UNIT-1

Introduction

What is Artificial Intelligence (AI)?

In today's world, technology is growing very fast, and we are getting in touch with different new technologies day by day.

Here, one of the booming technologies of computer science is Artificial Intelligence which is ready to create a new revolution in the world by making intelligent machines. The Artificial Intelligence is now all around us. It is currently working with a variety of subfields, ranging from general to specific, such as self-driving cars, playing chess, proving theorems, playing music, Painting, etc.

AI is one of the fascinating and universal fields of Computer science which has a great scope in future. AI holds a tendency to cause a machine to work as a human.

Artificial Intelligence is composed of two words **Artificial** and **Intelligence**, where Artificial defines *"man-made,"* and intelligence defines *"thinking power"*, hence AI means *"a man-made thinking power."*

So, we can define AI as:

"It is a branch of computer science by which we can create intelligent machines which can behave like a human, think like humans, and able to make decisions."

Artificial Intelligence exists when a machine can have human based skills such as learning, reasoning, and solving problems

With Artificial Intelligence you do not need to preprogram a machine to do some work, despite that you can create a machine with programmed algorithms which can work with own intelligence, and that is the awesomeness of AI.

It is believed that AI is not a new technology, and some people says that as per Greek myth, there were Mechanical men in early days which can work and behave like humans.

Why Artificial Intelligence?

Before Learning about Artificial Intelligence, we should know that what is the importance of AI and why should we learn it. Following are some main reasons to learn about AI:

- With the help of AI, you can create such software or devices which can solve real-world problems very easily and with accuracy such as health issues, marketing, traffic issues, etc.
- With the help of AI, you can create your personal virtual Assistant, such as Cortana, Google Assistant, Siri, etc.
- With the help of AI, you can build such Robots which can work in an environment where survival of humans can be at risk.

- AI opens a path for other new technologies, new devices, and new Opportunities.

Goals of Artificial Intelligence

Following are the main goals of Artificial Intelligence:

2. Replicate human intelligence
3. Solve Knowledge-intensive tasks
4. An intelligent connection of perception and action
5. Building a machine which can perform tasks that requires human intelligence such as:
    - Proving a theorem
    - Playing chess
    - Plan some surgical operation
    - Driving a car in traffic
6. Creating some system which can exhibit intelligent behavior, learn new things by itself, demonstrate, explain, and can advise to its user.

AI Problems

The following are a few major problems associated with Artificial Intelligence and its possible solutions.

1. Job Loss Problem

Job loss concerns related to Artificial Intelligence have been the subject of numerous business cases and academic studies. As per an Oxford Study, more than 47% of American jobs will be under threat due to automation by the mid-2030s. According to the World Economic Forum, Artificial Intelligence automation will replace over 75 million jobs by 2022. Some of the figures are even more daunting.

According to another Mckinsey report, AI-based robots could replace 30% of the global workforce. As per the AI expert and Venture Capitalist Kai-Fu Lee, 40% of the world's jobs will be replaced by AI-based bots in the next 10-15 years. Low-income and low-skilled workers will be the worst hit by this change. As AI becomes smarter by the day, even the High paid, High skilled workers become more vulnerable to job losses as, given the high cost of skilled workers, the companies get better margins by automating their work. However, these issues related to Job loss and wages can be addressed by focusing on the following measures.

- Overhauling the education system and giving more focus on skills like Critical Thinking, Creativity, and Innovation, as these skills are hard to replicate.
- To align it with industry demand, we can increase public and private investment in developing human capital.
- Improving the condition of the labor market by bridging the demand-supply gap and giving impetus to the gig economy.

2. Safety Problem

Safety issues associated with Artificial Intelligence have always stirred up a lot of furor. When experts like Elon Musk, Stephen Hawking, and Bill Gates, among others, express concern about AI safety, we should pay heed to its safety issues. Artificial Intelligence has encountered instances where Twitter Chatbot began spewing abusive and Pro-Nazi sentiments, and Facebook AI bots started interacting in an incomprehensible language, resulting in the project's shutdown.

There are grave concerns about Artificial Intelligence doing something harmful to humankind. A case in point is autonomous weapons which can be programmed to kill other humans. There are also imminent concerns with AI forming a "Mind of their Own" and doesn't value human life. Deploying such weapons will make undoing their repercussions very difficult. We can take the following measures to mitigate these concerns.

- We need to have strong regulations, especially when it comes to the creation or experimentation of Autonomous weapons.
- To ensure that no one gets involved in the rat race, we need global cooperation on issues concerning such kinds of weapons.
- Complete transparency in the system where such technologies have been experimented with is essential to ensure their safe usage.

3. Trust-Related Problem

As Artificial Intelligence algorithms become more powerful by the day, it also brings several trust-related issues on their ability to make fair decisions for the betterment of humankind. With AI slowly reaching human-level cognitive abilities, the trust issue becomes all the more significant. There are several applications where AI operates as a black box. For example- in High-Frequency trading, even the Program developers don't understand the basis on which AI executed the trade. Amazon's AI-based algorithm for same-day delivery inadvertently exhibited bias against black neighborhoods, providing a striking example. Another example was Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), where the Artificial Intelligence algorithm, while profiling suspects, was biased against the black community.

- We can take the following measures to bridge trust-related issues in Artificial Intelligence.
- All the major Artificial Intelligence providers need to set up guiding rules and principles related to trust and transparency in AI implementation. All stakeholders involved in Artificial Intelligence development and usage must religiously follow these principles.
- The stakeholders should be aware of the bias which inherently comes with AI algorithms and should have a robust bias detection mechanism and ways to handle it
- Awareness is another key factor that plays a major role in bridging the trust gap. We should sensitize users about AI operations, its capabilities, and even the shortcomings associated with Artificial Intelligence.

4. Computation Problem

Artificial Intelligence algorithm analyzes the humongous data that require immense computational power. So far, the problem has been dealt with with the help of Cloud Computing and Parallel Processing. However, as the amount of data increases and more complex deep learning algorithm comes into the mainstream, the present-day computational power will not be enough to cater to the complex requirement. We will need more storage and computational power to handle crunching exabytes and Zettabytes of data.

Quantum Computing can address the processing speed problem in the medium to long terms

It is based on Quantum theory concepts that might be the answer to solving computation power challenges. Quantum computing is 100 Million times faster than a normal computer we use at home. Although currently, it is in the research and experimental stage. Experts estimate that we can see its mainstream implementation in the next 10-15 years.

Foundation of Artificial Intelligence

Artificial Intelligence has become a buzzword in the technological world, with the potential to transform the way we live and work. The foundation of AI lies in the Base i, which refers to the fundamental building blocks that make up the technology. These building blocks include machine learning, natural language processing, computer vision, and robotics, among others. Together, these components form the backbone of AI, allowing machines to learn, adapt, and improve over time.

1. Machine learning: Machine learning is a subset of AI that focuses on the development of algorithms that enable machines to learn from data and make predictions or decisions without being explicitly programmed. For instance, an e-commerce platform uses machine learning algorithms to recommend products to customers based on their browsing history, purchases, and search history.

2. Natural language Processing: Natural Language Processing (NLP) refers to the ability of machines to understand, interpret, and generate human language. NLP is essential for chatbots, virtual assistants, and voice recognition systems that allow users to interact with machines using natural language.

3. Computer vision: Computer Vision involves training machines to interpret and understand visual data from the world around them. With computer vision, machines can recognize objects, faces, and even emotions, which is critical for applications such as facial recognition, surveillance, and self-driving cars.

4. Robotics: Robotics is the application of AI in the development of robots that can perform tasks autonomously. This includes everything from industrial robots used in manufacturing to autonomous drones and self-driving cars.

AI has the potential to transform industries and solve some of the world's most pressing problems. For instance, AI can help in the diagnosis and treatment of diseases, reduce carbon emissions, and improve crop yields. However, the success of AI depends on the strength of the Base i, which requires continued investment in research and development.

History of Artificial Intelligence

Artificial Intelligence is not a new word and not a new technology for researchers. This technology is much older than you would imagine. Even there are the myths of Mechanical men in Ancient Greek and Egyptian Myths. Following are some milestones in the history of AI which defines the journey from the AI generation to till date development.

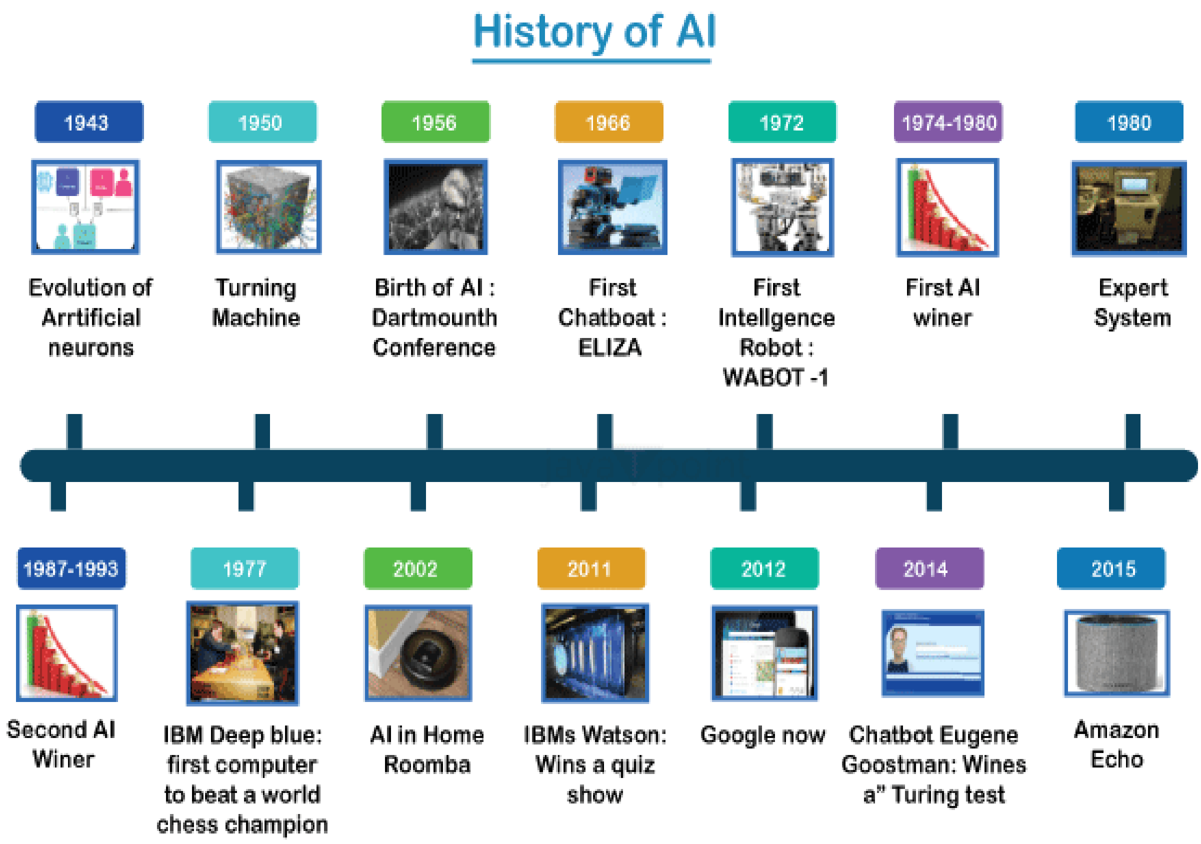

Maturation of Artificial Intelligence (1943-1952)

Between 1943 and 1952, there was notable progress in the expansion of artificial intelligence (AI). Throughout this period, AI transitioned from a mere concept to tangible experiments and practical applications. Here are some key events that happened during this period:

- **Year 1943:** The first work which is now recognized as AI was done by Warren McCulloch and Walter pits in 1943. They proposed a model of **artificial neurons**.
- **Year 1949:** Donald Hebb demonstrated an updating rule for modifying the connection strength between neurons. His rule is now called Hebbian learning.
- **Year 1950:** The Alan Turing who was an English mathematician and pioneered Machine learning in 1950. Alan Turing publishes **"Computing Machinery and Intelligence"** in which he proposed a test. The test can check the machine's ability to exhibit intelligent behavior equivalent to human intelligence, called a **Turing test**.
- **Year 1951:** Marvin Minsky and Dean Edmonds created the initial artificial neural network (ANN) named SNARC. They utilized 3,000 vacuum tubes to mimic a network of 40 neurons.

The birth of Artificial Intelligence (1952-1956)

From 1952 to 1956, AI surfaced as a unique domain of investigation. During this period, pioneers and forward-thinkers commenced the groundwork for what would ultimately transform into a revolutionary technological domain. Here are notable occurrences from this era:

- **Year 1952:** Arthur Samuel pioneered the creation of the Samuel Checkers-Playing Program, which marked the world's first self-learning program for playing games.
- **Year 1955:** An Allen Newell and Herbert A. Simon created the "first artificial intelligence program"Which was named as **"Logic Theorist"**. This program had proved 38 of 52 Mathematics theorems, and find new and more elegant proofs for some theorems.
- **Year 1956:** The word "Artificial Intelligence" first adopted by American Computer scientist John McCarthy at the Dartmouth Conference. For the first time, AI coined as an academic field.

At that time high-level computer languages such as FORTRAN, LISP, or COBOL were invented. And the enthusiasm for AI was very high at that time.

The golden years-Early enthusiasm (1956-1974)

The period from 1956 to 1974 is commonly known as the "Golden Age" of artificial intelligence (AI). In this timeframe, AI researchers and innovators were filled with enthusiasm and achieved remarkable advancements in the field. Here are some notable events from this era:

- **Year 1958:** During this period, Frank Rosenblatt introduced the perceptron, one of the early artificial neural networks with the ability to learn from data. This invention laid the foundation for modern neural networks. Simultaneously, John McCarthy developed the Lisp programming language, which swiftly found favor within the AI community, becoming highly popular among developers.
- **Year 1959:** Arthur Samuel is credited with introducing the phrase "machine learning" in a pivotal paper in which he proposed that computers could be programmed to surpass their creators in performance. Additionally, Oliver Selfridge made a notable contribution to machine learning with his publication "Pandemonium: A Paradigm for Learning." This work outlined a model capable of self-improvement, enabling it to discover patterns in events more effectively.
- **Year 1964:** During his time as a doctoral candidate at MIT, Daniel Bobrow created STUDENT, one of the early programs for natural language processing (NLP), with the specific purpose of solving algebra word problems.
- **Year 1965:** The initial expert system, Dendral, was devised by Edward Feigenbaum, Bruce G. Buchanan, Joshua Lederberg, and Carl Djerassi. It aided organic chemists in identifying unfamiliar organic compounds.
- **Year 1966:** The researchers emphasized developing algorithms that can solve mathematical problems. Joseph Weizenbaum created the first chatbot in 1966, which was named ELIZA. Furthermore, Stanford Research Institute created Shakey, the earliest mobile intelligent robot incorporating AI, computer vision, navigation, and NLP. It can be considered a precursor to today's self-driving cars and drones.

- **Year 1968:** Terry Winograd developed SHRDLU, which was the pioneering multimodal AI capable of following user instructions to manipulate and reason within a world of blocks.
- **Year 1969:** Arthur Bryson and Yu-Chi Ho outlined a learning algorithm known as backpropagation, which enabled the development of multilayer artificial neural networks. This represented a significant advancement beyond the perceptron and laid the groundwork for deep learning. Additionally, Marvin Minsky and Seymour Papert authored the book "Perceptrons," which elucidated the constraints of basic neural networks. This publication led to a decline in neural network research and a resurgence in symbolic AI research.
- **Year 1972:** The first intelligent humanoid robot was built in Japan, which was named WABOT-1.
- **Year 1973:** James Lighthill published the report titled "Artificial Intelligence: A General Survey," resulting in a substantial reduction in the British government's backing for AI research.

The first AI winter (1974-1980)

The initial AI winter, occurring from 1974 to 1980, is known as a tough period for artificial intelligence (AI). During this time, there was a substantial decrease in research funding, and AI faced a sense of letdown.

- The duration between years 1974 to 1980 was the first AI winter duration. AI winter refers to the time period where computer scientist dealt with a severe shortage of funding from government for AI researches.
- During AI winters, an interest of publicity on artificial intelligence was decreased.

A boom of AI (1980-1987)

Between 1980 and 1987, AI underwent a renaissance and newfound vitality after the challenging era of the First AI Winter. Here are notable occurrences from this timeframe:

- In 1980, the first national conference of the American Association of Artificial Intelligence was held at Stanford University.
- **Year 1980:** After AI's winter duration, AI came back with an "Expert System". Expert systems were programmed to emulate the decision-making ability of a human expert. Additionally, Symbolics Lisp machines were brought into commercial use, marking the onset of an AI resurgence. However, in subsequent years, the Lisp machine market experienced a significant downturn.
- **Year 1981:** Danny Hillis created parallel computers tailored for AI and various computational functions, featuring an architecture akin to contemporary GPUs.
- **Year 1984:** Marvin Minsky and Roger Schank introduced the phrase "AI winter" during a gathering of the Association for the Advancement of Artificial Intelligence. They cautioned the

business world that exaggerated expectations about AI would result in disillusionment and the eventual downfall of the industry, which indeed occurred three years later.

- **Year 1985:** Judea Pearl introduced Bayesian network causal analysis, presenting statistical methods for encoding uncertainty in computer systems.

The second AI winter (1987-1993)

- The duration between the years 1987 to 1993 was the second AI Winter duration.
- Again Investors and government stopped in funding for AI research as due to high cost but not efficient result. The expert system such as XCON was very cost effective.

The emergence of intelligent agents (1993-2011)

Between 1993 and 2011, there were significant leaps forward in artificial intelligence (AI), particularly in the development of intelligent computer programs. During this era, AI professionals shifted their emphasis from attempting to match human intelligence to crafting pragmatic, ingenious software tailored to specific tasks. Here are some noteworthy occurrences from this timeframe:

- **Year 1997:** In 1997, IBM's Deep Blue achieved a historic milestone by defeating world chess champion Gary Kasparov, marking the first time a computer triumphed over a reigning world chess champion. Moreover, Sepp Hochreiter and Jürgen Schmidhuber introduced the Long Short-Term Memory recurrent neural network, revolutionizing the capability to process entire sequences of data such as speech or video.
- **Year 2002:** for the first time, AI entered the home in the form of Roomba, a vacuum cleaner.
- **Year 2006:** AI came into the Business world till the year 2006. Companies like Facebook, Twitter, and Netflix also started using AI.
- **Year 2009:** Rajat Raina, Anand Madhavan, and Andrew Ng released the paper titled "Utilizing Graphics Processors for Extensive Deep Unsupervised Learning," introducing the concept of employing GPUs for the training of expansive neural networks.
- **Year 2011:** Jürgen Schmidhuber, Dan Claudiu Cire?an, Ueli Meier, and Jonathan Masci created the initial CNN that attained "superhuman" performance by emerging as the victor in the German Traffic Sign Recognition competition. Furthermore, Apple launched Siri, a voice-activated personal assistant capable of generating responses and executing actions in response to voice commands.

Deep learning, big data and artificial general intelligence (2011-present)

From 2011 to the present moment, significant advancements have unfolded within the artificial intelligence (AI) domain. These achievements can be attributed to the amalgamation of deep learning, extensive data application, and the ongoing quest for artificial general intelligence (AGI). Here are notable occurrences from this timeframe:

- **Year 2011:** In 2011, IBM's Watson won Jeopardy, a quiz show where it had to solve complex questions as well as riddles. Watson had proved that it could understand natural language and can solve tricky questions quickly.
- **Year 2012:** Google launched an Android app feature, "Google Now", which was able to provide information to the user as a prediction. Further, Geoffrey Hinton, Ilya Sutskever, and Alex Krizhevsky presented a deep CNN structure that emerged victorious in the ImageNet challenge, sparking the proliferation of research and application in the field of deep learning.
- **Year 2013:** China's Tianhe-2 system achieved a remarkable feat by doubling the speed of the world's leading supercomputers to reach 33.86 petaflops. It retained its status as the world's fastest system for the third consecutive time. Furthermore, DeepMind unveiled deep reinforcement learning, a CNN that acquired skills through repetitive learning and rewards, ultimately surpassing human experts in playing games. Also, Google researcher Tomas Mikolov and his team introduced Word2vec, a tool designed to automatically discern the semantic connections among words.
- **Year 2014:** In the year 2014, Chatbot "Eugene Goostman" won a competition in the infamous "Turing test." Whereas Ian Goodfellow and his team pioneered generative adversarial networks (GANs), a type of machine learning framework employed for producing images, altering pictures, and crafting deepfakes, and Diederik Kingma and Max Welling introduced variational autoencoders (VAEs) for generating images, videos, and text. Also, Facebook engineered the DeepFace deep learning facial recognition system, capable of identifying human faces in digital images with accuracy nearly comparable to human capabilities.
- **Year 2016:** DeepMind's AlphaGo secured victory over the esteemed Go player Lee Sedol in Seoul, South Korea, prompting reminiscence of the Kasparov chess match against Deep Blue nearly two decades earlier.Whereas Uber initiated a pilot program for self-driving cars in Pittsburgh, catering to a limited group of users.
- **Year 2018:** The "Project Debater" from IBM debated on complex topics with two master debaters and also performed extremely well.
- Google has demonstrated an AI program, "Duplex," which was a virtual assistant that had taken hairdresser appointments on call, and the lady on the other side didn't notice that she was talking with the machine.
- **Year 2021:** OpenAI unveiled the Dall-E multimodal AI system, capable of producing images based on textual prompts.
- **Year 2022:** In November, OpenAI launched ChatGPT, offering a chat-oriented interface to its GPT-3.5 LLM.

Now AI has developed to a remarkable level. The concept of Deep learning, big data, and data science are now trending like a boom. Nowadays companies like Google, Facebook, IBM, and Amazon are working with AI and creating amazing devices. The future of Artificial Intelligence is inspiring and will come with high intelligence.

Intelligent Agents:

An intelligent agent is an autonomous entity which acts upon an environment using sensors and actuators for achieving goals. An intelligent agent may learn from the environment to achieve their goals. A thermostat is an example of an intelligent agent.

Following are the main four rules for an AI agent:

- **Rule 1:** An AI agent must have the ability to perceive the environment.
- **Rule 2:** The observation must be used to make decisions.
- **Rule 3:** Decision should result in an action.
- **Rule 4:** The action taken by an AI agent must be a rational action.

Agents

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:

- Simple Reflex Agent
- Model-based reflex agent
- Goal-based agents
- Utility-based agent
- Learning agent

1. Simple Reflex agent:

- The Simple reflex agents are the simplest agents. These agents take decisions on the basis of the current percepts and ignore the rest of the percept history.
- These agents only succeed in the fully observable environment.
- The Simple reflex agent does not consider any part of percepts history during their decision and action process.
- The Simple reflex agent works on Condition-action rule, which means it maps the current state to action. Such as a Room Cleaner agent, it works only if there is dirt in the room.
- Problems for the simple reflex agent design approach:
- They have very limited intelligence
- They do not have knowledge of non-perceptual parts of the current state
- Mostly too big to generate and to store.
- Not adaptive to changes in the environment.

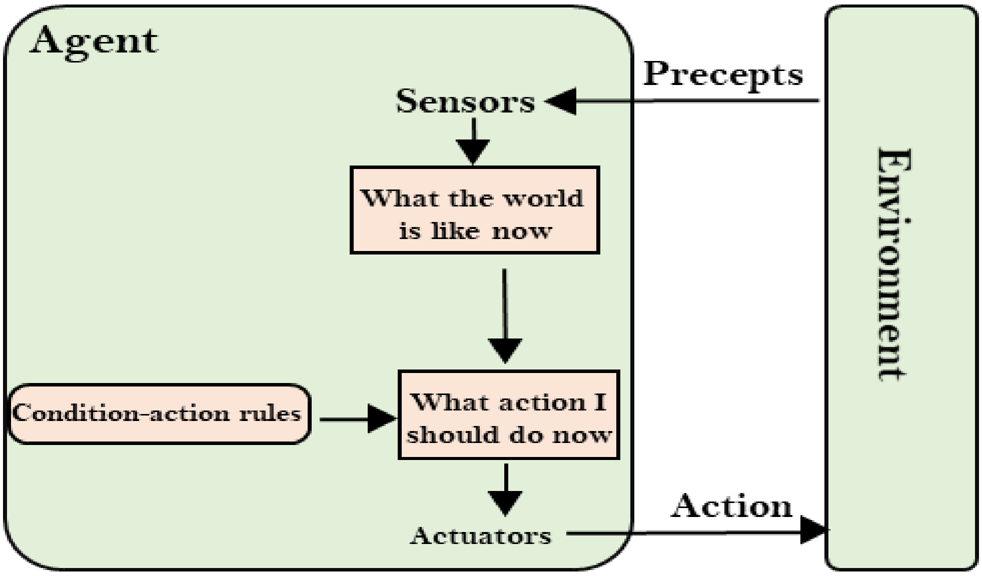

2. Model-based reflex agent

- The Model-based agent can work in a partially observable environment, and track the situation.
- A model-based agent has two important factors:
  - **Model:** It is knowledge about "how things happen in the world," so it is called a Model-based agent.
  - **Internal State:** It is a representation of the current state based on percept history.
- These agents have the model, "which is knowledge of the world" and based on the model they perform actions.
- Updating the agent state requires information about:
  a. How the world evolves
  b. How the agent's action affects the world.

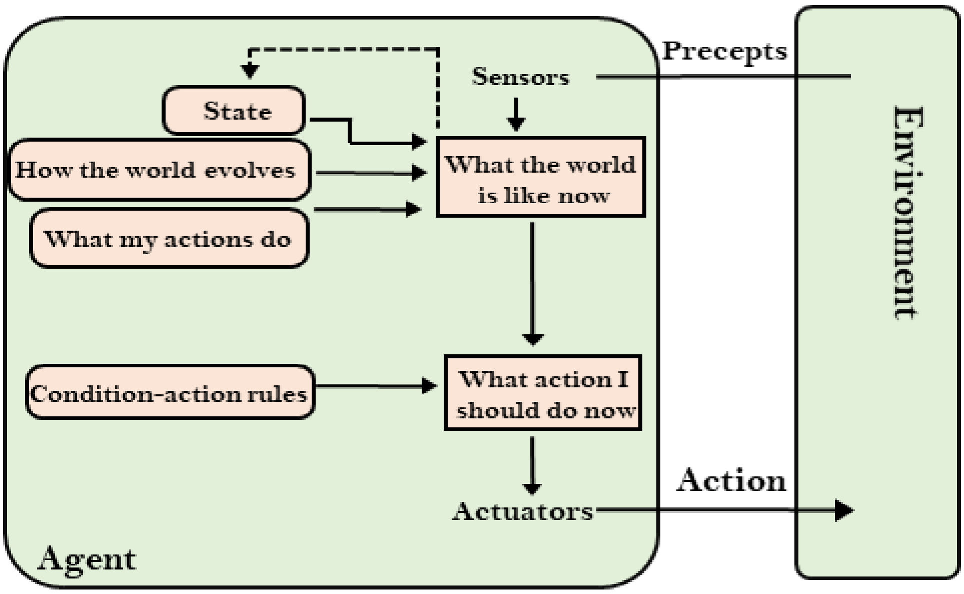

3. Goal-based agents

- The knowledge of the current state environment is not always sufficient to decide for an agent to what to do.
- The agent needs to know its goal which describes desirable situations.
- Goal-based agents expand the capabilities of the model-based agent by having the "goal" information.
- They choose an action, so that they can achieve the goal.
- These agents may have to consider a long sequence of possible actions before deciding whether the goal is achieved or not. Such considerations of different scenario are called searching and planning, which makes an agent proactive.

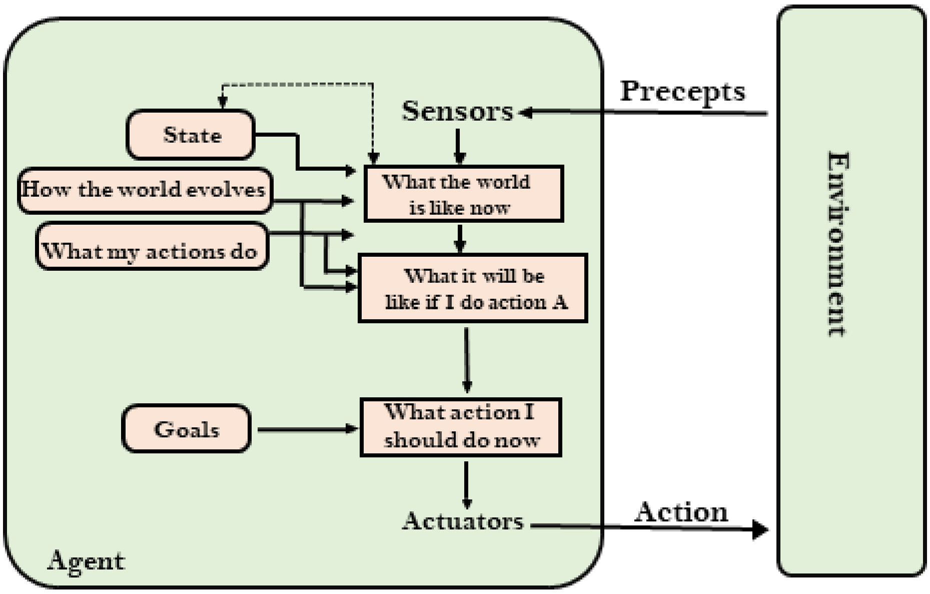

4. Utility-based agents

- These agents are similar to the goal-based agent but provide an extra component of utility measurement which makes them different by providing a measure of success at a given state.
- Utility-based agent act based not only goals but also the best way to achieve the goal.
- The Utility-based agent is useful when there are multiple possible alternatives, and an agent has to choose in order to perform the best action.
- The utility function maps each state to a real number to check how efficiently each action achieves the goals.

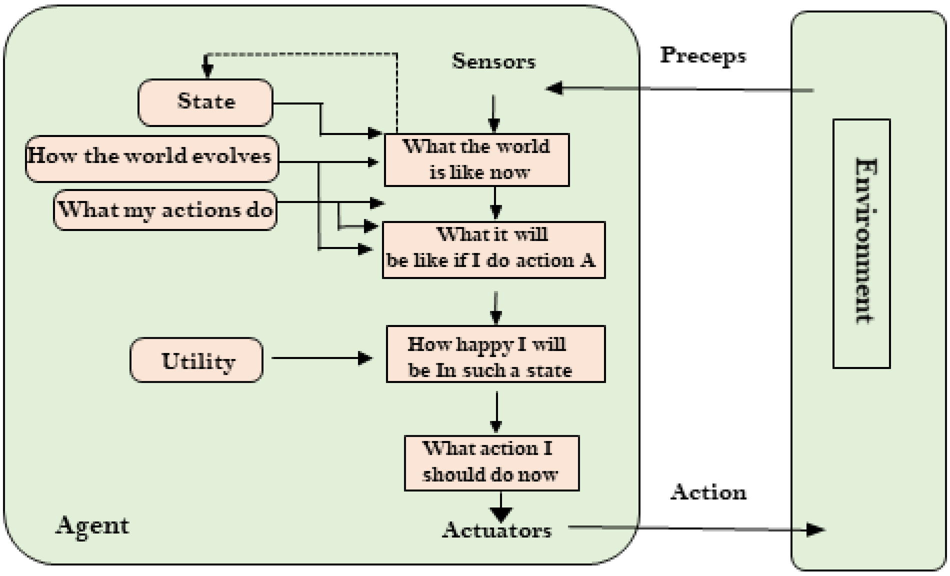

5. Learning Agents

- A learning agent in AI is the type of agent which can learn from its past experiences, or it has learning capabilities.
- It starts to act with basic knowledge and then able to act and adapt automatically through learning.
- A learning agent has mainly four conceptual components, which are:
  a. **Learning element:** It is responsible for making improvements by learning from environment
  b. **Critic:** Learning element takes feedback from critic which describes that how well the agent is doing with respect to a fixed performance standard.
  c. **Performance element:** It is responsible for selecting external action
  d. **Problem generator:** This component is responsible for suggesting actions that will lead to new and informative experiences.
- Hence, learning agents are able to learn, analyze performance, and look for new ways to improve the performance.

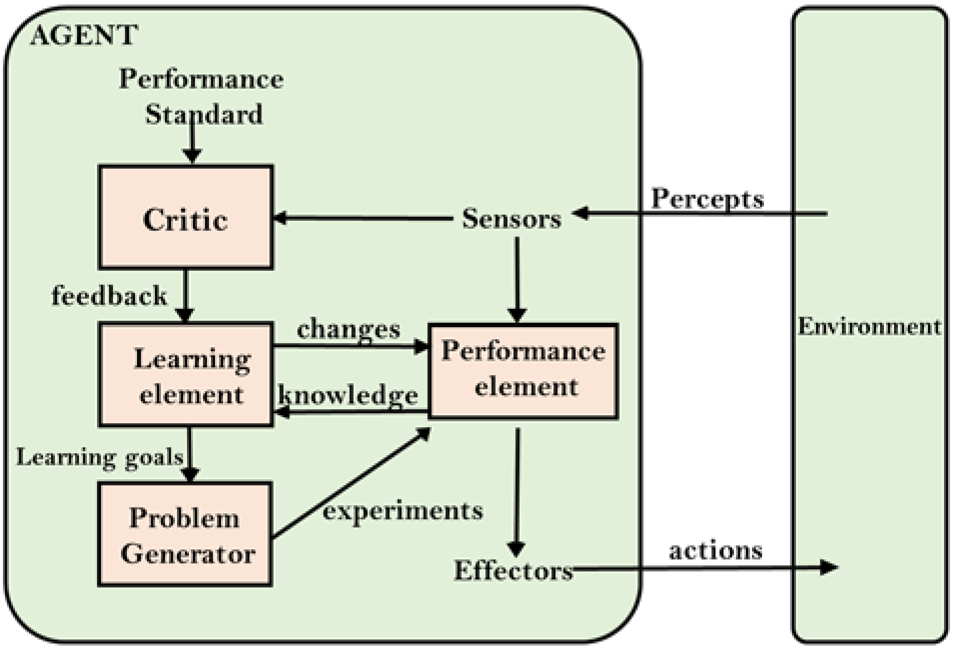

Agents in Artificial Intelligence

An AI system can be defined as the study of the rational agent and its environment. The agents sense the environment through sensors and act on their environment through actuators. An AI agent can have mental properties such as knowledge, belief, intention, etc.

What is an Agent?

An agent can be anything that perceiveits environment through sensors and act upon that environment through actuators. An Agent runs in the cycle of **perceiving**, **thinking**, and **acting**. An agent can be:

- **Human-Agent:** A human agent has eyes, ears, and other organs which work for sensors and hand, legs, vocal tract work for actuators.
- **Robotic Agent:** A robotic agent can have cameras, infrared range finder, NLP for sensors and various motors for actuators.
- **Software Agent:** Software agent can have keystrokes, file contents as sensory input and act on those inputs and display output on the screen.

Hence the world around us is full of agents such as thermostat, cellphone, camera, and even we are also agents.

Before moving forward, we should first know about sensors, effectors, and actuators.

Sensor: Sensor is a device which detects the change in the environment and sends the information to other electronic devices. An agent observes its environment through sensors.

Actuators: Actuators are the component of machines that converts energy into motion. The actuators are only responsible for moving and controlling a system. An actuator can be an electric motor, gears, rails, etc.

Effectors: Effectors are the devices which affect the environment. Effectors can be legs, wheels, arms, fingers, wings, fins, and display screen.

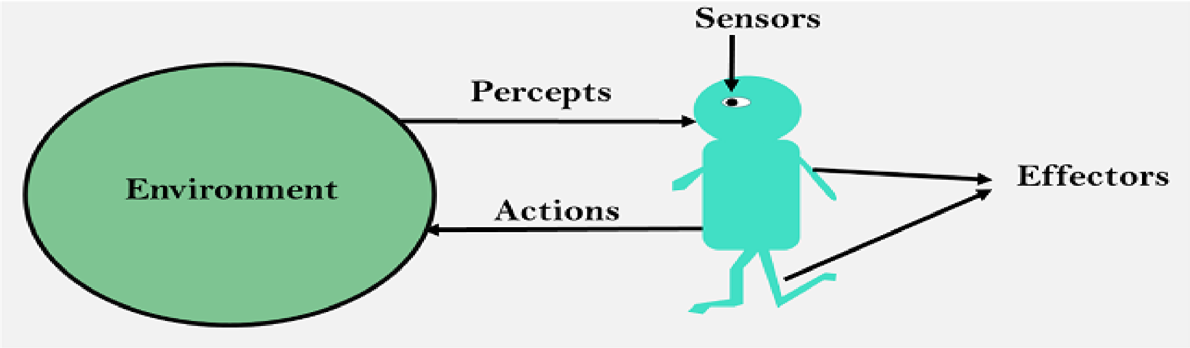

The concept of Rationality: The rationality of an agent is measured by its performance measure. Rationality can be judged on the basis of following points:

- Performance measure which defines the success criterion.
- Agent prior knowledge of its environment.
- Best possible actions that an agent can perform.
- The sequence of percepts.

A rational agent is an agent which has clear preference, models uncertainty, and acts in a way to maximize its performance measure with all possible actions.

A rational agent is said to perform the right things. AI is about creating rational agents to use for game theory and decision theory for various real-world scenarios.

For an AI agent, the rational action is most important because in AI reinforcement learning algorithm, for each best possible action, agent gets the positive reward and for each wrong action, an agent gets a negative reward.

The nature of Environments

The environment is the **Task Environment (problem)** for which the **Rational Agent is the solution.** Any task environment is characterised on the basis of PEAS.

1. **Performance** – What is the performance characteristic which would either make the agent successful or not. For example, as per the previous example clean floor, optimal energy consumption might be performance measures.
2. **Environment** – Physical characteristics and constraints expected. For example, wood floors, furniture in the way etc
3. **Actuators** – The physical or logical constructs which would take action. For example for the vacuum cleaner, these are the suction pumps
4. **Sensors** – Again physical or logical constructs which would sense the environment.

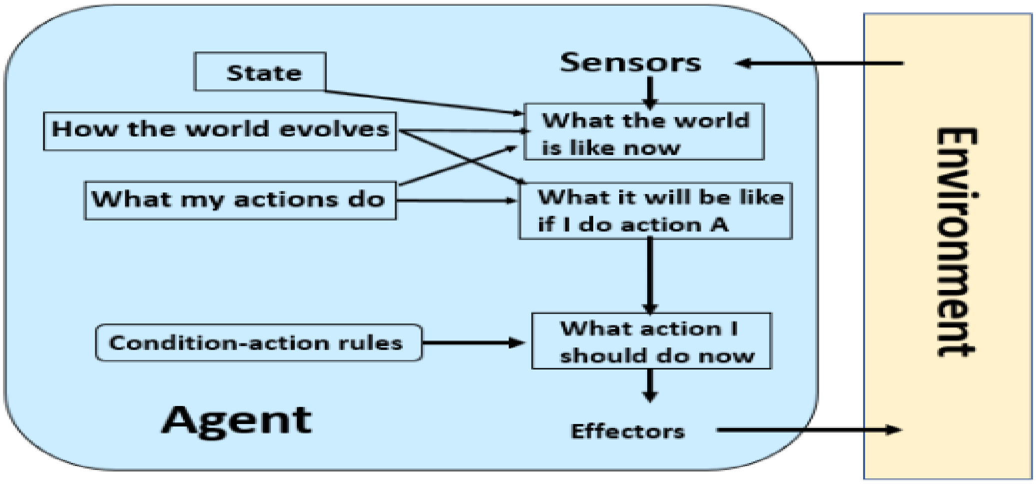

Rational Agents could be physical agents like the one described above or it could also be a program that operates in a non-physical environment like an operating system. Imagine a bot web site operator designed to scan Internet news sources and show the interesting items to its users, while selling advertising space to generate revenue.

Example, consider an online tutoring system

Agent	Performance	Environment	Actuator	Sensor
Math E learning system	SLA defined score on the test	Student, Teacher, parents	Computer display system for exercises, corrections, feedback	Keyboard, Mouse

Environments can further be classified into various buckets. This would help determine the intelligence which would need to be built in the agent. These are

- **Observable** – Full or Partial? If the agents sensors get full access then they do not need to pre-store any information. Partial may be due to inaccuracy of sensors or incomplete information about an environment, like limited access to enemy territory
- **Number of Agents** – For the vacuum cleaner, it works in a single agent environment but for driver-less taxis, every driver-less taxi is a separate agent and hence multi agent environment
- **Deterministic** – The number of unknowns in the environment which affect the predictability of the environment. For example, floor space for cleaning is mostly deterministic, the furniture is where it is most of the time but taxi driving on a road is non-deterministic.
- **Discrete** – Does the agent respond when needed or does it have to continuously scan the environment. Driver-less is continuous, online tutor is discrete
- **Static** – How often does the environment change. Can the agent learn about the environment and always do the same thing?
- **Episodic** – If the response to a certain precept is not dependent on the previous one i.e. it is stateless (static methods in Java) then it is discrete. If the decision taken now influences the future decisions then it is a sequential environment.

Hence to summarise

- An agent is something that perceives and acts in an environment.
- The performance measure evaluates the behaviour of the agent in an environment.

- A rational agent acts so as to maximise the expected value of the performance measure.
- A task environment specification includes the PEAS i.e. performance measure, the external environment, the actuators, and the sensors.
- Task environments vary along several significant dimensions. In designing an agent, the first step must always be to specify the task environment as fully as possible.

Features of Environment

As per Russell and Norvig, an environment can have various features from the point of view of an agent:

1. Fully observable vs Partially Observable
2. Static vs Dynamic
3. Discrete vs Continuous
4. Deterministic vs Stochastic
5. Single-agent vs Multi-agent
6. Episodic vs sequential
7. Known vs Unknown
8. Accessible vs Inaccessible

1. Fully observable vs Partially Observable:

- If an agent sensor can sense or access the complete state of an environment at each point in time then it is a fully observable environment, it is partially observable. For reference, Imagine a chess-playing agent. In this case, the agent can fully observe the state of the chessboard at all times. Its sensors (in this case, vision or the ability to access the board's state) provide complete information about the current position of all pieces. This is a fully observable environment because the agent has perfect information about the state of the world.
- A fully observable environment is easy as there is no need to maintain the internal state to keep track of the history of the world. For reference, Consider a self-driving car navigating a busy city. While the car has sensors like cameras, lidar, and radar, it can't see everything at all times. Buildings, other vehicles, and pedestrians can obstruct its sensors. In this scenario, the car's environment is partially observable because it doesn't have complete and constant access to all relevant information. It needs to maintain an internal state and history to make informed decisions even when some information is temporarily unavailable.
- An agent with no sensors in all environments then such an environment is called unobservable. For reference, think about an agent designed to predict earthquakes but placed in a sealed, windowless room with no sensors or access to external data. In this situation, the environment is unobservable because the agent has no way to gather information about the outside world. It can't sense any aspect of its environment, making it completely unobservable.

2. Deterministic vs Stochastic:

- If an agent's current state and selected action can completely determine the next state of the environment, then such an environment is called a deterministic environment. For reference, Chess is a classic example of a deterministic environment. In chess, the rules are well-defined, and each move made by a player has a clear and predictable outcome based on those rules. If you move a pawn from one square to another, the resulting state of the chessboard is entirely determined by that action, as is your opponent's response. There's no randomness or uncertainty in the outcomes of chess moves because they follow strict rules. In a deterministic environment like chess, knowing the current state and the actions taken allows you to completely determine the next state.
- A stochastic environment is random and cannot be determined completely by an agent. For reference, The stock market is an example of a stochastic environment. It's highly influenced by a multitude of unpredictable factors, including economic events, investor sentiment, and news. While there are patterns and trends, the exact behavior of stock prices is inherently random and cannot be completely determined by any individual or agent. Even with access to extensive data and analysis tools, stock market movements can exhibit a high degree of unpredictability. Random events and market sentiment play significant roles, introducing uncertainty.
- In a deterministic, fully observable environment, an agent does not need to worry about uncertainty.

3. Episodic vs Sequential:

- In an episodic environment, there is a series of one-shot actions, and only the current percept is required for the action. For example, Tic-Tac-Toe is a classic example of an episodic environment. In this game, two players take turns placing their symbols (X or O) on a 3x3 grid. Each move by a player is independent of previous moves, and the goal is to form a line of three symbols horizontally, vertically, or diagonally. The game consists of a series of one-shot actions where the current state of the board is the only thing that matters for the next move. There's no need for the players to remember past moves because they don't affect the current move. The game is self-contained and episodic.
- However, in a Sequential environment, an agent requires memory of past actions to determine the next best actions. For example, Chess is an example of a sequential environment. Unlike Tic-Tac-Toe, chess is a complex game where the outcome of each move depends on a sequence of previous moves. In chess, players must consider the history of the game, as the current position of pieces, previous moves, and potential future moves all influence the best course of action. To play chess effectively, players need to maintain a memory of past actions, anticipate future moves, and plan their strategies accordingly. It's a sequential environment because the sequence of actions and the history of the game significantly impact decision-making.

4. Single-agent vs Multi-agent:

- If only one agent is involved in an environment, and operating by itself then such an environment is called a single-agent environment. For example, Solitaire is a classic example of a single-agent environment. When you play Solitaire, you're the only agent involved. You make all the decisions and actions to achieve a goal, which is to arrange a deck of cards in a specific way. There are no other agents or players interacting with you. It's a solitary game where the outcome depends solely on your decisions and moves. In this single-agent environment, the agent doesn't need to consider the actions or decisions of other entities.
- However, if multiple agents are operating in an environment, then such an environment is called a multi-agent environment. For reference, A soccer match is an example of a multi-agent environment. In a soccer game, there are two teams, each consisting of multiple players (agents). These players work together to achieve common goals (scoring goals and preventing the opposing team from scoring). Each player has their own set of actions and decisions, and they interact with both their teammates and the opposing team. The outcome of the game depends on the coordinated actions and strategies of all the agents on the field. It's a multi-agent environment because there are multiple autonomous entities (players) interacting in a shared environment.
- The agent design problems in the multi-agent environment are different from single-agent environments.

5. Static vs Dynamic:

- If the environment can change itself while an agent is deliberating then such an environment is called a dynamic environment it is called a static environment.
- Static environments are easy to deal with because an agent does not need to continue looking at the world while deciding on an action. For reference, A crossword puzzle is an example of a static environment. When you work on a crossword puzzle, the puzzle itself doesn't change while you're thinking about your next move. The arrangement of clues and empty squares remains constant throughout your problem-solving process. You can take your time to deliberate and find the best word to fill in each blank, and the puzzle's state remains unaltered during this process. It's a static environment because there are no changes in the puzzle based on your deliberations.
- However, for a dynamic environment, agents need to keep looking at the world at each action. For reference, Taxi driving is an example of a dynamic environment. When you're driving a taxi, the environment is constantly changing. The road conditions, traffic, pedestrians, and other vehicles all contribute to the dynamic nature of this environment. As a taxi driver, you need to keep a constant watch on the road and adapt your actions in real time based on the changing circumstances. The environment can change rapidly, requiring your continuous

attention and decision-making. It's a dynamic environment because it evolves while you're deliberating and taking action.

6. Discrete vs Continuous:

- If in an environment, there are a finite number of percepts and actions that can be performed within it, then such an environment is called a discrete environment it is called a continuous environment.
- Chess is an example of a discrete environment. In chess, there are a finite number of distinct chess pieces (e.g., pawns, rooks, knights) and a finite number of squares on the chessboard. The rules of chess define clear, discrete moves that a player can make. Each piece can be in a specific location on the board, and players take turns making individual, well-defined moves. The state of the chessboard is discrete and can be described by the positions of the pieces on the board.
- Controlling a robotic arm to perform precise movements in a factory setting is an example of a continuous environment. In this context, the robot arm's position and orientation can exist along a continuous spectrum. There are virtually infinite possible positions and orientations for the robotic arm within its workspace. The control inputs to move the arm, such as adjusting joint angles or applying forces, can also vary continuously. Agents in this environment must operate within a continuous state and action space, and they need to make precise, continuous adjustments to achieve their goals.

7. Known vs Unknown:

- Known and unknown are not actually a feature of an environment, but it is an agent's state of knowledge to perform an action.
- In a known environment, the results of all actions are known to the agent. While in an unknown environment, an agent needs to learn how it works in order to perform an action.
- It is quite possible for a known environment to be partially observable and an Unknown environment to be fully observable.
- The opening theory in chess can be considered as a known environment for experienced chess players. Chess has a vast body of knowledge regarding opening moves, strategies, and responses. Experienced players are familiar with established openings, and they have studied various sequences of moves and their outcomes. When they make their initial moves in a game, they have a good understanding of the potential consequences based on their knowledge of known openings.
- Imagine a scenario where a rover or drone is sent to explore an alien planet with no prior knowledge or maps of the terrain. In this unknown environment, the agent (rover or drone) has to explore and learn about the terrain as it goes along. It doesn't have prior knowledge of the

landscape, potential hazards, or valuable resources. The agent needs to use sensors and data it collects during exploration to build a map and understand how the terrain works. It operates in an unknown environment because the results and consequences of its actions are not initially known, and it must learn from its experiences.

8. Accessible vs Inaccessible:

- If an agent can obtain complete and accurate information about the state's environment, then such an environment is called an Accessible environment else it is called inaccessible.
- For example, Imagine an empty room equipped with highly accurate temperature sensors. These sensors can provide real-time temperature measurements at any point within the room. An agent placed in this room can obtain complete and accurate information about the temperature at different locations. It can access this information at any time, allowing it to make decisions based on the precise temperature data. This environment is accessible because the agent can acquire complete and accurate information about the state of the room, specifically its temperature.
- For example, consider a scenario where a satellite in space is tasked with monitoring a specific event taking place on Earth, such as a natural disaster or a remote area's condition. While the satellite can capture images and data from space, it cannot access fine-grained information about the event's details. For example, it may see a forest fire occurring but cannot determine the exact temperature at specific locations within the fire or identify individual objects on the ground. The satellite's observations provide valuable data, but the environment it is monitoring (Earth) is vast and complex, making it impossible to access complete and detailed information about all aspects of the event. In this case, the Earth's surface is an inaccessible environment for obtaining fine-grained information about specific events.

Structure of Agents

Structure of an AI Agent

The task of AI is to design an agent program which implements the agent function. The structure of an intelligent agent is a combination of architecture and agent program. It can be viewed as:

Agent = Architecture + Agent program

Following are the main three terms involved in the structure of an AI agent:

Architecture: Architecture is machinery that an AI agent executes on.

Agent Function: Agent function is used to map a percept to an action.

$$f: P^* \rightarrow A$$

Agent program: Agent program is an implementation of agent function. An agent program executes on the physical architecture to produce function f.

Agent Environment in AI

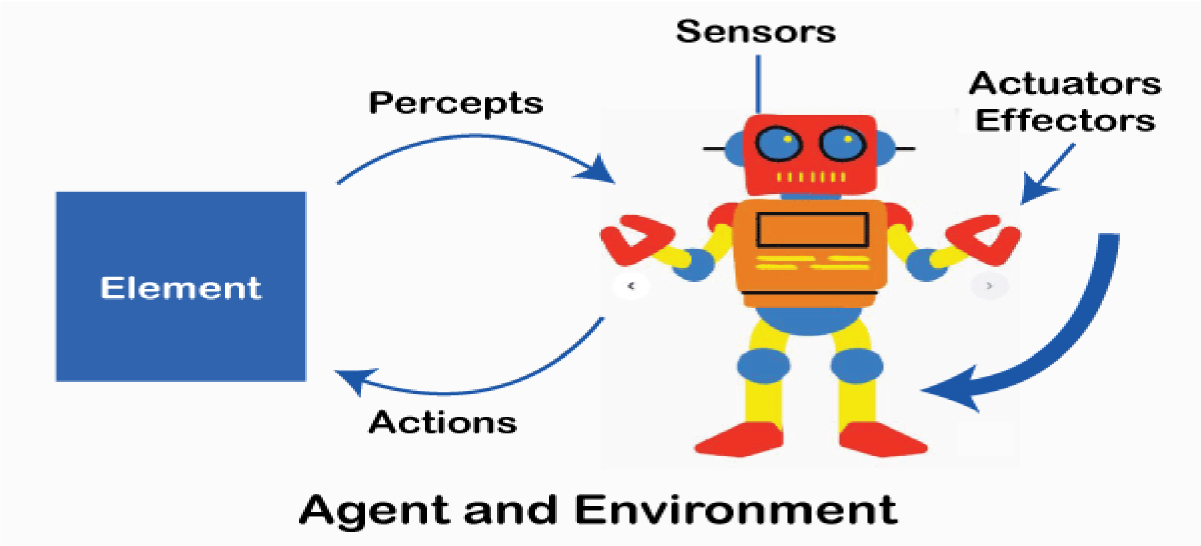

An environment is everything in the world which surrounds the agent, but it is not a part of an agent itself. An environment can be described as a situation in which an agent is present.

The environment is where agent lives, operate and provide the agent with something to sense and act upon it. An environment is mostly said to be non-feministic.

Problem-solving agents

In artificial intelligence, a problem-solving agent refers to a type of intelligent agent designed to address and solve complex problems or tasks in its environment. These agents are a fundamental concept in AI and are used in various applications, from game-playing algorithms to robotics and decision-making systems. Here are some key characteristics and components of a problem-solving agent:

1. **Perception**: Problem-solving agents typically have the ability to perceive or sense their environment. They can gather information about the current state of the world, often through sensors, cameras, or other data sources.
2. **Knowledge Base**: These agents often possess some form of knowledge or representation of the problem domain. This knowledge can be encoded in various ways, such as rules, facts, or models, depending on the specific problem.
3. **Reasoning**: Problem-solving agents employ reasoning mechanisms to make decisions and select actions based on their perception and knowledge. This involves processing information, making inferences, and selecting the best course of action.
4. **Planning**: For many complex problems, problem-solving agents engage in planning. They consider different sequences of actions to achieve their goals and decide on the most suitable action plan.
5. **Actuation**: After determining the best course of action, problem-solving agents take actions to interact with their environment. This can involve physical actions in the case of robotics or making decisions in more abstract problem-solving domains.
6. **Feedback**: Problem-solving agents often receive feedback from their environment, which they use to adjust their actions and refine their problem-solving strategies. This feedback loop helps them adapt to changing conditions and improve their performance.
7. **Learning**: Some problem-solving agents incorporate machine learning techniques to improve their performance over time. They can learn from experience, adapt their strategies, and become more efficient at solving similar problems in the future.

Problem-solving agents can vary greatly in complexity, from simple algorithms that solve straightforward puzzles to highly sophisticated AI systems that tackle complex, real-world problems. The design and implementation of problem-solving agents depend on the specific problem domain and the goals of the AI application.

Problem formulation

In the field of artificial intelligence, problem formulation plays a crucial role in the development and implementation of intelligent systems. Problem formulation refers to the process of accurately defining and stating the problems that need to be solved using artificial intelligence techniques. It involves identifying the specific task or goal, determining the available resources and constraints, and outlining the desired outcomes.

One example of problem formulation in artificial intelligence is in the domain of computer vision. The problem statement here might be to develop a system that can accurately recognize and identify objects in images or videos. In formulating this problem, one needs to consider factors such as the type and complexity of objects to be recognized, the quality of available data, and the computational resources available for processing the images.

Another example of problem formulation in artificial intelligence is in the field of natural language processing. In this case, the problem statement could be to build a system that can accurately understand and respond to human language. Formulating this problem involves defining the specific language tasks to be accomplished, such as text classification or machine translation. The formulation also involves considering the linguistic variations and complexities that exist in different languages.

Overall, problem formulation is a critical step in the development of artificial intelligence systems. The quality and accuracy of the problem formulation greatly influence the effectiveness and efficiency of the solutions that can be achieved. By carefully formulating the problems, researchers and developers can better design and implement intelligent systems that address the specific needs and challenges of various domains.

Formulating problems in AI examples

In the field of artificial intelligence, problem formulation is a crucial step in solving complex tasks. Before developing any AI system, the problem at hand needs to be carefully defined and structured. The formulation of a problem involves the creation of a clear and concise statement that describes the given task or challenge.

Artificial intelligence problems can take various forms, including classification, regression, optimization, and planning. The problem formulation provides a precise definition of the task, outlining the desired input-output relationship and the constraints or objectives involved.

Here are some examples of problem formulation in artificial intelligence:

Classification: Given a dataset of images, classify each image into predefined categories such as “cat” or “dog.” The problem formulation involves defining the features that will be used to distinguish between different categories and training a model to accurately classify the images.

Regression: Given a dataset of housing prices with various features (e.g., number of rooms, location), predict the price of a new house based on its features. The problem formulation involves identifying the relevant features, selecting a regression model, and training it on the available data.

Optimization: Given a set of constraints and objectives, find the best solution that maximizes or minimizes a predefined criterion. For example, in route optimization, the problem formulation involves defining the constraints (e.g., delivery time, fuel efficiency) and objectives (e.g., shortest distance, lowest cost) to find the optimal route.

Planning: Given a set of initial conditions, a goal state, and a set of actions, finds a sequence of actions that transform the initial state into the goal state. The problem formulation involves defining the available actions, their preconditions and effects, and designing a planning algorithm that can find the optimal sequence of actions.

These examples demonstrate the importance of problem formulation in artificial intelligence. By clearly defining the problem and its objectives, researchers and developers can design and implement effective AI systems that can tackle complex tasks and provide valuable solutions.