Unit 1 Introduction to Artificial Intelligence (AI)

Structure:

1.1 Introduction

Now you will be introduced to the world of Artificial Intelligence (AI). Artificial Intelligence (AI) is one of the newest disciplines, formally initiated in 1956 when the name was coined. However, the study of intelligence is one of the oldest disciplines being approximately 2000 years old. The advent of computers made it possible for the first time for people to test models they proposed for learning, reasoning, perceiving, etc.

Artificial Intelligence is the area of computer science focusing on creating machines that can engage on behaviors that humans consider intelligent. The ability to create intelligent machines has intrigued humans since ancient times and today with the advent of the computer and 50 years of research into AI programming techniques, the dream of smart machines is becoming a reality. Researchers are creating systems which can mimic human thought, understand speech, beat the best human chess player, and countless other feats never before possible. Find out how the military is applying AI logic to its hi-tech systems, and how in the near future Artificial Intelligence may impact our lives. In this unit you will be introduced to the definition and fundamentals of artificial intelligence. You will study the types of AI tasks, brief history of AI and Importance of AI.

Objectives:

After studying this unit, you should be able to

- define Artificial Intelligence (AI)
- list the modern founders of Artificial Intelligence

- list the advantages of symbolic and non-symbolic representations
- list the types of AI tasks
- explain the history and importance of AI

1.2 Concept of Artificial Intelligence (AI)

Artificial Intelligence (AI), as we know it today, is a relatively new field. Even though some groundwork had been laid earlier, AI began in earnest with the emergence of the modern computer during the 1940s and the 1950s. It was the ability of these new electronic machines to store large amounts of information and process it at very high speeds that gave researchers the vision of building systems which could emulate some human ability.

During the last fifty years, we have witnessed the realization of many of these early researchers vision. We have seen computer systems shrink in size and cost by several orders of magnitude. We have seen memories increase in storage capacity to the point where they equal a significant fraction of the human brain's storage capacity. We have seen the speed and reliability of systems improve dramatically. And, we have seen the introduction of many impressive software tools.

Given these new hardware and software systems and our improved understanding of Homo sapiens, we are on the threshold of witnessing a whole new horizon of exciting innovations. During the last few decades, we have witnessed the introduction of many intelligent computer systems never dreamed of by the early visionaries. We will see the introduction of systems which equal or exceed human abilities, and see them become an important part of most business and government operations as well as our daily activities.

Definitions of AI

As a beginning the following definition can be offered:

AI is the branch of computer science concerned with the study and creation of computer systems that exhibit some form of intelligence: systems that learn new concepts and tasks, systems that can reason and draw useful conclusions about the world around us, systems that can understand a natural language or perceive and comprehend a visual scene, and systems that perform other types of feats that require human types of intelligence.

An understanding of AI requires an understanding of related terms such as intelligence, knowledge, reasoning, thought, cognition, learning, and a number of computer-related terms (refer figure 1.1). Dictionaries define intelligence as the ability to acquire, understand and apply knowledge, or the ability to exercise thought and reason. Of course, intelligence is more than this.

Intelligence is the integrated sum of those feats which gives us the ability to remember a face not seen for thirty or more years, or to build and send rockets to the moon. The food for this intelligence is knowledge. Knowledge comes from processed information which at the root level comes from the data.

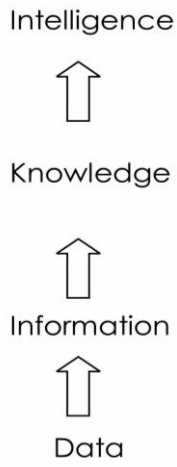

Figure 1.1: Relationship between Data and Intelligence

Another Definition of AI can be stated as:

AI is the study of how to make computers do things which, at the moment, people do better.

This definition is, of course, somewhat ephemeral because of its reference to the current state of computer science. And it fails to include some areas of potentially very large impact, namely problems that cannot now be solved well either by computers or people. But it provides a good outline of what constitutes AI, and it avoids the philosophical issues that dominate attempts to define the meaning of either artificial or intelligence.

A good general definition of AI could be:

AI is the part of computer science concerned with designing intelligent computer systems, that is, computer systems that exhibit the characteristics we associate with intelligence in human behavior – understanding language, learning, reasoning and solving problems.

What AI is not?

1) AI is not the study and creation of conventional computer systems.
2) AI is neither the study of the mind, nor of the body, nor of languages, as customarily found in the fields of psychology, physiology, cognitive science, or linguistics.

Ultimate goal of AI

The primary goal of AI research is to increase our understanding of perceptual, reasoning, learning, linguistic and creative processes. This understanding is no doubt helpful in the design and construction of useful new tools in science, industry, and culture; Just as the invention of the internal combustion engine and the development of machines like airplanes resulted in unprecedented enhancement of the mobility of our species, the tools resulting from AI research are already beginning to extend human intellectual and creative capabilities in ways that our predecessors could only dream about. Sophisticated understanding of the underlying mechanisms and the potential and limits of human as well as other forms of intelligence is also likely to shed new lights on the social, environmental, and cultural problems of our time and aid the search for solutions. So the ultimate goal of AI is to develop working computer systems that are truly capable of performing some tasks that require high level of intelligence.

Practical AI

AI, in short, is about the design and implementation of intelligent agents. This requires the use of AI tools and techniques for search, knowledge representation, and adaptation and learning and their application to problem-solving, planning, analysis, design, knowledge acquisition, discovery, etc. Some would argue that much of contemporary AI work essentially involves reducing problems requiring intelligence into search problems using appropriate ways of representing the knowledge necessary for the solution of such problems. A side-effect of this is a wide range of practically useful tools (theorem-provers, game-players, vision programs,

natural language interfaces, stock-market analysts, programmer's assistants, assembly line robots, physician's assistants, tutoring programs, architect's assistants, internet softbots, and so on).

Relation of AI to other disciplines

The invention of digital (and analog) computers in the 1940s and 1950s and the work in the theory of computation, information theory, and control that accompanied it provided the experimental tools and the theoretical underpinnings of AI research. Much related work has taken place in related fields addressing similar questions (e.g., bionics, cybernetics, neural networks, statistical pattern recognition, syntactic pattern recognition, expert systems, computer vision, robotics, computational linguistics, decision theory, cognitive psychology, artificial life, computational neuroscience, computational organization theory, etc.). AI, broadly interpreted, is closely intertwined with, and often subsumes, much of the work in most of these fields.

AI is often regarded as a branch of computer science. AI's special relationship with computer science is due to the fact that the language of computation is to the study of mind what calculus was to the study of physics. Calculus provided the mathematical tools for formulation of questions and search for answers in classical physics (It should come as no surprise that Newton and Leibnitz were among the inventors of calculus). But physics is more than calculus; it developed its own armamentarium of experimental methods to probe its domain | the physical universe. AI (and more recently, cognitive science) continue to develop their own experimental tools and theoretical frameworks. In the process, AI has contributed over the years, a wide variety of concepts and tools to Computer Science- LISP - one of the earliest high-level programming languages, the first multi-tasking operating system, logic programming, constraint programming, heuristic search, object-oriented programming, neural networks, computational learning theory, temporal logic, deductive databases, high-dimensional grammars, evolutionary programming | to name a few. AI problems have stimulated research in other areas of computer science | massively parallel architectures for vision, theoretical research in complexity of reasoning and learning, and so on. AI is occasionally viewed as a sibling of psychology. Psychology is concerned with the formulation and experimental verification of theories of behavior

with human or animal subjects. AI is concerned with computational models that exhibit aspects of intelligent behavior. It is not generally committed to any particular (e.g., human-like) set of mechanisms or any particular ways of implementing the chosen mechanisms. Yet, the information processing models coming out of AI research have strongly influenced contemporary research in human and animal psychology and neuroscience.

Insofar as intelligent behavior is normally associated with living systems, AI shares some of the concerns of the field of study that has been provocatively, and misleadingly, labeled artificial life.

AI can also be thought of as applied epistemology (the branch of philosophy that is concerned with the nature of knowledge). AI research has brought to light entirely new questions and new ways of looking at old problems in epistemology.

AI is often treated as a branch of engineering that is concerned with the design, implementation, and evaluation of intelligent artifacts. AI research has resulted in a number of useful practical tools (programs that configure computer systems, diagnose faults in engines, software agents that scour the Internet for information on demand, etc.).

AI attacks a long-standing mix of problems from a number of more established disciplines like philosophy, psychology, linguistics, anthropology, engineering, and neuroscience. While freely borrowing from these disciplines, it brings to the study of intelligent behavior, a unique approach, and a unique set of tools and in the process, sometimes raises entirely new questions due to its use of computation as a substrate for theory-construction and experimentation. This has led to arguably one of the most important scientific developments of this century, the birth of Cognitive Science (which attempts to integrate insights and results from its constituent disciplines better than most (though by no means all) of the work in AI). All of this gives us a new perspective on some of the long-standing questions about the nature of mind. But it does not make the questions themselves necessarily any easier! Every discipline has a domain of enquiry. For AI, it is the entire range of human and non-human intellectual enterprise spanning the entire space of actual and possible intelligent adaptive systems. As a result, AI gets deeply involved in the conceptual and methodological questions in any area in which it is applied: The use of AI in synthesis of

artistic objects (e.g., drawings and paintings) necessarily has to involve an understanding of the specification of ways of representing the knowledge used by an artist as well as theories about creativity in the domain of art; the use of AI tools to model the process of scientific exploration in some area (say molecular biology) necessarily entails an understanding of the scientific method and is likely to yield new insights on hypothesis formation, experimental design, and theory selection in that area. As a consequence, AI is one of the most interdisciplinary fields of study currently taught in our universities.

Modern founders of AI

The following list gives you the Modern founders of AI

- Alan Turing ("Computing Machinery and Intelligence"; Turing test)
- McCulloch & Pitts (neural nets)
- Norbert Wiener (cybernetics)
- John von Neumann (game theory)
- Claude Shannon (information theory)
- Newell & Simon (The Logic Theorist)
- John McCarthy (LISP, common sense reasoning)
- Marvin Minsky (Frames)
- Donald Michie (Freddy)

Achievements of AI

Let us have a look at the Achievements of AI

- Deep Thought is an international grand master chess player.
- Sphinx can recognize continuous speech without training for each speaker. It operates in near real time using a vocabulary of 1000 words and has 94% word accuracy.
- Navlab is a truck that can drive along a road at 55 kmph in normal traffic.
- Carlton and United Breweries use an AI planning system to plan production of their beer.
- Robots are used regularly in manufacturing.
- Natural language interfaces to databases can be obtained on a PC.
- Machine Learning methods have been used to build expert systems.
- Expert systems are used regularly in finance, medicine, manufacturing, and agriculture

How can we make computer systems more intelligent?

- *Autonomy* to perform tasks that currently require human operators without human intervention or monitoring.
- *Flexibility* in dealing with variability in the environment.
- *Ease of use*: computers that are able to understand what the user wants from limited instructions in natural languages.
- *Learning* from experience.

How do we build an intelligent agent?

- Must be able to perceive its environment.
- Must be able to affect its environment.
- Must be able to reason about observations and actions.
- Must be able to learn from observations and actions.
- Must have goals.

Symbolic representations

To construct intelligent systems it is necessary to employ internal representations of a symbolic nature, with cognitive activity corresponding to computational manipulation of these symbolic representations. The symbolic representations must refer to the external world. The figure 1.2 shows the Symbolic Representation.

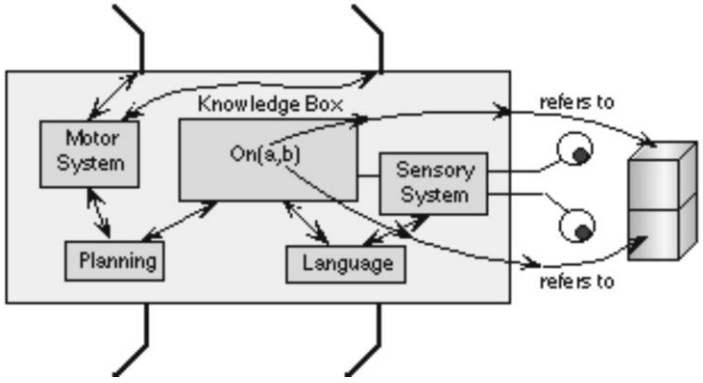

Figure 1.2: Symbolic Representation

Advantages of symbolic representation:

The following are the advantages of symbolic representation.

- The system builder can read what the system knows.
- Knowledge is represented by sentences in a formal language.
- It is possible to read the representation and understand the meaning of the knowledge.

Non-symbolic representations

Knowledge is represented by weights on connections in a network as shown in figure 1.3.

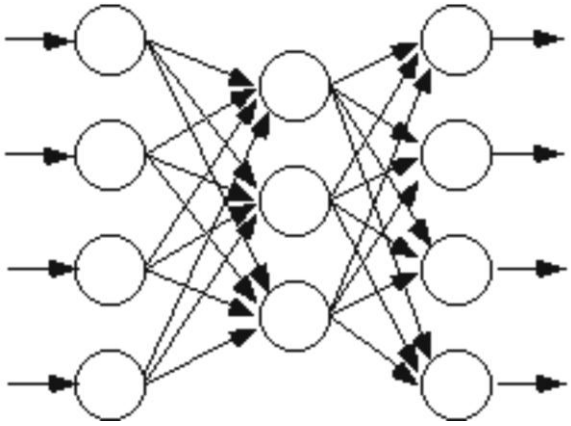

Figure 1.3: Non-symbolic Representation

Advantages of non-symbolic representation:

The Advantages of non-symbolic representation are:

- Can deal with combinations of attributes such as an image.
- Noise tolerant

Self Assessment Questions

1. Artificial Intelligence (AI) is one of the newest disciplines, formally initiated in _____ when the name was coined.
2. AI is the part of computer science concerned with designing _____ computer systems.
3. _____ are used regularly in manufacturing.
4. Expert systems are not used regularly in finance, medicine (True/False?)

1.3 Types of AI Tasks

We can classify AI tasks is into 3 classes. They are: Mundane tasks, Formal tasks and Expert tasks.

Mundane tasks

- Perception
- Vision
- Speech
- Natural Language understanding, generation and translation

- Common-sense Reasoning
- Simple reasoning and logical symbol manipulation
- Robot Control

Formal tasks

- Games
  - Chess
    Deep Blue recently beat Gary Kasparov
  - Backgammon
  - Draughts
  - GO

To solve these problems we must explore a large number of solutions quickly and choose the Best One.

- Mathematics
  - Geometry and Logic
    Logic Theorist: It proved mathematical theorems. It actually proved several theorems from Classical Math Textbooks
  - Integral Calculus
    Programs such as Mathematical and Mathcad and perform complicated symbolic integration and differentiation.
- Proving Properties of Programs e.g. correctness

Manipulate Symbols and reduce problem (usually recursively), until the answer is obvious. That is, it can be looked up in a table.

Expert tasks

- Engineering
  - Design
  - Fault finding
  - Manufacturing
- Planning
- Scientific Analysis
- Medical Diagnosis
- Financial Analysis

1.4 A brief History of AI

Now you will be introduced to the History of AI.

Work before 1950

AI began to emerge as a separate field of study during 1940s and 1950s when the computer became a commercial reality. Prior to this time, a number of important areas of research that would later shape early work on AI were beginning to mature.

Firstly, during 1920s and 1930s there was the work of logicians such as Alonzo Church, Kurt Godel. It helped to produce formalized methods for reasoning, the form of logic known as propositional and predicate calculus. It demonstrated that facts and ideas from a language such as English could be formally described and manipulated mechanically in meaningful ways. Turing, sometimes regarded as the father of AI, also demonstrated, as early as 1936, that a simple computer processor (later named as Turing machine) could manipulate symbols as well as numbers.

During 1940s and 1950s, Norbert Wiener, coined the term Cybernetics. Cybernetics is the study of communication in human and machine and it brought together many parallels between human and machine. Cybernetics combines the concepts from information theory, feedback control systems and electronic computers. Finally, during the 1950s, the electronic stored program digital computer became a commercial reality. This followed several years of prototype systems including the Mark I Harvard relay computer (1944), the university of Pennsylvania Moore School of Electrical Engineering's ENIAC electronic computer (1947), and subsequent development of the Aberdeen Providing Ground's EDVAC and Sperry-Rand's UNIVAC.

Other important developments during this early period which helped to launch AI include the introduction of Boolean algebra, switching theory, and even statistical decision theory.

Work after 1950

Although the computer provided the technology necessary for AI, it was not until the early 1950's that the link between human intelligence and machines was really observed. Norbert Wiener was one of the first Americans to make observations on the principle of feedback theory. The most familiar example of feedback theory is the thermostat: It controls the temperature of an

environment by gathering the actual temperature of the house, comparing it to the desired temperature, and responding by turning the heat up or down. What was so important about his research into feedback loops was that Wiener theorized that all intelligent behavior was the result of feedback mechanisms. Mechanisms that could possibly be simulated by machines. This discovery influenced much of the early development of AI.

In late 1955, Newell and Simon developed *The Logic Theorist*, considered by many to be the first AI program. The program, representing each problem as a tree model, would attempt to solve it by selecting the branch that would most likely result in the correct conclusion. The impact that the logic theorist made on both the public and the field of AI has made it a crucial stepping stone in developing the AI field.

In 1956, John McCarthy organized a conference to draw the talent and expertise of others interested in machine intelligence for a month of brainstorming. He invited them to Vermont for "The Dartmouth summer research project on artificial intelligence." From that point on, because of McCarthy, the field would be known as Artificial intelligence. Although not a huge success, the Dartmouth conference did bring together the founders in AI, and served to lay the groundwork for the future of AI research.

Some significant AI events after 1950 include the following:

1950: Alan Turing publishes, "Computing Machinery and Intelligence."

1956: John McCarthy coins the term, "Artificial Intelligence" at a Dartmouth computer conference.

1956: Demonstration of the first running AI program at Carnegie Mellon University.

1958: John McCarthy invents the Lisp language, an AI programming language, at Massachusetts Institute of Technology (MIT).

1964: Danny Bobrow shows that computers can understand natural language enough to solve algebra word programs (MIT).

1965: Joseph Weizenbaum builds ELIZA, an interactive program that carries on a dialogue in English on any topic (MIT).

1969: Shakey, a robot, combines locomotion, perception and problem solving (Stanford Research Institute).

1979: The first computer-controlled autonomous vehicle, the Stanford Cart, is built.

1983: Danny Hillis co-founds Thinking Machines, the first company to produce massively parallel computers.

1985: The drawing program, Aaron, created by Harold Cohen, is demonstrated at AI conference.

1990s: Major advances in all areas of AI. Significant demonstrations in machine learning, intelligent tutoring, case-based reasoning, multi-agent planning, scheduling, uncertain reasoning, data mining, natural landscape understanding and translation, vision, virtual reality and games.

1997: IBM computer Deep Blue beats world champion Garry Kasparov in chess match.

2000: Interactive robot pets become commercially available. MIT displays Kismet, a robot with a face that expresses emotions. Carnegie Mellon robot Nomad explores remote regions of Antarctica and locates meteorites.

Self Assessment Questions

5. During 1940s and 1950s, Norbert Wiener, coined the term _____.
6. In late_____, Newell and Simon developed *The Logic Theorist*, considered by many to be the first AI program.
7. In _____ the first computer-controlled autonomous vehicle, the Stanford Cart, is built.

1.5 Importance of AI

AI may be one of the most important developments of this century. It will affect the lives of most individuals in civilized countries. And countries leading in the development of AI by then will emerge as the dominant economic powers of the world.

The Japanese launched a very ambitious program in AI research and development. Known as the fifth generation, this plan was officially announced in October 1981. It calls for the implementation of a ten-year plan to develop intelligent supercomputers. It is a cooperative effort between government and private companies having an interest in the manufacture of computer products, robotics, and related fields. With the combined budget of

about one billion dollars, the Japanese are determined they will realize many of their goals, namely, to produce systems that can converse in a natural language, understand speech and visual scenes, learn and refine their knowledge, make decisions, and exhibit other human traits.

The British initiated Alvey Project but their goals were not as ambitious as the Japanese but are set to help the British keep abreast and remain in the race. The European Common Market countries have jointly initiated a separate cooperative plan named the ESPRIT program. Other countries including Canada, Italy, Austria and Singapore have made some commitment in funded research and development.

In the United States, firstly, there was the formation of a consortium of private companies in 1983 to develop advanced technologies that apply AI techniques (like VLSI). The consortium is known as the Microelectronics and Computer Technology Corporation (MCC) and is headquartered in Austin, Texas. Secondly, the Department of Defense Advanced Research Project Agency (DARPA) has increased its funding for AI, including development support in three significant programs:

1) Development of an autonomous land vehicle (ALV) (a driverless military vehicle),
2) Development of a pilot's associate (an expert system which provides assistance to fighter pilots), and
3) The Strategic Computing Program (an AI based military supercomputer project).

In addition IBM, DEC, AT&T, HP, Texas Instruments, and Xerox have their own research program in AI.

Self Assessment Questions

8. The European Common Market countries have jointly initiated a separate cooperative plan named the _____ program.
9. ALV stands for _____.
10. Which consortium was formed to develop advanced technologies that apply AI techniques?

1.6 Summary

In this unit, we learnt the various definitions of Artificial Intelligence (AI) and types of Artificial Intelligence tasks. Also, we discussed history and

importance of AI. Artificial Intelligence is the area of computer science focusing on creating machines that can engage on behaviors that humans consider intelligent. AI began in earnest with the emergence of the modern computer during the 1940s and the 1950s. AI is not the study and creation of conventional computer systems. AI tasks are classified into 3 classes: Mundane tasks, Formal tasks and Expert tasks. In the year 1956 demonstration of the first running AI program was done at Carnegie Mellon University. AI may be one of the most important developments of this century. In the United States, firstly, there was the formation of a consortium of private companies in 1983 to develop advanced technologies that apply AI techniques (like VLSI).

1.7 Terminal Questions

1. What is artificial intelligence?
2. Name the Modern Founders of AI.
3. List the advantages of symbolic and non-symbolic representations.
4. Give the classification of different types of tasks of AI.
5. Give a brief history of AI starting from early 1950 to till date.

1.8 Answers

Self Assessment Questions

1. 1956
2. intelligent
3. Robots
4. False
5. Cybernetics
6. 1955
7. 1979
8. ESPRIT
9. Autonomous Land Vehicle
10. Microelectronics and Computer Technology Corporation (MCC)

Terminal Questions

1. AI is the branch of computer science concerned with the study and creation of computer systems that exhibit some form of intelligence (Refer section 1.2 for detail)

2. Alan Turing ("Computing Machinery and Intelligence"; Turing test), McCulloch & Pitts (neural nets) (Refer section 1.2)
3. Advantages of symbolic representation:
   - The system builder can read what the system knows.

   Advantages of Non-symbolic representation:
   - Can deal with combinations of attributes such as an image (Refer section 1.2).
4. One possible classification of AI tasks is into 3 classes: Mundane tasks, Formal tasks and Expert tasks. (Refer section 1.3).
5. AI began to emerge as a separate field of study during 1940s and 1950s when the computer became a commercial reality. (Refer section 1.4)