

Unit-1: Introduction to AI, Part-1

Concept of AI:

Artificial intelligence, or AI, refers to the simulation of human intelligence by software-coded heuristics. Nowadays this code is prevalent in everything from cloud-based, enterprise applications to consumer apps and even embedded firmware.

The year 2022 brought AI into the mainstream through widespread familiarity with applications of Generative Pre-Training Transformer. The most popular applications are OpenAI's DALL-E text-to-image tool and ChatGPT. The widespread fascination with ChatGPT made it synonymous with AI in the minds of most consumers. However, it represents only a small portion of the ways that AI technology is being used today.

The ideal characteristic of artificial intelligence is its ability to rationalize and take actions that have the best chance of achieving a specific goal. A subset of artificial intelligence is machine learning (ML), which refers to the concept that computer programs can automatically learn from and adapt to new data without being assisted by humans. Deep learning techniques enable this automatic learning through the absorption of huge amounts of unstructured data such as text, images, or video.

When most people hear the term artificial intelligence, the first thing they usually think of is robots. That's because big-budget films and novels weave stories about human-like machines that wreak havoc on Earth. But nothing could be further from the truth.

Artificial intelligence is based on the principle that human intelligence can be defined in a way that a machine can easily mimic it and execute tasks, from the most simple to those that are even more complex. The goals of artificial intelligence include mimicking human cognitive activity. Researchers and developers in the field are making surprisingly rapid strides in mimicking activities such as learning, reasoning, and perception, to the extent that these can be concretely defined. Some believe that innovators may soon be able to develop systems that exceed the capacity of humans to learn or reason out any subject. But others remain skeptical because all cognitive activity is laced with value judgments that are subject to human experience.

As technology advances, previous benchmarks that defined artificial intelligence become outdated. For example, machines that calculate basic functions or recognize text through optical character recognition are no longer considered to embody artificial intelligence, since this function is now taken for granted as an inherent computer function.

AI is continuously evolving to benefit many different industries. Machines are wired using a cross-disciplinary approach based on mathematics, computer science, linguistics, psychology, and more.

History of AI:

The first work that is now generally recognized as AI was done by Warren McCulloch and Walter Pills (1943). They drew on three sources: knowledge of the basic physiology and function of neurons in the brain; a formal analysis of propositional logic due to Russell and

Whitehead; and Turing's theory of computation. They proposed a model of artificial neurons in which each neuron is characterized as being "on" or "off" with a switch to "on" occurring in response to stimulation by a sufficient number of neighboring neurons. The state of a neuron was conceived of as "factually equivalent to a proposition which proposed its adequate stimulus." They showed, for example, that any computable function could be computed by some network of connected neurons, and that all the logical connectives (and, or not, etc.) could be implemented by simple net structures. McCulloch and Pints also suggested that suitably defined networks could learn. Donald Hebb (1949) demonstrated a simple updating rule for modifying the connection strengths between neurons. His rule, now called Hebbian learning, remains an influential model to this day.

Two undergraduate students at Harvard, Marvin Minsky and Dean Edmonds, built the first neural network computer in 1950. The SNARC, as it was called, used 3000 vacuum tubes and a surplus automatic pilot mechanism from a B-24 bomber to simulate a network of 40 neurons. Later, at Princeton, Minsky studied universal computation in neural networks. His Ph.D. committee was skeptical about whether this kind of work should be considered mathematics, but von Neumann reportedly said, "If it isn't now, it will be someday." Minsky was later to prove influential theorems showing the limitations of neural network research.

There were a number of early examples of work that can be characterized as AI, but Alan Turing's vision was perhaps the most influential. He gave lectures on the topic as early as 1947 at the London Mathematical Society and articulated a persuasive agenda in his 1950 article "Computing Machinery and Intelligence." Therein, he introduced the Turing Test, machine learning, genetic algorithms, and reinforcement learning. He proposed the Child Programme idea, explaining "Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulated the child's?"

Current Status:

AI is a vast and ever-evolving domain, and one way to categorize its progress is by classifying it into three main categories: Artificial Narrow Intelligence, General Intelligence, and Super Intelligence. The classification will help us better understand the current state of AI because:

1. All current AI solutions belong to the Artificial Narrow Intelligence class.
2. ANI systems are good at specific tasks, like analyzing data, generating texts, recognizing images, and translating languages. ANI is already being used to automate tasks and provide insights in many industries.
3. General Intelligence is achieved when AI is able to independently reason and solve problems for which it was not even originally designed or intended, in other words, at that point, AI begins to match human intelligence. Though AI systems are capable of performing many tasks that were once thought to be the exclusive domain of humans, it's still far from achieving General Intelligence.
4. Superintelligence would be AI that is smarter than humans in every way. It would be able to reason, be creative, and adapt to new situations. Superintelligence is still far away, but it is a goal that some AI researchers are working towards.

AI technologies have made significant advancements in various fields, including natural language processing, image recognition, and automation. These developments have led to the creation of intelligent systems that can perform tasks that were once thought to be exclusive to human capabilities.

The current status of Artificial Intelligence is marked by a substantial increase in adoption among large companies. According to the latest Artificial Intelligence Index report, AI adoption among large companies has surged by an impressive 47% when compared to data from 2018.

This surge in adoption underscores the growing recognition of AI's potential to transform various industries and enhance business operations. Large companies are increasingly leveraging AI technologies to gain competitive advantages, streamline processes, and make data-driven decisions.

Scope:

1. AI in Science and Research

AI is making lots of progress in the scientific sector. Artificial Intelligence can handle large quantities of data and processes it quicker than human minds. This makes it perfect for research where the sources contain high data volumes.

2. AI in Cyber Security

Cybersecurity is another field that's benefitting from AI. As organizations are transferring their data to IT networks and cloud, the threat of hackers is becoming more significant.

3. AI in Data Analysis

Data analysis can benefit largely from AI and ML. AI algorithms are capable of improving with iterations, and this way, their accuracy, and precision increase accordingly. AI can help data analysts with handling and processing large datasets.

4. AI in Transport

The transport sector has been using AI for decades. Although the presence and scope of artificial intelligence have been theoretically existent for some time now, only a few people are aware that we use it on a regular basis. The scope of AI is quite advanced because it can learn automated manual tasks just like humans. Since automation is becoming more prevalent, time-consuming tasks are being taken over by AI. Any system that has AI shows the same amount of potential that human intelligence shows. These AI-powered machines can carry out jobs that require critical thinking and decision-making processes all by themselves. Several businesses in the transportation sector have been taken over by AI so that they can stay ahead in the market.

5. AI in Home

AI has found a special place in people's homes in the form of Smart Home Assistants. Amazon Echo and Google Home are popular smart home devices that let you perform various tasks with just voice commands.

6. AI in Healthcare

The medical sector is also using this technology for its advantages. AI is helping medical researchers and professionals in numerous ways. Wrong diagnoses are a significant problem in the medical sector. AI can help doctors in avoiding these errors by providing them with relevant databases and recommendations. It can analyze the database of patients with similar symptoms and suggest the treatment that was the most successful in those cases.


7. AI in Education

The importance of education in this world has been prevalent, but it continues to grow even today. With a large part of the country's population being the youth, it is important that they receive a good quality education. Along with that, it is also necessary that they understand AI and its benefits. Just like all the other sectors, it is critical for the education sector to keep up with AI as well as the artificial intelligence scope keeps increasing to fuel the education sector.

Agents:

An agent is anything that can perceive its environment through sensors and acts upon that environment through effectors.

1. A human agent has sensory organs such as eyes, ears, nose, tongue and skin parallel to the sensors, and other organs such as hands, legs, mouth, for effectors.
2. A robotic agent replaces cameras and infrared range finders for the sensors, and various motors and actuators for effectors.
3. A software agent has encoded bit strings as its programs and actions.



Agent Terminology:

1. Performance Measure of Agent – It is the criteria, which determines how successful an agent is.
2. Behavior of Agent – It is the action that agent performs after any given sequence of percepts.
3. Percept – It is agent's perceptual inputs at a given instance.
4. Percept Sequence – It is the history of all that an agent has perceived till date.
5. Agent Function – It is a map from the precept sequence to an action.

Rationality:

Rationality is nothing but status of being reasonable, sensible, and having good sense of judgment. Rationality is concerned with expected actions and results depending upon what the agent has perceived. Performing actions with the aim of obtaining useful information is an important part of rationality.

A rational agent always performs right action, where the right action means the action that causes the agent to be most successful in the given percept sequence. The problem the agent solves is characterized by Performance Measure, Environment, Actuators, and Sensors (PEAS).

The Structure of Intelligent Agents:

Agent's structure can be viewed as –

1. Agent = Architecture + Agent Program
2. Architecture = the machinery that an agent executes on.
3. Agent Program = an implementation of an agent function.

Simple Reflex Agents:

They choose actions only based on the current percept. They are rational only if a correct decision is made only on the basis of current precept. Their environment is completely observable.

Model Based Reflex Agents:

They use a model of the world to choose their actions. They maintain an internal state.

1. Model – knowledge about “how the things happen in the world”.
2. Internal State – It is a representation of unobserved aspects of current state depending on percept history.

Goal Based Agents:

They choose their actions in order to achieve goals. Goal-based approach is more flexible than reflex agent since the knowledge supporting a decision is explicitly modeled, thereby allowing for modifications.

Utility Based Agents:

They choose actions based on a preference (utility) for each state.

Goals are inadequate when –

1. There are conflicting goals, out of which only few can be achieved.
2. Goals have some uncertainty of being achieved and you need to weigh likelihood of success against the importance of a goal.

Environments:

An environment in artificial intelligence is the surrounding of the agent. The agent takes input from the environment through sensors and delivers the output to the environment through actuators. There are several types of environments:

1. Fully Observable vs Partially Observable:

When an agent sensor is capable to sense or access the complete state of an agent at each point in time, it is said to be a fully observable environment else it is partially observable.

Maintaining a fully observable environment is easy as there is no need to keep track of the history of the surrounding. An environment is called unobservable when the agent has no sensors in all environments.

Examples:

Chess – the board is fully observable, and so are the opponent's moves.

Driving – the environment is partially observable because what's around the corner is not known.

2. Deterministic vs Stochastic:

When a uniqueness in the agent's current state completely determines the next state of the agent, the environment is said to be deterministic. The stochastic environment is random in nature which is not unique and cannot be completely determined by the agent.

Examples:

Chess – there would be only a few possible moves for a coin at the current state and these moves can be determined.

Self-Driving Cars- the actions of a self-driving car are not unique, it varies from time to time.

3. Competitive vs Collaborative:

An agent is said to be in a competitive environment when it competes against another agent to optimize the output. The game of chess is competitive as the agents compete with each other to win the game which is the output. An agent is said to be in a collaborative environment when multiple agents cooperate to produce the desired output. When multiple self-driving cars are found on the roads, they cooperate with each other to avoid collisions and reach their destination which is the output desired.

4. Single-agent vs Multi-agent:

An environment consisting of only one agent is said to be a single-agent environment. A person left alone in a maze is an example of the single-agent system. An environment involving more than one agent is a multi-agent environment. The game of football is multi-agent as it involves 11 players in each team.

5. Dynamic vs Static:

An environment that keeps constantly changing itself when the agent is up with some action is said to be dynamic. A roller coaster ride is dynamic as it is set in motion and the environment keeps changing every instant. An idle environment with no change in its state is called a static environment. An empty house is static as there's no change in the surroundings when an agent enters.

6. Discrete vs Continuous:

If an environment consists of a finite number of actions that can be deliberated in the environment to obtain the output, it is said to be a discrete environment. The game of chess is discrete as it has only a finite number of moves. The number of moves might vary with every game, but still, it's finite. The environment in which the actions are performed cannot be numbered i.e. is not discrete, is said to be continuous. Self-driving cars are an example of continuous environments as their actions are driving, parking, etc. which cannot be numbered.

7. Episodic vs Sequential:

In an Episodic task environment, each of the agent's actions is divided into atomic incidents or episodes. There is no dependency between current and previous incidents. In each incident, an agent receives input from the environment and then performs the corresponding action.

In a Sequential environment, the previous decisions can affect all future decisions. The next action of the agent depends on what action he has taken previously and what action he is supposed to take in the future.

Example:

Checkers- Where the previous move can affect all the following moves.

8. Known vs Unknown

In a known environment, the output for all probable actions is given. Obviously, in case of unknown environment, for an agent to make a decision, it has to gain knowledge about how the environment works.

Problem Formulations:

Problem formulation is the process of deciding what actions and states to consider, given a goal. The process of looking for a sequence of actions that reaches the goal is called search. A search algorithm takes a problem as input and returns a solution in the form of an action sequence.

Problems and Solutions:

A problem can be defined formally by five components:

1. The **initial state** that the agent starts in.
2. A description of the possible **actions** available to the agent.
3. A description of what each action does; the formal name for this is the **transition model**.

Together, the initial state, actions, and transition model implicitly define the state space of the problem—the set of all states reachable from the initial state by any sequence of actions. The state space forms a directed network or graph in which the nodes are states and the links between nodes are actions.

4. The **goal test**, which determines whether a given state is a goal state. Sometimes there is an explicit set of possible goal states, and the test simply checks whether the given state is one of them

5. A **path cost** function that assigns a numeric cost to each path. The problem-solving agent chooses a cost function that reflects its own performance measure.


Example Problems:

The problem-solving approach has been applied to a vast array of task environments. A toy problem is intended to illustrate or exercise various problem-solving methods. It can be given a concise, exact description and hence is usable by different researchers to compare the performance of algorithms. Let us consider the example of **vacuum world**.

This can be formulated as a problem as follows:

States: The state is determined by both the agent location and the dirt locations. The agent is in one of two locations, each of which might or might not contain dirt. Thus,

there are $2 \times 2 \times 2 = 8$ possible world states. A larger environment with n locations has $n \cdot 2^n$ states.



Initial state: Any state can be designated as the initial state.

Actions: In this simple environment, each state has just three actions: Left, Right, and Suck. Larger environments might also include Up and Down.

Transition model: The actions have their expected effects, except that moving Left in the leftmost square, moving Right in the rightmost square, and Sucking in a clean square have no effect. The complete state space is shown in the below Figure.

Goal test: This checks whether all the squares are clean.

Path cost: Each step costs 1, so the path cost is the number of steps in the path.

