m(p) \end{bmatrix}$$

and k is an $n(1 + m) \times 1$ vector of unknown consequent parameters,

$$k = [k_{10} \ k_{11} \ k_{12} \ \dots \ k_{1m} \ k_{20} \ k_{21} \ k_{22} \ \dots \ k_{2m} \ \dots \ k_{n0} \ k_{n1} \ k_{n2} \ \dots \ k_{nm}]^T$$

How does an ANFIS learn?

- Usually the number of input-output patterns P used in training is greater than the number of consequent parameters.
- It means that we are dealing here with an over determined problem, and thus exact solution may not even exist.
- Instead, we solve for K numerically.

How does an ANFIS learn?

- In the ANFIS training algorithm suggested by Jang, both antecedent parameters and consequent parameters are optimised.
- In the forward pass, the consequent parameters are adjusted while the antecedent parameters remain fixed.
- In the backward pass, the antecedent parameters are tuned while the consequent parameters are kept fixed.

Function approximation using the ANFIS model

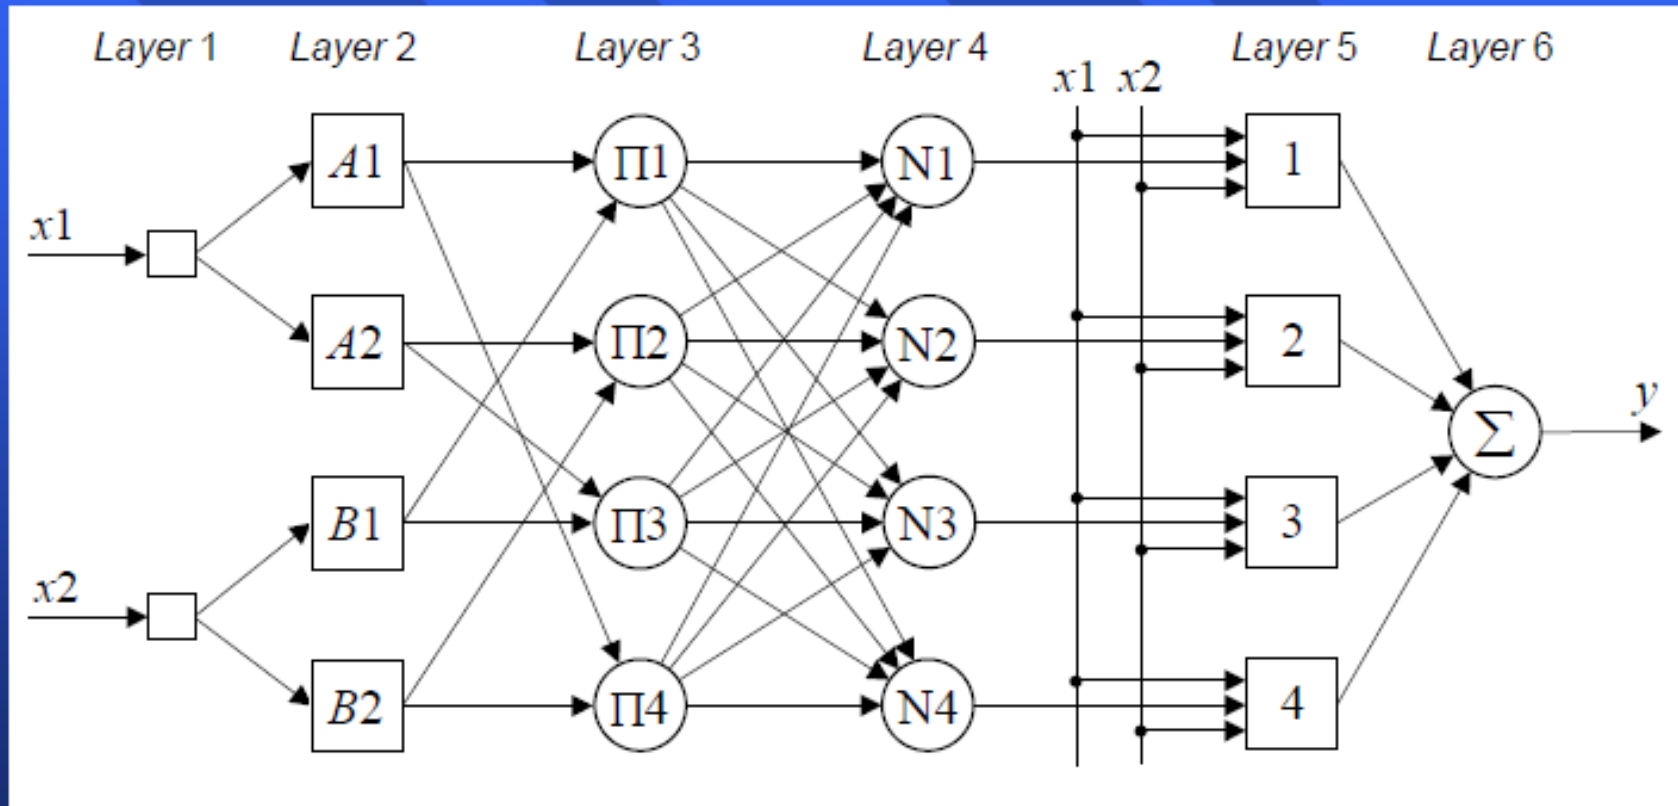
- In this example, an ANFIS is used to follow a trajectory of the non-linear function defined by the equation

$$y = \frac{\cos(2 x_1)}{e^{x_2}}$$

- First, we choose an appropriate architecture for the ANFIS. An ANFIS must have two inputs – x_1 and x_2 – and one output – y .
- Thus, in our example, the ANFIS is defined by four rules, and has the structure shown below.


Example (matlab command line)

An ANFIS model with four rules




Example (matlab command line)

Learning in an ANFIS with two membership functions assigned to each input (one epoch)



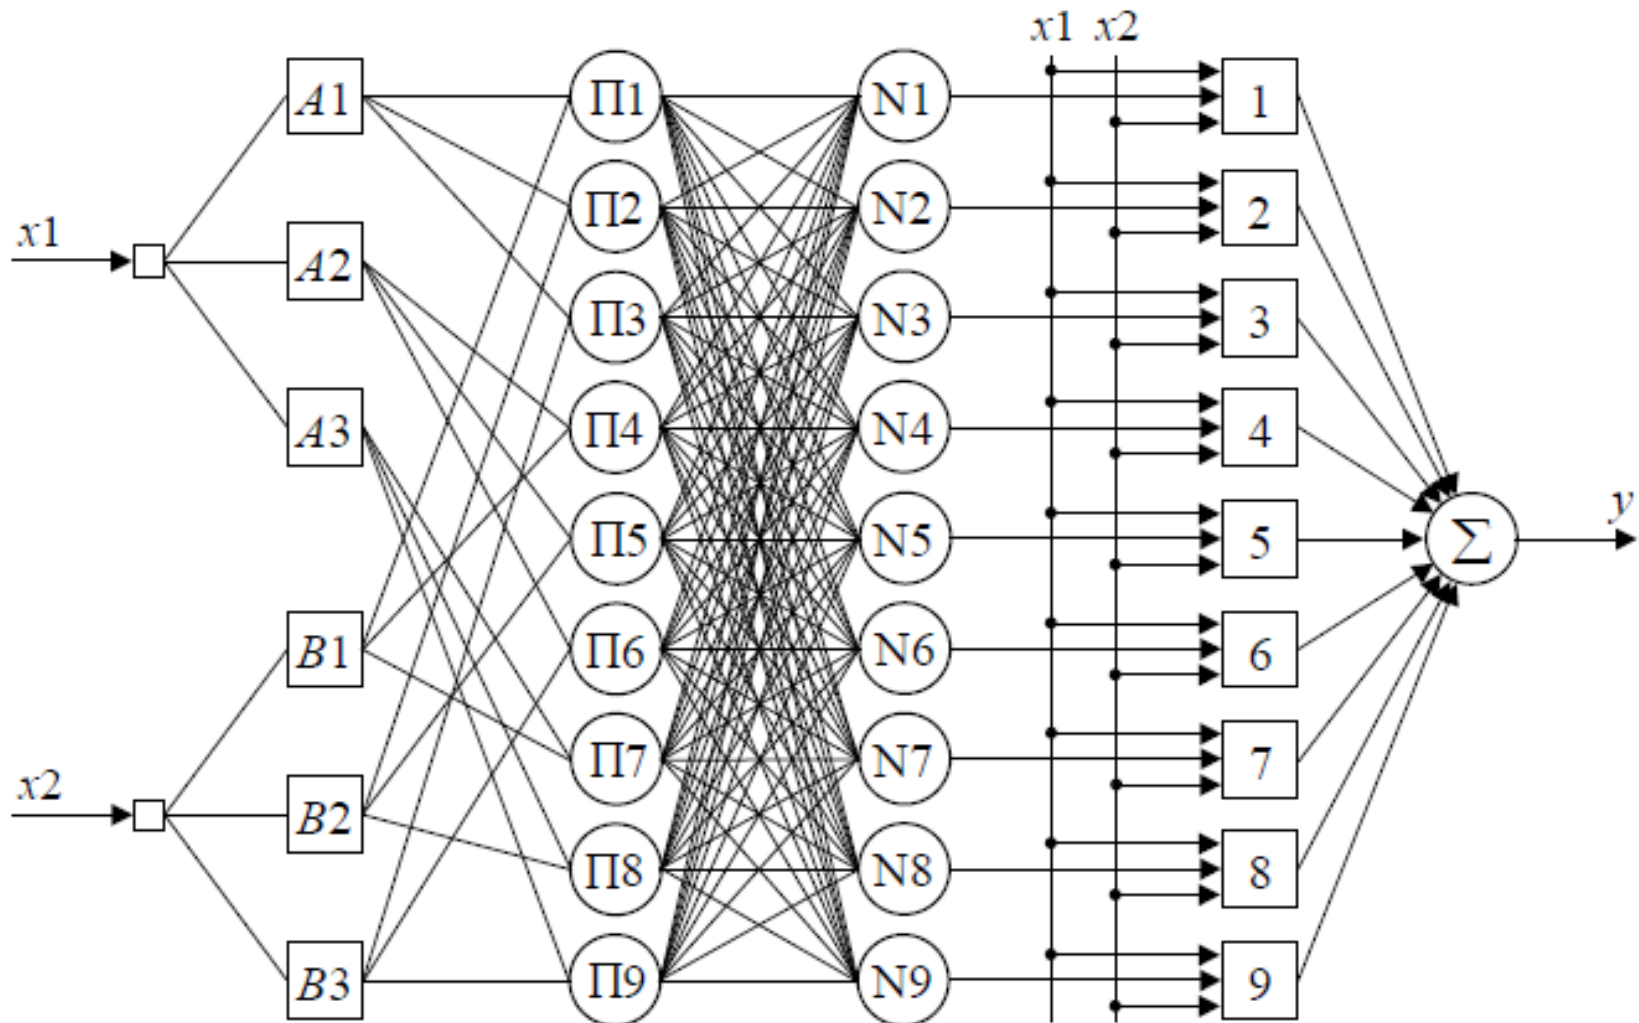
Example (matlab command line)

Learning in an ANFIS with two membership functions assigned to each input (100 epochs)




Example (matlab command line)

An ANFIS model with nine rules




Example (matlab command line)

Learning in an ANFIS with three membership functions assigned to each input (one epoch)





Example (matlab command line)

Initial and final membership functions of the ANFIS



(a) Initial membership functions.



(b) Membership functions after 100 epochs of training.

