1 What is Artificial Intelligence (AI):

AI is a branch of computer science devoted to the study of computer designed to imitate the thinking and reasoning power of the human mind.

The definitions of artificial intelligence according to eight recent textbooks are shown in **fig.(1.1).** These definitions vary along two main dimensions. The ones on top are concerned with **Thought Processes** and **Reasoning**, whereas the ones on the bottom address **behavior**. Also, the definitions on the left measure success in terms of **human performance**, whereas the ones on the right measure against an ideal **concept of intelligence**, which we will call rationality. **A system is rational if it does the right thing**. This gives us four possible goals to keep track the artificial intelligence, as seen in the **figure (1.1).**

Historically, all four approaches have been followed. As one might expect, a tension exists between approaches centered around humans and approaches centered around rationality. A human-centered approach must be an **empirical science**, involving hypothesis and experimental confirmation. A rationalist approach involves a combination of mathematics and engineering.

	Human	Rational
Thought	**Thinking Humanly** "The exciting new effort to make computers think … *machines with minds*, in the full and literal sense." **(Haugeland, 1985)** "[The automation of] activities that we associate with human thinking, activities such as decision- making, problem solving, learning" **(Bellman, 1978)**	**Thinking Rationally** "The study of mental faculties through the use of computational models." **(Charniak and McDermott, 1985)** "The study of the computations that make it possible to perceive, reason, and act." **(Winston, 1992)**
Behavior	**Acting Humanly** "The art of creating machines that perform functions that require intelligence when performed by people" **(Kurzweil, 1990)** "The study of how to make computers do things at which, at the moment, people are better" **(Rich and Knight, 1991)**	**Acting Rationally** "Computational Intelligence is the study of the design of intelligent agents." **(Poole *et al.*, 1998)** "AI . . . is concerned with intelligent behavior in artifacts." **(Nilsson, 1998)**

Figure (1.1): Some definitions of AI (organized into four categories)

Acting Humanly: The Turing Test approach

The Turing Test, proposed by Alan Turing (1950), was designed to provide a satisfactory operational definition of intelligence. Turing defined intelligent behavior as the ability to achieve human-level performance in all cognitive tasks, sufficient to fool an **interrogator**. Roughly speaking, the test he proposed is that the computer should be interrogated by a human via a teletype, and passes the test if the **interrogator** cannot tell if there is a computer or a human at the other end, and whether or not a computer is really intelligent if it passes.

Thinking Humanly: The cognitive modelling approach

When we have a precise theory of the mind, it becomes possible to express the theory as a computer program. If the program's input/output and timing behavior matches human behavior, that is evidence that some of the program's mechanisms may also be operating in humans. **For example**, Newell and Simon, who developed GPS, the "***General Problem Solver***" were not content to have their program correctly solve problems. They were more concerned with comparing the trace of its reasoning steps to traces of human subjects solving the same problems. This is in contrast to other researchers such as Wang (I960), who were concerned with getting the right answers regardless of how humans might do it.

Thinking Rationally: The laws of thought approach

The **law of thought** approach concerned with **correct inferences, regardless the ways**, but there were two main obstacles to create intelligent systems **First**, it is not easy to take informal knowledge and state it in the formal terms required by logical notation, **Second**, there is a big difference between being able to solve a problem "in principle" and doing so in practice. Even problems with just a few dozen facts can exhaust the computational resources of any computer unless it has some guidance as to which reasoning steps to try first. Although both of these obstacles apply to any attempt to build computational reasoning systems, they appeared first in the logicist tradition because the power of the representation and reasoning systems are well-defined and fairly well understood.

Acting Rationally: The rational agent approach

In this approach, AI is viewed as the study and construction of rational agents. In the "***laws of thought***" approach concerned with correct inferences. Making correct inferences is sometimes part of being a rational agent, because one way to act rationally is to reason logically to the conclusion that a given action will achieve one's goals, and then to act on that conclusion. On the other hand, correct inference is not all of rationality, because there are often situations where there is no provably correct thing to do, yet something must still be done. There are also ways of acting rationally that cannot be reasonably said to involve inference. **For example**, *pulling one's hand off of **a hot stove** is a reflex action that is more successful than a **slower** action taken after careful deliberation*. All the "cognitive skills" needed for the Turing Test are there to allow rational actions. Thus, we need the ability to represent knowledge and reason with it because this enables us to reach good decisions in a wide variety of situations. The study of AI as rational agent design therefore has **two advantages**. **First**, it is more general than the "***laws of thought***" approach, because correct inference is only a useful mechanism for achieving rationality, and not a necessary one. **Second**, it is easier to scientific development than approaches based on human behavior or human thought, because the standard of rationality is clearly defined and completely general.

1.2 The Disciplines of AI:

The subject of AI deals with the various kinds of knowledge representation schemes, different techniques of intelligent search, various methods for resolving uncertainty of data and knowledge, different schemes for automated machine learning and many others. The subject of AI has been enriched with a wide discipline of knowledge from Philosophy, Psychology, Cognitive Science, Computer Science, Mathematics and Engineering. Thus in **fig.(1.2),** they have been referred to as the parent disciplines of AI. **Fig.(1.2)** also reveals the subject area of AI and its application areas.

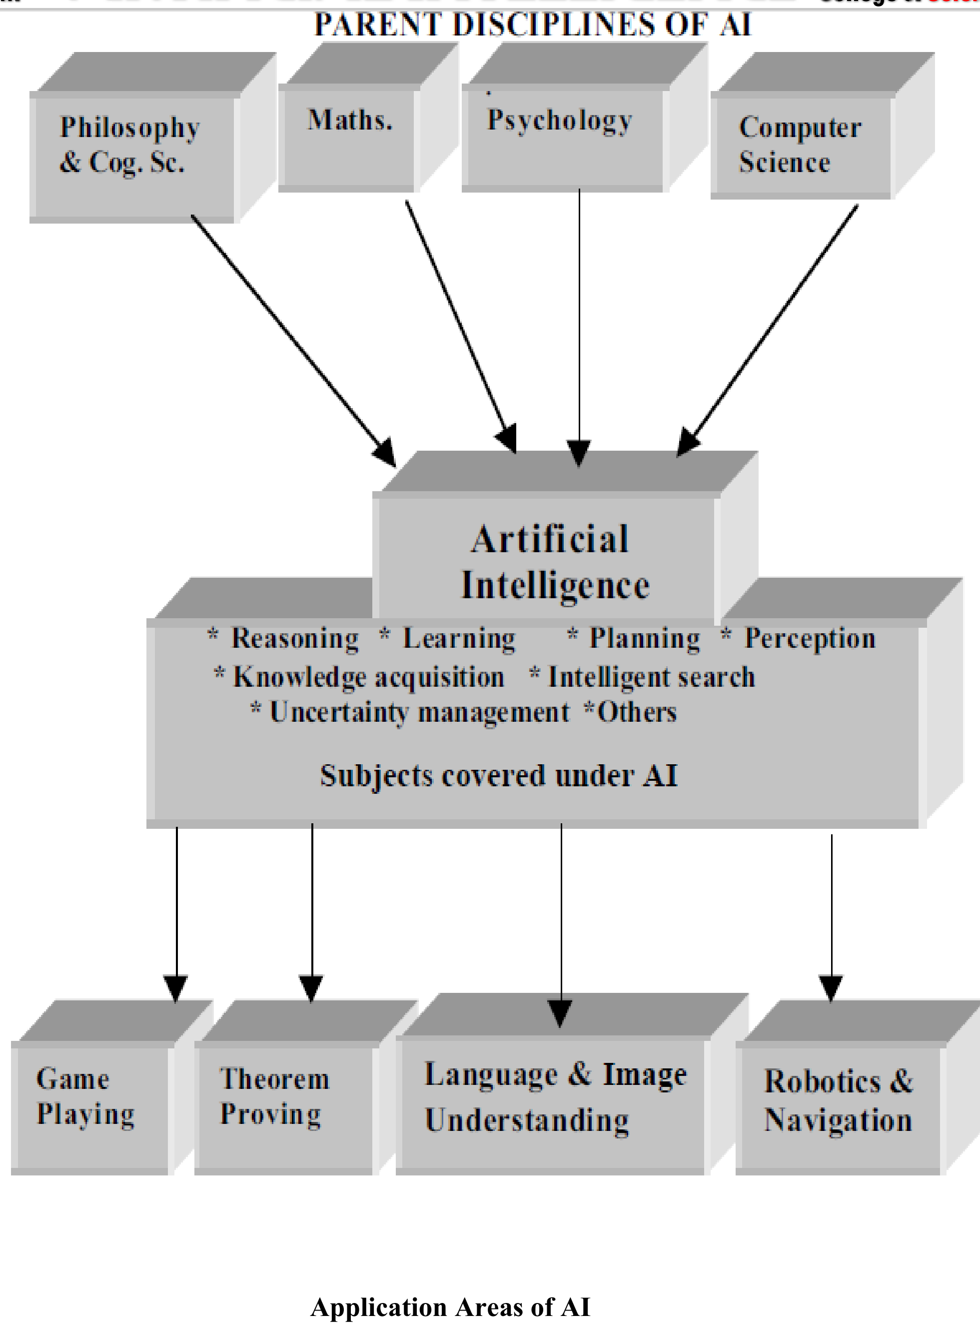

Fig.(1.2): AI, its parent disciplines and application areas.

1.3 The Subject of AI

The subject of AI was originated with game-playing and theorem-proving programs and was gradually enriched with theories from a number of parent disciplines. As a young discipline of science, the significance of the topics covered under the subject changes considerably with time. At present, the topics which we find significant and useful to understand the subject are outlined below:

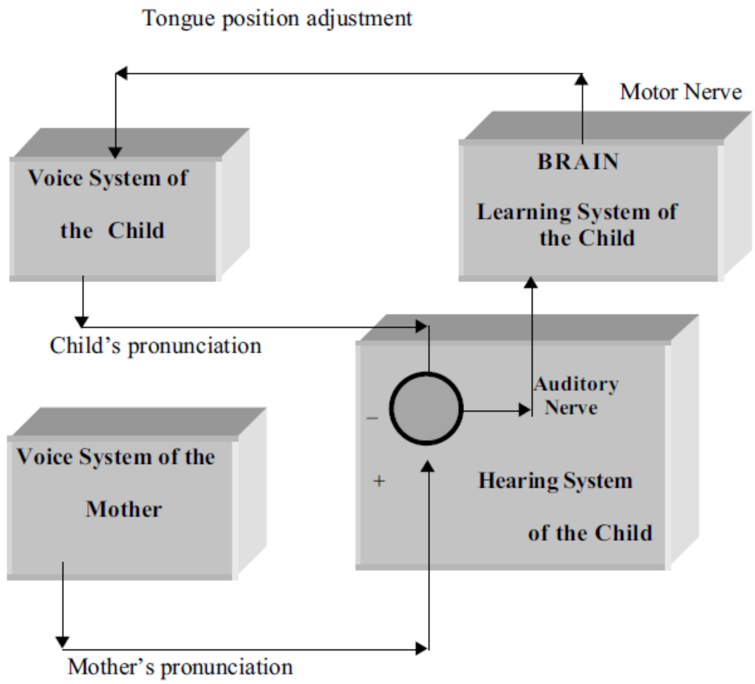

Fig. (1.3): Pronunciation learning of a child from his mother.

[1] Learning Systems: Among the subject areas covered under AI, learning systems needs special mention. The concept of learning is illustrated here with reference to a natural problem of learning of pronunciation by a child from his mother **fig.(1.3).** The hearing system of the child receives the pronunciation of the character "A" and the voice system attempts to imitate it. The difference of the mother's and the child's pronunciation, hereafter called the error signal, is received by the child's learning system through the auditory nerve, and an actuation signal is generated by the learning system through a motor nerve for adjustment of the pronunciation of the child. The adaptation of the child's voice system is continued until the amplitude of the error signal is insignificantly low. Each time the voice system passes through an adaptation cycle, the resulting tongue position of the child for speaking "A" is saved by the learning process. The learning problem discussed above is an example of the well-known *parametric learning*, where the adaptive learning process adjusts the parameters of the child's voice system autonomously to keep its response close enough to the "*sample training pattern*". The artificial neural networks, which represent the electrical analogue of the biological nervous systems, are gaining importance for their increasing applications in supervised (parametric) learning problems. Besides this type, the other common learning methods, which we do unknowingly, are inductive and analogy-based learning. In inductive learning, the learner makes generalizations from examples. For instance, noting that "cuckoo flies", "parrot flies" and "sparrow flies", the learner generalizes that "birds fly". On the other hand, in analogy-based learning, the learner, for example, learns the motion of electrons in an atom analogously from his knowledge of planetary motion in solar systems.

[2] Knowledge Representation and Reasoning: In a reasoning problem, one has to reach a pre-defined goal state from one or more given initial states. So, the lesser the number of transitions for reaching the goal state, the higher the efficiency of the reasoning system. Increasing the efficiency of a reasoning system thus requires minimization of intermediate states, which indirectly calls for an organized and complete knowledge base. A complete and organized storehouse of knowledge needs minimum search to identify the appropriate knowledge at a given problem state and thus yields the right next state on the leading edge of the problem-solving process. Organization of knowledge, therefore, is of paramount importance in knowledge engineering. A variety of knowledge representation techniques are in use in Artificial Intelligence. ***Production rules, semantic nets, frames, filler and slots, and predicate logic*** are only a few to mention. The selection of a particular type of representational scheme of knowledge depends both on the nature of applications and the choice of users.

[3] Planning: Another significant area of AI is *planning*. The problems of reasoning and planning share many common issues, but have a basic difference that originates from their definitions. The reasoning problem is mainly concerned with the testing of the satisfiability of a goal from a given set of data and knowledge. The planning problem, on the other hand, deals with the determination of the methodology by which a successful goal can be achieved from the known initial states.

[4] Knowledge Acquisition: Acquisition (Elicitation) of knowledge is equally hard for machines as it is for human beings. It includes generation of new pieces of knowledge from given knowledge base, setting dynamic data structures for existing knowledge, learning knowledge from the environment and refinement of knowledge. Automated acquisition of knowledge by machine learning approach is an active area of current research in Artificial Intelligence.

[5] Intelligent Search: Search problems, which we generally encounter in Computer Science, are of a deterministic nature, i.e., the order of visiting the elements of the search space is known. For example, in depth first and breadth first search algorithms, one knows the sequence of visiting the nodes in a tree. However, search problems, which we will come across in AI, are non-deterministic and the order of visiting the elements in the search space is completely dependent on data sets.

[6] Soft Computing: Soft computing, according to Prof. Zadeh, is "an emerging approach to computing, which parallels the remarkable ability of the human mind to reason and learn in an environment of uncertainty and imprecision". It, in general, is a collection of computing tools and techniques, shared by closely related disciplines that include fuzzy logic, artificial neural nets, genetic algorithms, belief calculus, and some aspects of machine learning like inductive logic programming. These tools are used independently as well as jointly depending on the type of the domain of applications. The scope of the first three tools in the broad spectrum of AI is outlined below.

- **Fuzzy Logic**: Fuzzy logic deals with fuzzy sets and logical connectives for modeling the human-like reasoning problems of the real world. A fuzzy set, unlike conventional sets, includes all elements of the universal set of the domain but with varying membership values in the interval [0,1]. It may be noted that a conventional set contains its members with a value of membership equal to one and disregards other elements of the universal set, for they have zero membership. The most common operators applied to fuzzy sets are AND (minimum), OR (maximum) and negation (complementation), where AND and OR have binary arguments, while negation has unary argument. The logic of fuzzy sets was proposed by Zadeh, who introduced the concept in systems theory, and later extended it for approximate reasoning in expert systems.

Artificial Neural Nets: Artificial neural nets (ANN) are electrical analogues of the biological neural nets. Biological nerve cells, called neurons, receive signals from neighboring neurons or receptors through dendrites, process the received electrical pulses at the cell body and transmit signals through a large and thick nerve fiber, called an axon. The electrical model of a typical biological neuron consists of a linear activator, followed by a non-linear inhibiting function. The linear activation function yields the sum of the weighted input excitation, while the non-linear inhibiting function attempts to arrest the signal levels of the sum. The resulting signal, produced by an electrical neuron, is thus bounded (amplitude limited). An artificial neural net is a collection of such electrical neurons connected in different topology. The most common application of an artificial neural net is in machine learning. In a learning problem, the weights and/or non-linearities in an artificial neural net undergo an adaptation cycle. The adaptation cycle is required for updating these parameters of the network, until a state of equilibrium is reached, following which the parameters no longer change further. The ANN support both supervised and unsupervised types of machine learning. The supervised learning algorithms realized with ANN have been successfully applied in ***control, automation, robotics and computer vision***. The unsupervised learning algorithms built with ANN, on the other hand, have been applied in ***scheduling, knowledge acquisition, planning and analog to digital conversion of data.***

Genetic Algorithms: A genetic algorithm (GA) is a stochastic algorithm that mimics the natural process of biological evolution. It follows the principle of Darwinism, which rests on the fundamental belief of the "survival of the fittest" in the process of natural selection of species. GAs find extensive applications in intelligent search, machine learning and optimization problems. The problem states in a GA are denoted by chromosomes, which are usually represented by binary strings. The most common operators used in GA are crossover and mutation. The processes of execution of crossover and mutation are illustrated in **fig.(1.4)** and **fig.(1.5)** respectively.

The evolutionary cycle in a GA consists of the following three sequential steps.

(a) Generation of population (problem states represented by chromosomes).

(b) Genetic evolution through crossover followed by mutation.

(c) Selection of better candidate states from the generated population.

In step (a) of the above cycle, a few initial problem states are first identified. The step (b) evolves new chromosomes through the process of crossover and mutation. In step (c) a fixed number of better candidate states are selected from the generated population. The above steps are repeated a finite number of times for obtaining a solution for the given problem.

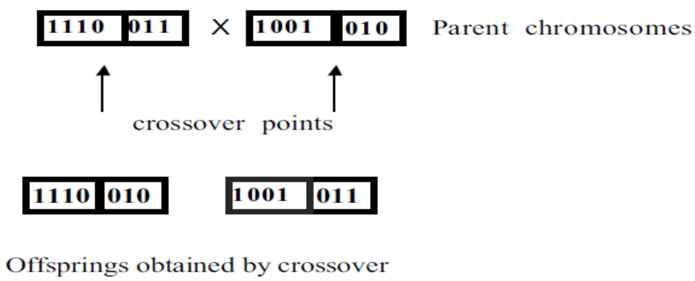

Fig.(1.4): Exchange of genetic information by crossover operation.

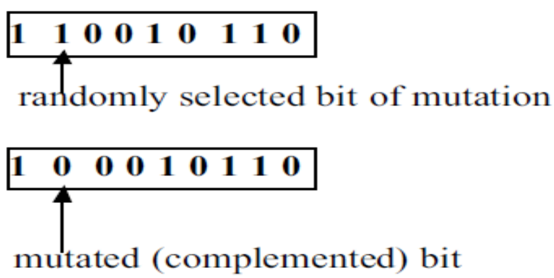

Fig.(1.5): The mutation operation: randomly selected bits are complemented.

[7] Management of Imprecision and Uncertainty: Data and knowledge bases in many typical AI problems, such as reasoning and planning, are often contaminated with various forms of incompleteness. The incompleteness of data, hereafter called **imprecision**, generally appears in the database for: i) **Lack of appropriate data**. ii) **Poor authenticity level of the sources**.

The incompleteness of knowledge, often referred to as **uncertainty**, originates in the knowledge base due to lack of certainty of the pieces of knowledge. Reasoning in the presence of imprecision of data and uncertainty of knowledge is a complex problem. Various tools and techniques have been devised for reasoning under incomplete data and knowledge. Some of these techniques employ: i) **stochastic** ii) **fuzzy** and iii) **belief network** models.

In a stochastic reasoning model, the system can have transition from one given state to a number of states, such that the sum of the probability of transition to the next states from the given state is strictly unity. In a fuzzy reasoning system, the sum of the membership value of transition from the given state to the next state may be greater than or equal to one. The belief network model updates the stochastic/fuzzy belief assigned to the facts embedded in the network until a condition of equilibrium is reached, following which there would be no more change in beliefs. Recently, fuzzy tools and techniques have been applied in a specialized belief network, called a **fuzzy Petri net,** for handling both imprecision of data and uncertainty of knowledge by a unified approach.