
UNIT 1 INTRODUCTION TO ARTIFICIAL INTELLIGENCE

Structure

- 1.1 Introduction
- 1.2 Objectives
- 1.3 Basics of Artificial Intelligence (AI)?
- 1.4 Brief history of Artificial Intelligence
- 1.5 Components of Intelligence
- 1.6 Approaches to Artificial Intelligence
- 1.7 Comparison between Artificial Intelligence (AI), Machine Learning (ML) and DeepLearning (DL).
- 1.8 Application Areas of Artificial Intelligence Systems
- 1.9 Intelligent Agents
 - 1.9.1 Stimulus - Response Agents
- 1.10 Summary
- 1.11 Solutions/Answers
- 1.12 Further Readings

1.1 INTRODUCTION

Today, artificial intelligence is used in a wide variety of applications, including engineering, technology, the military, opinion mining, sentiment analysis, and many more. It is also used in more advanced domains, such as language processing and applications for aerospace.

AI is everywhere in today's world, and people are gradually becoming accustomed to its presence. It is utilised in systems that recognise both voices and faces. In addition to this, it can provide you with shopping recommendations that are tailored to your own purchasing preferences. Finding spam and preventing fraudulent use of credit cards is made much easier when you have this skill. The most cutting-edge technology currently on the market are virtual assistants like Apple's Siri, Amazon's Alexa, Microsoft's Cortana, and Google's own Google Assistant. It's possible that you're already familiar with the technology involved in artificial intelligence (AI). Are you?

AI has become very popular all over the world today. It imitates human intelligence in machines by programming them to do the same things people do. As a technology, AI is going to have a bigger impact on how people live their daily lives. Everyone wants to connect to Artificial Intelligence as a technology these days. Before we can understand AI, we need to know and talk about some basic things. For example, what is the difference between knowledge and intelligence? The key to starting this unit is the answer to this question.

The accumulation of information and abilities that a person has gained through their life experiences is known as knowledge. While intelligence refers to one's capacity to put one's knowledge into practise. To put it simply, knowledge is what we have learned over the years, and it expands as time passes. Because of this, it represents the culmination of everything that we have realised over the course of our lives. It is important to highlight that having information does not automatically make one intelligent; rather, intelligence is what makes one smart.

There is a well-known proverb that says "marks are not the measure of intelligence." This is due to the fact that intelligence is not a measurement of how much information one possesses. In fact, it is the measure of how much we comprehend and put into practise. People who are knowledgeable may gain a lot of information, but an intelligent person understands how to comprehend, analyze, and use the information. You could have a lot of knowledge but still be the least intelligent person in the room. Knowledge and intelligence are inextricably linked, and each contributes to the other's development. Knowledge enables one to learn the understandings that others have of things, whereas intelligence is the foundation for one's ability to grasp the things themselves.

Now that we have an understanding of the distinction between intelligence and knowledge, our next issue is: what exactly is artificial intelligence? Incorporating intelligence into a machine is related to the field of Artificial Intelligence, whereas both concepts, namely the representation of knowledge and its engineering, are the basis of traditional AI research. This topic will be discussed in section 1.3 of this unit, but in a nutshell, incorporating intelligence into a machine is related to the field of Artificial Intelligence. Knowledge engineering is a subfield of artificial intelligence (AI) that applies rules to data in order to simulate the way in which an expert would think about the information. It does this by analysing the structure of a job or a decision in order to figure out how one arrives at a conclusion.

The subsequent units of this course, you will learn about some of the concepts that are essential for knowledge representation, such as frames, scripts, and other related topics. In addition, this course will address the issues that are associated with the knowledge representation for uncertain situations, such as employing the method of fuzzy logic, rough sets, and the Dempster Shafer theory, among other relevant topics.

Some of the prerequisites to get started with this subject.

- a) Strong understanding of Basic concepts of Mathematics viz. Algebra, Calculus, probability and Statistics.
- b) Experience in programming using Python or Java.
- c) A good understanding of algorithms.
- d) Reasonably good data analytics skills.
- e) Fundamental knowledge in discrete mathematics.
- f) Finally, internal will to learn

1.2 OBJECTIVES

After going through this unit, you will be able to:

- Understand the difference between knowledge and intelligence
- Answer – what is AI?
- Identify various approaches of AI
- Compare Artificial Intelligence (AI), Machine Learning (ML) & Deep Learning (DL).
- Understand the concept of agents in AI

1.3 BASICS OF ARTIFICIAL INTELLIGENCE?

Knowledge and intelligence are two important concepts, and we were able to gain an understanding of the fundamental distinction between the two terms. Now that we have your attention, let's talk about what artificial intelligence actually is. The meaning of Artificial Intelligence will be covered in this section; however, before we get started, it is important to note that the field of Artificial Intelligence is related to the process of incorporating intelligence into machines; the specifics of this process, as well as the mechanism itself, will be covered in this unit as well as the subsequent units of this training.

The following is a list of eight definitions of artificial intelligence that have been provided by well-known authors of artificial intelligence textbooks.

- 1) According to Haugeland in 1985, "The Exciting New Effort to Make Computers Think... Machines with Minds, in the Full and Literal Sense,"
- 2) According to Bellman, "the automation of behaviours that we connect with human thinking, activities such as decision-making, problem-solving, and learning..." 1978
- 3) "The study of mental capabilities through the application of computer models," (also known as "The Study of Mental Capabilities"), Charniak and McDermott's 1985.
- 4) According to Winston (1992), "the study of the calculations that make it possible to perceive, reason, and act."
- 5) "The art of building machines that execute functions that demand intellect when performed by people," as defined by Kurzweil in the year 1990.
- 5) "The art of building machines that execute functions that demand intellect when performed by people," as defined by Kurzweil in the year 1990. To the Rich and the Knight, 1991
- 7) According to Schalkoff (1990), "a field of study that aims to explain and replicate intelligent behaviour in terms of computational processes."
- 8) According to Luger and Stubblefield (1993), "the discipline of computer science that is concerned with the automation of intelligent behaviour."

According to the concepts presented earlier, there are four distinct objectives that might be pursued in the field of artificial intelligence. These objectives are as follows:

- The creation of systems that think in the same way as people do.
- The creation of systems that are capable of logical thought.
- The creation of machines that can mimic human behaviour.
- The creation of systems that behave in a logical manner

In addition, we discovered through our conversation in the earlier section 1.1 of this Unit, that Artificial Intelligence (AI) is the intelligence that is incorporated into machines; in other words, AI is the ability of a machine to display human-like capabilities such as reasoning, learning, planning, and creativity. We learned this information because AI is the ability of a machine to display human-like capabilities. Taking into mind the Emerging AI technologies to sense, interpret, and act according to the circumstances, relevant exemplary solutions are summarised in Figure 1, which attempts to encapsulate the understanding of the question "What is Artificial Intelligence?"




Figure-1: What is Artificial Intelligence? – Emerging AI Technologies

The applications of artificial intelligence (AI) that are shown in figure-1 do not all require the same kinds of AI. In a general sense, one may say that AI can be categorised into the following levels:

- software level (i.e., Embodied AI). Where, the software level consists of things like search engines, virtual assistants, speech and facial recognition systems, picture analysis tools, and other things like that.
- Hardware level (Embedded) includes Robots, autonomous vehicles, drones, the Internet of Things, and other technologies fall under the category of embedded artificial intelligence (AI).

On the basis of the functionalities the AI can be classified based on Type 1 and Type 2




Figure 1(a) – Classification of Artificial Intelligence

Here's a brief introduction the first type of AI i.e., Type 1 AI. Following are the three stages of Type 1 - Artificial Intelligence:

- Artificial Narrow Intelligence-(ANI)
- Artificial General Intelligence-(AGI)
- Artificial Super Intelligence-(ASI)

Types of Artificial Intelligence




Figure 1(b) –Three Stages of Type-I Artificial Intelligence

The various categories of Artificial Intelligence are discussed as follows:

- Artificial Narrow Intelligence (ANI), also called Weak AI or Narrow AI: Weak AI is a term for thinking that is "simulated." Such systems seem to act intelligently, but they don't have any awareness of what they are doing. For example, a chatbot might talk to you in a way that seems natural, but it doesn't know who it is or why it's talking to you. Artificial intelligence is a system that was built to do a certain job.
- Artificial General Intelligence (AGI): Strong or General Artificial Intelligence, also called "actual" thinking. That is, acting like a smart human and thinking like one with a conscious, subjective mind. For instance, when two humans talk, they probably both know who they are, what they're doing, and why.

Systems with strong or general artificial intelligence can do things that people can do. These systems tend to be harder to understand and more complicated. They are set up to handle situations where they might need to solve problems on their own without help from a person. Uses for these kinds of systems include self-driving cars and operating rooms in hospitals.

- c) Artificial Super Intelligence (ASI) - Super intelligence: The term "super intelligence" usually refers to a level of general and strong AI that is smarter than humans, if that's even possible. The ASI is seen as the logical next step after the AGI because it can do more than humans can. This includes making decisions, making rational decisions, and even things like building emotional relationships. There is a marginal difference between AGI and ASI.

☞ Check Your Progress 1

- Q1 How Knowledge differs from intelligence? What do you understand by the term Artificial Intelligence (AI)? List the various technologies and their corresponding illustrative solutions.

.....

.....

.....

- Q2 Classify AI on the basis of the functionalities of AI

.....

.....

.....

- Q3 Compare ANI, AGI and ASI, in context of AI

.....

.....

.....

1.4 BRIEF HISTORY-ARTIFICIAL INTELLIGENCE

AI's ideas come from early research into how people learn and think. Also very old is the idea that a computer could act like a person. Greek mythology is where the idea of machines that can think for themselves comes from.

- Aristotle, who lived from 384 BC to 322 BC, made a syllogistic logic system that was not formal. This is where the first formal system of deductive reasoning got its start.

At the start of the 17th century, Descartes said that animal bodies are just complex machines.

- Pascal made the first mechanical digital calculator in the year 1642.

In the 1800s, George Boole came up with a number system called "binary algebra" that showed (some) "laws of thought."

- Charles Babbage and Ada Byron worked on programmable mechanical calculators.

In the late 19th century and early 20th century, mathematicians and philosophers like Gottlob Frege, Bertram Russell, Alfred North Whitehead, and Kurt Godel built on Boole's first ideas about logic to make mathematical representations of logic problems.

When electronic computers came along, it was a big step forward in how we could study intelligence.

McCulloch and Pitts made a Boolean circuit model of the brain in 1943. They wrote about how neural networks can be used to do math in the paper "A Logical Calculus of Ideas Immanent in Nervous Activity."

- In 1950, Turing wrote a paper called "Computing Machines and Intelligence." This article gave a good overall picture of AI. To learn more about Alan Turing, go to <http://www.turing.org.uk/turing>.

Turing's paper talked about a lot of things, one of which was how to solve problems by using heuristics as guides to look through the space of possible solutions. He used the game of chess to explain how his ideas about how machines can think work. He even said that the machine could change its own instructions so that machines could learn from what they do.

The SNARC was built by Marvin Minsky and Dean Edmonds in 1951. It was the first randomly wired neural network learning machine (SNARC stands for Stochastic Neural Analog Reinforcement Computer). It was a computer with a network of 40 neurons and 3000 vacuum tubes.

Samuel made a number of programmes to help people play checkers between 1952 and 1956.

In 1956, Dartmouth was the site of a well-known meeting. At the conference, the people who came up with the idea of AI met for the first time. At this meeting, the name "Artificial Intelligence" was chosen.

- Newell and Simon's book *The Logic Theorist* came out. Many people think it was the first show to use artificial intelligence.

In 1959, Gelernter made a Geometry Engine. In 1961, James Slagle's Ph.D. dissertation at MIT was a programme called SAINT. It was written in LISP, and a first-year college student could use it to solve calculus problems.

Thomas Evan made a programme called "Analogy" in 1963 that could solve analogy problems like those on an IQ test. The first collection of articles about artificial intelligence was put together by Edward A. Feigenbaum and Julian Feldman. It was called "Computers and Thought." It was released in 1963.

In 1965, J. Allen Robinson came up with a way to prove things mechanically. He called it the Resolution Method. This made it possible for formal logic to work well as a language for representing programmes. In 1967, Feigenbaum,

Lederberg, Buchanan, and Sutherland at Stanford showed how the Dendral programme could be used to understand the mass spectra of organic chemical compounds. This was the first programme that worked well and was based on scientific knowledge. The SRI robot Shakey showed in 1969 that it was possible to move, see, and solve problems all at the same time.

From 1969 to 1979, the first systems based on knowledge were set up in place.

- In 1974, MYCIN showed how powerful rule-based systems can be for representing and drawing conclusions about knowledge in medical diagnosis and treatment. For the Knowledge Representation Version 2 CSE IIT, Kharagpur, plans were made. There were also some frames that Minski had made. There are logic-based programming languages like Prolog and Planner. In the 1980s, Lisp Machines was made and sold. In 1985, neural networks were once again all the rage. In 1988, probabilistic and decision-theoretic methods were used again.

Early AI was based on general systems that didn't know much. AI researchers realised that for machines to be able to reason about complex tasks, they need to know a lot about a narrow field.

Dean Pomerleau made ALVINN at CMU in 1989. Autonomous Land Vehicle in a Neural Network is what ALVINN stands for. This is a system that learns to drive by watching someone else do it. It has a neural network that gets an image from a two-dimensional camera that is 30x32 units. The output layer tells the vehicle where it needs to go. The system drove a car from the East Coast to the West Coast of the United States, which is about 2,850 miles. A person only drove about 50 of these miles. The system took care of the rest.

In the 1990s, AI made a lot of progress, especially in machine learning, data mining, intelligent tutoring, case-based reasoning, multi-agent planning and scheduling, uncertain reasoning, understanding and translating natural language, vision, virtual reality, games, and other areas.

Rod Brooks' COG Project at MIT made a lot of progress toward making a humanoid robot with the help of a lot of people.

In the 1990s,

- 1997 was the year of the first official Robo-Cup soccer game. It was played on a tabletop with 40 teams of robots talking to each other.
- As more and more people use the web, web crawlers and other AI-based programmes that pull information from it are becoming more and more important.
- Deep Blue In 1997, Gary Kasparov, who was the world champion at the time, lost to IBM's Deep Blue chess programme.

In 2000,

- The Nomad robot goes to remote parts of Antarctica to look for meteorite samples.
- Space probes that are made of robots can work on their own to learn more

about space. They keep an eye on what's going on around them, make decisions, and take action to get where they want to go. In April 2004, the first three-month missions of NASA's Mars rovers went well. The Spirit rover was looking at a group of hills on Mars that took two months to reach. It is finding strangely eroded rocks that might be new pieces of the puzzle that is the history of the area. Spirit's twin sister, Opportunity, was looking at the layers of rock in a crater.

- Internet agents: As the Internet grows quickly, more people want to use Internet agents to keep track of what users are doing, find the information they need, and figure out which information is the most useful. The reader can learn more about AI by reading about it in the news.

1.5 COMPONENTS OF INTELLIGENCE

According to the dominant school of thought in psychology, human intelligence should not be viewed as a singular talent or cognitive process but rather as a collection of distinct components. The majority of attention in the field of artificial intelligence research has been paid to the following aspects of intelligence: learning, reasoning, problem-solving, perception, and language comprehension.

Learning: There are numerous approaches to develop a learning system. Making mistakes is the simplest way to learn. A basic software that solves "mate in one" chess issues, for example, might test different moves until it finds one that answers the problem. The programme remembers which move worked so that the next time the computer is given the identical situation, it can provide an immediate response. The simple act of memorising things like answers to problems, words in a vocabulary list, and so on is known as "rote learning" or memorization.

We'll talk about another classification that doesn't depend on the way knowledge is represented or how it is represented. According to this system, there are five ways to learn:(ii)

- (i) Rote Learning or memorising.
- (ii) Learning by Instructions
- (iii) Learning by analogy.
- (iv) Learning by Induction
- (v) Learning by deduction.

Rote Learning: Rote learning is the simplest method of learning since it involves the least amount of interpretation. The information is simply copied into a database in this method of learning. This is the technique for memorising multiplication tables.

On a computer, rote learning is relatively simple to implement. The difficulty of implementing what is known as generalisation is more difficult. Generalized learning allows the student to perform better in situations they haven't faced before. A programme that learns the past tenses of regular English verbs by rote will not be able to form the past tense of "jump" until it has been presented

with "jumped" at least once, whereas a programme that can generalise from examples will be able to learn the "added" rule and thus form the past tense of "jump" even if it has never encountered this verb before. Modern techniques allow programmes to generalise complex rules based on data.

Learning by Instructions: The next method of learning is to be instructed. Because new knowledge must be contributed to an existing knowledge base in order to be useful, this type of learning necessitates greater inference. When a teacher instructs a student, this is the type of learning that occurs.

Learning by analogy: When you learn by analogy, you generate new ideas by connecting previously learned concepts. Textbooks frequently employ this method of instruction. For example, in the text, some problems are solved as examples, and students are subsequently given problems that are comparable to the examples. This type of learning also occurs when someone who can drive a light car attempts to drive a heavy vehicle.

Learning by Induction: Learning through induction is the most common method of learning. This is a method of learning that employs inductive reasoning, which is a style of reasoning that involves drawing a conclusion from a large number of good instances. If we encounter a lot of cows, we might notice that they have four legs, are white, and have two horns in the same position on their head, for example. Even though inductive reasoning frequently leads to valid conclusions, the conclusions are not always unarguable. For example, with the above-mentioned concept of cow, we might come across a black cow, a three-legged cow who has lost one leg in an accident, or a single-horn cow.

Learning by deduction: Finally, we discuss deductive learning, which is founded on deductive inference, a non-debatable mode of thinking. By irrefutable method of reasoning, we mean that if the hypotheses (or given facts) are accurate, the conclusion arrived through deductive (i.e., any irrefutable) reasoning is always correct. This is the most common method of thinking in mathematics.

Inductive learning is a crucial component of an agent's learning architecture. An agent learns based on:

- What it is learning, such as concepts, problem-solving techniques, or game-playing techniques, etc.
- The representation, predicate calculus, frame, script, and other elements that were employed.
- The critic, who expresses their opinion of the agency in general.

Learning based on feedback is normally categorized as:

- Supervised learning
- Unsupervised learning
- Reinforcement Learning.

Supervised Learning: It has a function that learns from inputs and outputs that are shown as examples. Some examples of this kind of learning are figuring

out useful things about the world from what you see, making a map from the current state's conditions to actions, and learning how the world changes over time.

Unsupervised Learning: There is no way to know what the inputs are and what the expected outputs are in this type of learning. So, the learning system has to figure out on its own which properties of objects it doesn't know about are important. For example, figuring out the shortest way to get from one city to another in a country you know nothing about.

Reinforcement (Rewards) for Learning: In some problems, the task or problem can only be seen, not said. Also, the job may be an ongoing one. The user tells the agent how happy or unhappy he or she is with the agent's work by sometimes giving the agent positive or negative rewards (i.e., reinforcements). The agent's job is to get as many rewards (or reinforcements) as possible. In a simple goal-attainment problem, the agent can be rewarded when it reaches the goal and punished when it doesn't.

You need an action plan to get the most out of this kind of task. But when it comes to tasks that never end, the future reward might be endless, making it hard to decide how to get the most out of it. One way to move forward in this kind of situation is to ignore future rewards after a certain point. That is, the agent may want rewards that will come soon more than rewards that will come a long time from now.

Delayed-reinforcement learning is the process of determining how to behave in situations when rewards are contingent on previous actions.

Reasoning: To reason is to draw conclusions that are appropriate for the situation. Both deductive and inductive reasoning can be used to make conclusions. An example of a deductive inference is, "Fred is either in the museum or in the cafe. He isn't in the cafe, so he must be in the museum." An example of inductive inference is, "In the past, accidents just like this one have been caused by instrument failure." The difference between the two is that in the deductive case, the truth of the premises guarantees the truth of the conclusion, while in the inductive case, the truth of the premises supports the conclusion that instrument failure caused the accident, but more research could show that the conclusion is actually false, even though the premises are true.

Programming computers to make inferences, especially deductive inferences, has had a lot of success. But you can't say that a programme can reason just because it can draw conclusions. To reason, you have to draw conclusions that make sense for the task or situation at hand. Giving computers the ability to tell what is important and what isn't is one of the hardest problems AI has to face.

Problem-solving: Problems usually go like this: given these data, find x. AI is used to solve a very wide range of problems. Some examples are finding the best way to win a board game, figuring out who someone is from a picture, and planning a series of steps that will allow a robot to do a certain task.

Methods for solving problems can be either specific or general. A special-purpose method is made to solve a specific problem and often takes advantage of very specific parts of the situation where the problem is happening. A general

method can be used to solve many different kinds of problems. The difference between the current state and the goal state can be reduced step by step with means-end analysis, which is a technique used in AI. The programme chooses actions from a list of ways, which for a simple robot might include pick up, put down, move forward, move back, move left, and move right, until the current state is changed into the goal state.

Perception: Perception involves scanning the surroundings with numerous sense organs, genuine or artificial, and internal processes for analyzing the scene into objects, their features, and relationships. The fact that one and the same item can have various appearances depending on the angle from which it is viewed, whether or not parts of it are projecting shadows, and so on, complicates analysis.

Artificial perception has progressed to the point where a self-controlled car-like device can drive at modest speeds on the open road, and a mobile robot can search a suite of bustling offices for and remove empty drink cans. FREDDY, a stationary robot with a moving TV 'eye' and a pincer 'hand,' was one of the first systems to merge perception and action (constructed at Edinburgh University during the period 1966-1973 under the direction of Donald Michie). FREDDY could recognise a wide range of items and could be taught to create simple artefacts from a jumble of parts, such as a toy automobile.

Language-understanding: A language is a set of signs with predetermined meaning. For example, traffic signs establish a mini-language; it is a matter of convention that the hazard-ahead sign signifies trouble ahead. This language-specific meaning-by-convention is distinct from what is known as natural meaning, as evidenced by phrases like "Those clouds signify rain" and "The drop in pressure suggests the valve is malfunctioning."

The productivity of full-fledged human languages, such as English, separates them from other types of communication, such as bird sounds and traffic sign systems. A productive language is one that is rich enough to allow for the creation of an infinite number of different sentences.

1.6 APPROACHES TO ARTIFICIAL INTELLIGENCE

In the previous sections of this unit, we learned about various concepts of Artificial Intelligence but now the question is “how do we measure if Artificial Intelligence is making a machine to behave or act or perform like human being or not?”

Perhaps, in the future, we will reach a point where AI can behave like humans, but what guarantees do we have that this will continue? Is it possible to make a system that acts like a human to test the certainty of Artificial Intelligence? " The following approaches constitute the foundation for evaluating an AI entity's human-likeness:

- Turing Test
- Approach of The Cognitive Modelling

- Approach of The Law of Thought
- Approach of The Rational Agent

Let's take a look at how these approaches perform:

In the past, researchers have worked hard to reach all four of these goals. But it is hard to find a good balance between approaches that focus on people and approaches that focus on logic. People are often "irrational" in the sense of being "emotionally unstable," so it's important to tell the difference between human and rational behaviour.

Researchers have found through their studies that a human-centered approach must be an empirical science with hypotheses and experiments to prove them. In a rationalist approach, math and engineering are used together. People in each group sometimes say bad things about the work done by the other groups, but the truth is that each way has led to important discoveries. Let's take a closer look at each one. Acting humanly: The approach of the Turing Test: Alan Turing, who is the most well-known name among the pioneers, thought about how to test an A.I. product to see if it was intelligent. Alan Turing came up with the idea for the Turing Test in 1950 (Turing).

Let's try to answer the question, **"What is the Turing Test in AI?"** Alan Turing, who is the most well-known name among the pioneers, thought about how to test an A.I. product to see if it was intelligent. Turing came up with a test, which is now known as the Turing Test, to see if something is smart. Below is a summary of the Turing test. See figure 2 for more details.




Figure 2: Turing Test

For the purpose of the test, There are three rooms that will be used for the test. In one of the rooms, there is a computer system that is said to be smart. In each of the other two rooms, there is one person sitting. One of the people, who we'll call C, is supposed to ask questions of the computer and the other person, who we'll call B, without knowing who each question is for and, of course, with the goal of figuring out who the computer is. On the other hand, the computer would reply in a way that would keep C from finding out who it is.

The only way for the three of them to talk to each other is through computer terminals. This means that the identity of the computer or person B can only be determined by how intelligent or not the responses are, not by any other human or machine traits. If C can't figure out who the computer is, then the computer must be smart. More accurately, the computer is smart if it can hide its identity from C.

Note that for a computer to be considered smart, it should be smart enough not to answer too quickly, at least not in less than a hundredth of a second, even if it can do something like find the sum of two numbers with more than 20 digits each.

Criticism to Turing Test: There have been a number of criticisms of the Turing test as a machine intelligence test. The Chinese Room Test, developed by John Searle, is one of the most well-known Criticism. The crux of the Chinese Room Test, which we'll discuss below, is that convincing a system, say A, that it possesses qualities of another system, say B, does not suggest that system A actually possesses those qualities. For example, a male human's ability to persuade people that he is a woman does not imply that he is capable of bearing children like a woman.

The scenario for the Chinese Room Test takes place in a single room with two windows. A Shakespeare scholar who knows English but not Chinese is sitting in the room with a kind of Shakespeare encyclopaedia. The encyclopaedia is printed so that for every pair of pages next to each other, one page is written in Chinese characters and the other page is an English translation of the Chinese page. Through one of the windows, Chinese characters with questions about Shakespeare's writing are sent to the person inside. The person looks through the encyclopaedia and, when he or she finds the exact copy of the sequence of characters sent in, reads the English translation, thinks of the answer, and writes it down in English for his or her own understanding. The person then looks in the encyclopaedia for the corresponding sequence of Chinese characters and sends the sequence of Chinese characters through the other window. Now, Searle says that even though the scholar acts as though he or she knows Chinese, this is not the case. Just because a system can mimic a quality doesn't mean that it has that quality.

Thinking humanly: It is the cognitive modeling approach to thinking like a human, from this point of view, the Artificial Intelligence model is based on Human Cognition, which is the core of the human mind. This is done through three approaches, which are as follows:

- Introspection, which means to look at our own thoughts and use those thoughts to build a model.
- Psychological Experiments, which means running tests on people and looking at how they act.
- Brain imaging, which means to use an MRI to study how the brain works in different situations and then copy that through code.

Thinking rationally i.e. The laws of thought approach: This approach Relates to use the laws of thought to think logically: The Laws of Thought are a long list of logical statements that tell our minds how to work. This method, called "Thinking Rationally," is based on these laws. By putting in place algorithms for artificial intelligence, these laws can be written down and made to work. But solving a problem by following the law is very different from solving a problem in the real world. Here are the biggest problems with this approach.

Acting rationally i.e. The rational agent approach: In every situation, a rational agent approach tries to find the best possible outcome. This means that it tries to make the best decision it can given the circumstances. It means that the agent approach is much more flexible and open to change. The Laws of Thought approach, on the other hand, says that a thing must act in a way that makes sense. But there are some situations where there is no logically right thing to do and there are more than one way to solve the problem, each with different results and trade-offs. At that point, the rational agent method works well.

☞ Check Your Progress 2

Q4 Briefly discuss the various components of intelligence

.....

.....

.....

Q5 How do we measure if Artificial Intelligence is making a machine to behave or act or perform like human being or not?"

.....

.....

.....

Q6 What is Turing Test? What is the Criticism to the Turing Test?

.....

.....

.....

1.7 COMPARISON - ARTIFICIAL INTELLIGENCE, MACHINE LEARNING & DEEP LEARNING

Artificial intelligence is a big field that includes a lot of different ways of doing things, from top-down (knowledge representation) to bottom-up (machine learning). In recent years, people have often talked about three related ideas: artificial intelligence (AI), machine learning (ML), and deep learning (DL) (DL). AI is the most general term, machine learning is a part of AI, and deep learning is a type of machine learning. Figure 5 shows how these three ideas are related to each other. Figure 2(b) shows that AI is a broad field with many different subdomains. However, AI's recent rise in popularity is largely due to how well machine learning, especially deep learning, works. So, this entry will talk about these two areas of AI: Machine Learning (ML) and Deep Learning (DL) Figure 2 (a)




Fig 2(a):AI, ML, DL




Figure 2(b) : Various Sub Domains of Artificial Intelligence

To make a system that is artificially intelligent, you have to carefully do Reverse-Engineering of human traits and machine abilities. Also, to understand how an AI system really works, you have to get a good grasp of the different parts of AI and how they can be used in different industries or industrial fields.

Introduction to Machine Learning (ML): Machine learning is a branch of artificial intelligence (AI). It explains one of the most important ideas in AI, which has to do with learning through experience and not through being taught. One of the most recent advances in AI, this way of learning is made possible by applying machine learning to very large data sets. Machine learning algorithms find patterns and learn how to make predictions and recommendations by using data and experiences instead of explicit programming instructions. The algorithms also change as they get new data and learn from their experiences. This makes them more effective over time.

Algorithms for machine learning are based on ways that people learn from their experiences. This means that they are programmed to learn from what they do and get better at what they do. They don't need to be told what to do to get the desired results. They are set up so that people can learn by looking at the data sets they can access and comparing what they see to examples of the final results. They also look for patterns in the output and try to figure out how to use the different parts to get the output they want.

ML shows a machine how to draw conclusions and make decisions based on what it has learned in the past. It looks for patterns and analyses past data to figure out what these patterns mean so that a possible conclusion can be reached without the need for human experience. Businesses save time and make better decisions by using automation to evaluate data and come to conclusions.

Machine learning provides predictions and prescriptions
Types of analytics (in order of increasing complexity)




Figure 3 :Machine Learning : Descriptive, Predictive and Prescriptive Analytics

Introduction to Deep Learning (DL) Deep Learning is a subfield of machine learning that focuses on algorithms called "Artificial Neural Networks" (ANN) that are based on how the brain is built and how it works. Deep Learning is a type of machine learning that can handle a wider range of data sources, needs less pre-processing of data, and often gives more accurate results than traditional machine learning methods.

A neural network is made up of layers of software-based calculators called "neurons" that are linked together. This neural network can take in a huge amount of data and process it through many layers. At each layer, the network learns more complex features of the data. The network can then decide what to do with the data, find out if it was right, and use what it has learned to decide what to do with new data. Deep learning is a way of programming computers that uses the way neural networks work to teach computers to do things that humans do naturally. So, Deep Learning is a way to teach a computer model to run classification algorithms based on an image, text, or sound. Once a neural network knows what an object looks like, it can spot that object in a new picture.

Deep learning is becoming more popular because its models can get better results. It uses large sets of labeled data and neural network architectures to train the models..

☞ Check Your Progress 3

Q7 Compare Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL).

.....

.....

.....

1.8 APPLICATION AREAS OF ARTIFICIAL INTELLIGENCE SYSTEMS

Artificial intelligence is the most important factor in the transformation of economies straight from the ground up, and it is contributing as an efficient alternative. It has a lot of potential to perform optimization in any industry, whether it smart cities or the health sector or agriculture or any other prospective sector of relevance, and below we have included a few of the systems in which AI is functioning as the major source of competitive advantage:

a) Healthcare: The application of AI in healthcare can help address issues of high barriers to access to healthcare facilities, particularly in rural areas that suffer from poor connectivity and a limited supply of healthcare professionals. This is especially true in areas where the supply of healthcare professionals is limited. The deployment of use cases like as AI-driven diagnostics, personalised treatment, early diagnosis of potential pandemics, and imaging diagnostics, amongst others, is one way to accomplish this goal.




Figure 4: Potential Use of AI in Health Care

b) Agriculture: AI has the potential to bring in a food revolution while simultaneously satisfying the ever-increasing need for food (global need to produce 50 percent more food and cater to an additional 2 billion people by 2050 as compared to today). It also has the ability to resolve issues such as inadequate demand prediction, a lack of secure irrigation, and the abuse or misuse of pesticides and fertilizers. These are only some of the problems that it could solve. The increase of crop output through real-time advising is one example of a use case. Other use cases include the advanced detection of pest infestations and the forecast of crop prices to advise sowing methods.




Figure 5: AI for Precision Farming -

All of the stages of the agricultural value chain indicated above in figure 5 have the potential for the application of artificial intelligence and other associated technologies to have an impact on the levels of production and efficiency at those stages.

c) Smart Mobility, including Transports and Logistics: Autonomous fleets for ride sharing, semi-autonomous features such as driver assistance, and predictive engine monitoring and maintenance are all possible use cases for smart mobility, which includes transportation and logistics. Other areas where AI can have a positive impact include self-driving trucks and delivery, as well as better traffic control.

d) Retail: The retail industry was one of the first to use AI solutions. For example, personalised suggestions, browsing based on user preferences, and image-based product search have all been used to improve the user experience. Other use cases include predicting what customers will want, keeping track of inventory better, and managing deliveries more efficiently.

e) Manufacturing: AI-based solutions are expected to help the manufacturing industry the most. This will make possible the "Factory of the Future" by allowing flexible and adaptable technical systems to automate processes and machinery that can respond to new or unexpected situations by making smart decisions. Impact areas include engineering (AI for R&D), supply chain management (predicting demand), production (AI can cut costs and increase efficiency), maintenance (predictive maintenance and better use of assets), quality assurance (e.g., vision systems with machine learning algorithms to find flaws and differences in product features), and in-plant logistics and warehousing.

f) Energy: In the energy sector, possible use cases include modelling and forecasting the energy system to make it less unpredictable and make balancing

and using power more efficient. In renewable energy systems, AI can help store energy through smart metres and intelligent grids. It can also make photovoltaic energy more reliable and less expensive. AI could also be used to predict maintenance of grid infrastructure, just like it is in manufacturing.

g) Smart Cities: Integrating AI into newly built smart cities and infrastructure could also help meet the needs of a population that is moving to cities quickly and improve the quality of life for those people. Some possible use cases include controlling traffic to reduce traffic jams and managing crowds better to improve security.

h) Education and Skilling: Quality and access problems in the education sector might be fixed by AI. Possible uses include adding to and improving the learning experience through personalized learning, automating and speeding up administrative tasks, and predicting when a student needs help to keep them from dropping out or to suggest vocational training.

i) Financial industry: The financial industry also uses AI. For example, it helps the fraud department of a bank find and flag suspicious banking and finance activities like unusual debit card use and large account deposits. AI is also used to make trading easier and more efficient. This is done by making it easier to figure out how many securities are being bought and sold and how much they cost.

Top Used Applications of Artificial Intelligence

- Tools and checkers for plagiarism
- Recognizing faces;
- Putting an AI autopilot on commercial planes
- Applications for sharing rides (E.g.: Uber, Lyft)
- E-mail spam filters; voice-to-text features; search suggestions
- Google's predictions based on AI (E.g.: Google Maps)
- Protecting against and stopping fraud.
- Smart personal assistants (E.g.: Siri, Alexa)

There are various ways to use artificial intelligence. The technology can be used in different industries and sectors, but the adoption of AI by different sectors has been affected by technical and regulatory challenges, but the biggest factor has been how it will affect business.

1.9 INTELLIGENT AGENTS

An agent may be thought of as an entity that acts, generally on behalf of someone else. More precisely, an agent is an entity that perceives its environment through sensors and acts on the environment through actuators. Some experts in the field require an agent to be additionally autonomous and goal directed also.

A percept may be thought of as an input to the agent through its sensors, over a unit of time, sufficient enough to make some sense from the input.

Percept sequence is a sequence of percepts, generally long enough to allow the agent to initiate some action.

In order to further have an idea about what a computer agent is, let us consider one of the first definitions of agent, which was coined by John McCarthy and his friends at MIT.

A software agent is a system which, when given a goal to be achieved, could carry out the details of the appropriate (computer) operations and further, in case it gets stuck, it can ask for advice and can receive it from humans, may even evaluate the appropriateness of the advice and then act suitably.

Essentially, a computer agent is a computer software that additionally has the following attributes:

- (i) it has autonomous control i.e., it operates under its own control
- (ii) it is perceptive, i.e., it is capable of perceiving its own environment
- (iii) it persists over a long period of time
- (iv) it is adaptive to changes in the environment and
- (v) it is capable of taking over others' goals.

As the concept of Intelligent Agents is of relatively new, different pioneers and other experts have been conceiving and using the term in different ways. There are two distinct but related approaches for defining an agent. The first approach treats an agent as an ascription i.e., the perception of a person (which includes expectations and points of view) whereas the other approach defines an agent on the basis of the description of the properties that the agent to be designed is expected to possess.

Let us first discuss the definition of agent according to first approach. Among the people who consider an agent as an ascription, a popular slogan is "Agent is that agent does". In everyday context, an agent is expected to act on behalf of someone to carry out a particular task, which has been delegated to it. But to perform its task successfully, the agent must have knowledge about the domain in which it is operating and also about the properties of its current user in question. In the course of normal life, we hire different agents for different jobs based on the required expertise for each job. Similarly, a non-human intelligent agent also is imbedded with required expertise of the domain as per requirements of the job under consideration. For example, a football-playing agent would be different from an email-managing agent, although both will have the common attribute of modeling their user.

According to the second approach, an agent is defined as an entity, which functions continuously and autonomously, in a particular environment, which may have other agents also. By continuity and autonomy of an agent, it is meant that the agent must be able to carry out its job in a flexible and intelligent fashion and further is expected to adapt to the changes in its environment without requiring constant human guidance or intervention. Ideally, an agent that functions continuously in an environment over a long period of time would also learn from its experience. In addition, we expect an agent, which lives in a

multi-agent environment, to be able to communicate and cooperate with them, and perhaps move from place to place in doing so.

According to the second approach to defining agent, an agent is supposed to possess some or all of the following properties:

- **Reactivity:** The ability of sensing the environment and then acting accordingly.
- **Autonomy:** The ability of moving towards its goal, changing its moves or strategy, if required, without much human intervention.
- **Communicating ability:** The ability to communicate with other agents and humans.
- **Ability to coexist by cooperating:** The ability to work in a multi-agent environment to achieve a common goal.
- **Ability to adapt to a new situation:** Ability to learn, change and adapt to the situations in the world around it.
- **Ability to draw inferences:** The ability to infer or conclude facts, which may be useful, but are not available directly.
- **Temporal continuity:** The ability to work over long periods of time.
- **Personality:** Ability to impersonate or simulate someone, on whose behalf the agent is acting.
- **Mobility:** Ability to move from one environment to another.

Task environments or problem environments are the environments, which include all the elements involved in the problems for which agents are thought of as solutions. Task environments will vary with every new task or problem for which an agent is being designed. Specifying the task environment is a long process which involves looking at different measures or parameters. Next, we discuss a standard set of measures or parameters for specifying a task environment under the heading PEAS.

PEAS (Performance, Environment, Actuators, Sensors)

For designing an agent, the first requirement is to specify the task environment to the maximum extent possible. The task environment for an agent to solve one type of problems, may be described by the four major parameters namely, performance (which is actually the expected performance), environment (i.e., the world around the agent), actuators (which include entities through which the agent may perform actions) and sensors (which describes the different entities through which the agent will gather information about the environment).

The four parameters may be collectively called as PEAS. We explain these parameters further, through an example of an automated agent, which we will preferably call automated public road transport driver. This is a much more complex agent than the simple boundary following robot which we have already discussed.

Example (An Automated Public Road Transport Driver Agent)

We describe the task environment of the agent on the basis of PEAS.

Performance Measures: Some of the performance measures which can easily be perceived of an automated public road transport driver would be:

- Maximizing safety of passengers
- Maximizing comfort of passengers
- Ability to reach correct destination
- Ability to minimize the time to reach the destination Obeying traffic rules
- Causing minimum discomfort or disturbance to other agents
- Minimizing costs, etc.

Environment (or the world around the agent) We must remember that the environment or the world around the agent is extremely uncertain or open ended. There are unlimited combinations of possibilities of the environment situations, which such an agent could face. Let us enumerate some of the possibilities or circumstances which an agent might face:

- Variety of roads e.g., from 12-lane express-ways, freeways to dusty rural bumpy roads; different road rules including the ones requiring left-hand drive in some parts of the world and right-hand drive in other parts.
- The degree of knowledge of various places through which and to which driving is to be done.
- Various kinds of passengers, including high cultured to almost ruffians etc.
- All kind of other traffic possibly including heavy vehicles, ultra-modern cars, three-wheelers and even bullock carts.

Actuators: These include the following:

- Handling steering wheel, brakes, gears and accelerator
- Understanding the display screen
- A device or devices for all communication required

Sensors: The agent acting as automated public road transport driver must have some way of sensing the world around it i.e., the traffic around it, the distance between the automobile and the automobiles ahead of it and its speed, the speeds of neighboring vehicles, the condition of the road, any turn ahead etc. It may use sensors like odometer, speedometer, sensors telling the different parameters of the engine, Global Positioning System (GPS) to understand its current location and the path ahead. Also, there should be some sort of sensors to calculate its distance from other vehicles etc.

We must remember that the agent example the automated public road transport driver, which we have considered above, is quite difficult to implement. However, there are many other agents, which operate in comparatively simpler and less dynamic environments, e.g., a game playing robot, an assembly line robot control, and an image processing agent etc.

In respect of the design and development of intelligent agents, with the passage of time, the momentum seems to have shifted from hardware to software, the latter being thought of as a major source of intelligence. But, obviously, some sort of hardware is essentially needed as a home to the intelligent agent.

There are two parts of an agent or its structure:

- A (hardware) device with sensors and actuators in which that agent will reside, called the architecture of the agent.
- An agent program that will convert or map the percepts in to actions.

Also, the agent program and its architecture are related in the sense that for a different agent architecture a different type of agent program is required and vice-versa. For example, in case of a boundary following robot, if the robot does not have the capability of sensing adjacent cells to the right, then the agent program for the robot has to be changed.

Next, we discuss different categories of agents, which are differentiated from each other on the basis of their agent programs. Capability to write efficient agent programs is the key to the success for developing efficient rational agents. Although the table driven approach (in which an agent acts on the basis of the set of all possible percepts by storing these percepts in tables) to design agents is possible yet the approach of developing equivalent agent programs is found much more efficient.

Next, we discuss some of the general categories of agents based on their agents' programs, Agents can be grouped into five classes based on their degree of perceived intelligence and capability. All these agents can improve their performance and generate better action over the time. These are given below:

- *SR (Simple Reflex) agents*
- *Model Based reflex agents*
- *Goal-based agents*
- *Utility based agents*
- *Stimulus-Response Agents*
- *Learning agents*

SR (Simple Reflex) agents: These are the agents or machines that have no internal state (i.e., they don't remember anything) and simply react to the current percepts in their environments. An interesting set of agents can be built, the behaviour of the agents in which can be captured in the form of a simple set of functions of their sensory inputs. One of the earliest implemented agents of this category was called *Machina Speculatrix*. This was a device with wheels, motor, photo cells and vacuum tubes and was designed to move in the direction of light of less intensity and was designed to avoid the direction of the bright light. **A boundary following robot is also an SR agent.** For an automobile-driving agent also, some aspects of its behavior like applying brakes immediately on observing either the vehicle immediately ahead applying brakes or a human being coming just in front of the automobile suddenly, show the simple reflex

capability of the agent. Such a simple reflex action in the agent program of the agent can be implemented with the help of simple condition-action rules.

For example : **IF** *a human being comes in front of the automobile suddenly*
THEN *apply breaks immediately.*

Although implementation of SR agents is simple **yet on the negative side** this type of agents has very limited intelligence because ***they do not store or remember anything***. As a consequence, they cannot make use of any previous experience. In summary, they do not learn. Also **they are capable of operating correctly only if the environment is fully observable.**

Model Based Reflex agents: Simple Reflex agents are not capable of handling task environments that are not fully observable. In order to handle such environments properly, in addition to reflex capabilities, the agent should, maintain some sort of internal state in the form of a function of the sequence of percepts recovered up to the time of action by the agent. Using the percept sequence, the internal state is determined in such a manner that it reflects some of the aspects of the unobservable environment. Further, in order to reflect properly the unobserved environment, the agent is expected to have a model of the task environment encoded in the agent's program, where the model has the knowledge about–

- (i) the process by which the task environment evolves independent of the agent and
- (ii) effects of the actions of the agent have on the environment.

Thus, in order to handle properly the partial observability of the environment, the agent should have a model of the task environment in addition to reflex capabilities. Such agents are called **Model-based Reflex Agents**.

Goal Based Agents: In order to design appropriate agent for a particular type of task, we know the nature of the task environment plays an important role. Also, it is desirable that the complexity of the agent should be minimum and just sufficient to handle the task in a particular environment. In this regard, first we discussed the simplest type of agents, viz., Simple Reflex Agents. The action of this type of agent is decided by the current precept only. Next, we discussed the Model-Based Reflex Agents, for which an action is decided by taking into consideration not only the latest precept, but the whole precept history summarized in the form of internal state. Also, action for this type of agent is also decided by taking into consideration the knowledge of the task environment, represented by a model of the environment and encoded into the agent's program. However, in respect of a number of tasks, even this much knowledge may not be sufficient for appropriate action. For example, when we are going from city A to city B, in order to take appropriate action, it is not enough to know the summary of actions and path which has taken us to some city C between A and B. We also have to remember the goal of reaching to city B.

Goal based agents are **driven by the goal** they want to achieve, i.e., **their actions are based on the information regarding their goal, in addition to, of course, other information in the current state.** This goal information is

also a part of the current state description and it describes everything that is desirable to achieve the goal. As mentioned earlier, an example of a goal-based agent is an agent that is required to find the path to reach a city. In such a case, if the agent is an automobile driver agent, and if the road is splitting ahead into two roads, then the agent has to decide which way to go to achieve its goal of reaching its destination. Further, if there is a crossing ahead then the agent has to decide, whether to go straight, to go to the left or to go to the right. In order to achieve its goal, the agent needs some information regarding the goal which describes the desirable events and situations to reach the goal. The agent program would then use this goal information to decide the set of actions to take in order to reach its goal.

Another desirable capability which a good goal-based agent should have been that if an agent finds that a part of the sequence of the previous steps has taken the agent away from its goal then it should be able to retract and start its actions from a point which may take the agent toward the goal.

In order to take appropriate action, decision-making process in goal-based agents may be simple or quite complex depending on the problem. Also, **the decision-making required by the agents of this kind needs some sort of looking into the future**. For example, it may analyze the possible outcome of a particular action before it actually performs that action. In other words, we can say that *the agent would perform some sort of reasoning of if-then-else type*, e.g., an automobile driver agent having one of its goals as not to hit any vehicle in front of it, when finds the vehicle immediately ahead of it slowing down may not apply brakes with full force and instead may apply brakes slowly so that the vehicles following it may not hit it.

As the goal-based agents may have to reason before they take an action, these agents might be slower than other types of agents but will be more flexible in taking actions as their decisions are based on the acquired knowledge which can be modified also. Hence, **as compared to SR agents** which may require rewriting of all the condition-action rules in case of change in the environment, the goal-based agents can adapt easily when there is any change in its goal.

Utility Based Agents: Goal based agent's success or failure is judged in terms of its capability for achieving or not achieving its goal. A goal-based agent, for a given pair of environment state and possible input, only knows whether the pair will lead to the goal state or not. Such an agent will not be able to decide in which direction to proceed when there are more than one conflicting goals. Also, in a goal-based agent, there is no concept of partial success or somewhat satisfactory success. Further, if there are more than one method of achieving a goal, then no mechanism is incorporated in a Goal-based agent of choosing or finding the method which is faster and more efficient one, out of the available ones, to reach its goal.

A more general way to judge the success or happiness of an agent may be, through assigning to each state a number as an approximate measure of its success in reaching the goal from the state. In case, the agent is embedded with such a capability of assigning such numbers to states, then it can choose, out of the reachable states in the next move, the state with the highest assigned

number, out of the numbers assigned to various reachable states, indicating possibly the best chance of reaching the goal.

It will allow the goal to be achieved more efficiently. Such an agent will be more useful, i.e., will have more utility. A utility-based agent uses a utility function, which maps each of the world states of the agent to some degree of success. If it is possible to define the **utility function** accurately, then the agent will be able to reach the goal quite efficiently. Also, a utility-based agent is able to make decisions in case of conflicting goals, generally choosing the goal with higher success rating or value. Further, in environments with multiple goals, the utility-based agent quite likely chooses the goal with least cost or higher utility goal out of multiple goals.

Stimulus-Response Agents: A stimulus response agent (or a reactive agent) take input from the world through sensors, and then take action based on those inputs through actuators. Between the stimulus and response, there is a processing unit that can be arbitrarily complex. An example of such an agent is one that controls a vehicle in a racing game: the agent “looks” at the road and nearby vehicles, and then decides how much to turn and break. Such Agents (Stimulus-Response Agents are the Reactive agents) represents a special category of agents, which do not possess internal, symbolic models of their environments; instead, they act/respond in a stimulus-response manner to the present state of the environment in which they are embedded. These agents are relatively simple and they interact with other agents in basic ways. Nevertheless, complex patterns of behavior emerged from the interactions when the ensemble of agents is viewed globally

Learning Agents: It is not possible to encode all the knowledge in advance, required by a rational agent for optimal performance during its lifetime. This is especially true of the real life, and not just theoretical, environments. These environments are **dynamic** in the sense that the environmental conditions change, not only due to the actions of the agents under considerations, but due to other environmental factors also. For example, all of a sudden, a pedestrian comes just in front of the moving vehicle, even when there is green signal for the vehicle. In a multi-agent environment, all the possible decisions and actions an agent is required to take, are generally unpredictable in view of the decisions taken and actions performed simultaneously by other agents. Hence, **the ability of an agent to succeed in an uncertain and unknown environment depends on its learning capability** i.e., its capability to change approximately its knowledge of the environment. For an agent with learning capability, some initial knowledge is coded in the agent program and after the agent starts operating, it learns from its actions the evolving environment, the actions of its competitors or adversaries etc. so as to improve its performance in ever-changing environment. If approximate learning component is incorporated in the agent, then the knowledge of the agent gradually increases after each action starting from its initial knowledge which was manually coded into it at the start.

Conceptually the learning agent consists of four components:

- (i) **Learning Component:** It is the component of the agent, which on the basis of the percepts and the feedback from the environment, gradually improves the performance of the agent.

- (ii) **Performance Component:** It is the component from which all actions originate on the basis of external percepts and the knowledge provided by the learning component.

The design of learning component and the design of performance element are very much related to each other because a learning component is of no use unless the performance component can be designed to convert the newly acquired knowledge into better useful actions.

- (iii) **Critic Component:** This component finds out how well the agent is doing with respect to a certain fixed performance standard and it is also responsible for any future modifications in the performance component. The critic is necessary to judge the agent's success with respect to the chosen performance standard, especially in a dynamic environment. Foreexample, in order to check whether a certain job is accomplished, the critic will not depend on external percepts only but it will also compare the current state to the state, which indicates the completion of that task.

- (iv) **Problem Generator Component:** This component is responsible for suggesting actions (some of which may not be optimal) in order to gain some fresh and innovative experiences. Thus, this component allows the agent to experiment a little by traversing sometimes uncharted territories by choosing some new and suboptimal actions. This may be useful, because the actions which may seem suboptimal in a short run, may turn out to be much better in the long run.

In the case of an automobile driver agent, this agent would be of little use if it does not have learning capability, as the environment in which it has to operate is totally dynamic and unpredictable in nature. **Once the automobile driver agent starts operating, it keeps on learning from its experiences, both positive and negative.** If faced with a totally new and previously unknown situation, e.g., encountering a vehicle coming from the opposite direction on a one-way road, the problem generator component of the driver agent might suggest some innovative action to tackle this new situation. Moreover, the learning becomes more difficult in the case of an automobile driver agent, because the environment is only partially observable.

Different Forms of Learning in Agents: The purpose of embedding learning capability in an agent is that it should not depend totally on the knowledge initially encoded in it and on the external percepts for its actions. The agent learns by evaluating its own decisions and/or making observations of new situations it encounters in the ever-changing environment.

There may be various criteria for developing learning taxonomies. The criteria may be based on –

- The type of knowledge learnt, e.g., concepts, problem-solving or game playing,
- The type of representation used, e.g., predicate calculus, rules or frames,
- The area of application, e.g., medical diagnosis, scheduling or prediction.

Check Your Progress 4

Q9 What are Intelligent agents in AI? Briefly discuss the properties of Agents.

.....

.....

.....

Q10 What are Task environments? Briefly discuss the standard set of measures or parameters for specifying a task environment under the heading PEAS.

.....

.....

.....

1.10 SUMMARY

In this unit we learned about the difference between knowledge and intelligence and also pointed out the meaning of Artificial Intelligence (AI), along with the application of AI systems in various fields. The unit also covers the historical development of the field of AI systems. Along with the development of the AI as a discipline, the need of classification of AI systems was felt, and hence the unit discussed the classification of the AI systems in detail. Further, the unit discussed about the concepts of Artificial Intelligence (AI), Machine Learning (ML) and Deep Learning (DL). Finally, the unit discussed the components of Intelligence, which was extended for the understanding of the concepts of Intelligent Agents, with special emphasis on Stimulus - Response Agents

1.11 SOLUTIONS/ANSWERS

Check Your Progress 1

Q1 How Knowledge differs from intelligence? What do you understand by the term Artificial Intelligence (AI) ? List the various technologies and their corresponding illustrative solutions.

Sol- Refer section 1.3

Q2 Classify AI on the basis of the functionalities of AI

Sol- Refer section 1.3

Q3 Compare ANI, AGI and ASI, in context of AI

Sol- Refer section 1.3

Check Your Progress 2

Q4 Briefly discuss the various components of intelligence

Sol – Refer Section 1.5

Q5 How do we measure if Artificial Intelligence is making a machine to behave or act or perform like human being or not ?”

Sol – Refer Section 1.6

Q6 What is Turing Test? What is the Criticism to the Turing Test ?

Sol – Refer Section 1.6

☞ **Check Your Progress 3**

Q7 Compare Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL).

Sol – Refer Section 1.7

Q8 Compare Descriptive, Predictive and Prescriptive analytics performed under Machine Learning.

Sol – Refer Section 1.7

☞ **Check Your Progress 4**

Q9 What are Intelligent agents in AI? Briefly discuss the properties of Agents.

Sol – Refer Section 1.9

Q10 What are Task environments? Briefly discuss the standard set of measures or parameters for specifying a task environment under the heading PEAS.

Sol – Refer Section 1.9

1.12 FURTHER READINGS

1. Ela Kumar, “Artificial Intelligence”, IK International Publications
2. E. Rich and K. Knight, “Artificial intelligence”, Tata Mc Graw Hill Publications
3. N.J. Nilsson, “Principles of AI”, Narosa Publ. House Publications
4. John J. Craig, “Introduction to Robotics”, Addison Wesley publication
5. D.W. Patterson, “Introduction to AI and Expert Systems" Pearson publication