	Pithapur Rajahs Government College(A) Kakinada **Department of Computer Science**	**Program:** III B.Sc. (AI) **Semester :** V
Course 13	Course Name: ALGORITHMS FOR INTELLIGENT SYSTEMS	
	Hours Allocated: 3hrs/week	Credits: 3

UNIT I SYLLABUS

Introduction to Artificial Intelligence: What is Artificial Intelligence?, Definitions and goals of AI**,** Applications of AI in real life**. A Brief History of AI -** Evolution from early symbolic AI to modern deep learning**,** Key milestones and AI winters**. The State of the Art in AI:** Recent advancements (e.g., deep learning, reinforcement learning, generative AI)**,** Modern AI tools and platforms (e.g., OpenAI, TensorFlow, etc.)**,** Ethical issues and societal impacts of AI.

What is Artificial Intelligence

Artificial Intelligence (AI) refers to the technology that allows machines and computers to replicate human intelligence.

- Enables systems to perform tasks that require human-like decision-making, such as learning from data, identifying patterns, making informed choices and solving complex problems.
- Improves continuously by utilizing methods like machine learning and deep learning.
- Used in healthcare for diagnosing diseases, finance for fraud detection, e-commerce for personalized recommendations and transportation for self-driving cars. It also powers virtual assistants like Siri and Alexa, chatbots for customer support and manufacturing robots that automate production processes.

Definitions and goals of AI

AI's primary goals revolve around solving real-world problems, improving human decision-making capabilities, automating tasks, and driving innovation. These goals make AI integral to enhancing efficiency across industries, fostering creativity, and enabling groundbreaking advancements. Below, we delve into the major goals of AI, their applications, and how they bring value to modern society.

1. Problem-Solving and Decision-Making

One of the fundamental goals of Artificial Intelligence is to enable systems to solve complex problems and automate decision-making processes. AI algorithms use data-driven insights to analyze vast amounts of information, identify patterns, and make accurate, informed decisions in real time. By automating these processes, AI enhances operational efficiency and ensures faster outcomes.

Applications:

- **Chess Engines:** AI-powered engines like *Deep Blue* demonstrate decision-making capabilities by evaluating countless possible moves to find the optimal solution.
- **Predictive Analytics:** In finance, AI forecasts stock market trends, optimizes investments, and identifies potential risks.
- **Logistics Optimization:** AI tools streamline supply chains by minimizing delivery times, optimizing routes, and predicting maintenance needs.

Benefits: AI's ability to solve problems efficiently allows businesses to reduce human errors, make data-backed decisions, and boost productivity. By analyzing vast datasets, AI uncovers insights that human decision-makers might overlook, ensuring real-time solutions in critical situations.

2. Natural Language Processing (NLP)

Natural Language Processing is another critical goal of AI. It focuses on enabling machines to understand, interpret, and generate human language. NLP bridges the communication gap between humans and machines, making interactions more seamless and technology more accessible.

Applications:

- **Chatbots:** AI-powered chatbots, such as those used in customer service, provide instant and accurate responses to user queries, reducing response time and improving user satisfaction.
- **Translation Tools:** Applications like *Google Translate* overcome language barriers by accurately translating text and speech into multiple languages.
- **Voice Assistants:** AI assistants like *Siri*, *Alexa*, and *Google Assistant* use NLP to understand voice commands and perform tasks such as setting reminders, providing information, or controlling smart devices.

Benefits: NLP enhances communication between humans and machines, simplifying complex interactions. It plays a vital role in automating customer service, enabling cross-language communication, and offering personalized user experiences.

3. Machine Learning and Deep Learning

Machine Learning (ML) and Deep Learning (DL) are at the heart of AI systems, allowing machines to learn from data and improve their performance over time without human intervention. ML focuses on building models that identify patterns and make predictions, while DL uses neural networks to solve more complex problems.

Applications:

- **Image Recognition:** AI tools identify objects, people, and activities in images, which is widely used in healthcare for diagnosing diseases and in security systems for facial recognition.
- **Recommendation Systems:** Platforms like *Netflix*, *YouTube*, and *Amazon* analyze user behavior to recommend movies, videos, or products, improving customer engagement and satisfaction.
- **Fraud Detection:** In the finance industry, ML algorithms analyze transactions to detect anomalies and flag fraudulent activities.

Benefits: ML and DL automate insights, predictions, and decision-making processes across various domains. They power applications ranging from autonomous vehicles to predictive healthcare, transforming industries by providing innovative, data-driven solutions.

4. Robotics and Automation

Robotics and automation aim to combine AI with machines to perform tasks with precision, consistency, and autonomy. AI-powered robots are designed to handle labor-intensive, repetitive, or hazardous tasks, enabling improved operational efficiency and workplace safety.

Applications:

- **Manufacturing Automation:** AI-driven robots assemble products, conduct quality inspections, and manage production lines, increasing accuracy and reducing production costs.
- **Autonomous Vehicles:** Companies like *Tesla* and *Waymo* develop self-driving cars that navigate roads using AI, minimizing human intervention and reducing accidents.
- **Delivery Drones:** Logistics companies use drones to ensure faster and more efficient package deliveries, particularly in remote or inaccessible areas.

Benefits: Robotics and automation streamline workflows, enhance productivity, and reduce human involvement in dangerous tasks. Industries such as manufacturing, logistics, and healthcare rely on AI-powered robots to drive innovation and ensure operational excellence.

5. Enhancing Healthcare and Medicine

AI is revolutionizing the healthcare sector by enabling better diagnosis, treatment, and research capabilities. AI systems analyze vast medical datasets to uncover insights, improving patient outcomes while reducing healthcare costs.

Applications:

- **Medical Imaging:** AI-powered tools like *DeepMind* analyze radiology scans to detect diseases such as cancer at early stages, improving diagnosis accuracy.
- **Personalized Treatments:** AI tailors treatment plans based on patient data, ensuring customized and effective care.

- **Drug Discovery:** AI accelerates drug research and development by predicting the efficacy of compounds, reducing the time and cost of clinical trials.

Benefits: By leveraging AI, healthcare providers can diagnose diseases faster, reduce costs, and deliver personalized treatments. AI's integration into healthcare ensures improved patient care and better resource management.

6. Fostering Creativity and Innovation

Contrary to the perception that AI focuses solely on analytical tasks, one of its emerging goals is fostering creativity and innovation. AI tools assist humans in generating creative outputs across industries, enabling professionals to push the boundaries of art, design, and technology.

Applications:

- **Generative AI:** Tools like *DALL·E* and *MidJourney* create unique artworks, images, and designs based on textual descriptions, enabling artists and designers to experiment with new ideas.
- **Music Composition:** AI-powered platforms generate original music compositions, helping artists streamline their creative processes.
- **Content Creation:** AI tools assist in writing articles, brainstorming ideas, and automating storytelling, boosting productivity for writers and marketers.

Benefits: AI enhances human creativity by automating repetitive tasks, offering fresh perspectives, and accelerating innovation. It empowers professionals to focus on high-value, imaginative work, transforming industries like art, media, and entertainment.

AI's primary goals encompass solving complex challenges, improving decision-making, and driving innovation across diverse sectors. From automating workflows and enhancing healthcare to fostering creativity, AI continues to reshape industries and improve lives worldwide. Each goal plays a critical role in harnessing the true potential of Artificial Intelligence, making it a transformative force in modern society.

AI applications are rapidly expanding across various sectors, impacting daily life in numerous ways. From personalized recommendations on streaming services to fraud detection in finance and self-driving cars, AI is becoming increasingly integrated into our routines. AI-powered tools like chatbots and virtual assistants are also transforming customer service and personal assistance.

Applications of AI in Real Life

Here's a more detailed look at some key areas:

1. E-commerce and Recommendation Systems:

- AI algorithms analyze user data and browsing habits to provide personalized product recommendations on platforms like Amazon, Netflix, and YouTube.

- This enhances the user experience by suggesting relevant items and content, leading to increased engagement and sales.

2. Healthcare:

- AI is used in medical imaging analysis for early disease detection and diagnosis.
- AI-powered tools assist in drug discovery and development, potentially accelerating the process of finding new treatments.
- Robotics and automation are used in surgical procedures and patient care.

3. Finance:

- AI algorithms detect fraudulent transactions and suspicious activity in financial systems.
- Personalized financial advice and investment recommendations are provided by AI-driven platforms.
- Algorithmic trading uses AI to analyze market data and execute trades in real-time.

4. Transportation:

- Self-driving cars utilize AI for navigation, object recognition, and decision-making in complex environments.
- AI optimizes traffic flow and improves transportation efficiency in ridesharing and public transit systems.
- AI-powered navigation apps like Google Maps provide real-time updates on traffic, road closures, and alternative routes.

5. Customer Service and Virtual Assistants:

- AI-powered chatbots handle customer inquiries, provide support, and resolve issues efficiently.
- Virtual assistants like Siri and Alexa respond to voice commands, manage tasks, and provide information.
- AI learns from interactions to improve customer experiences and tailor responses based on individual preferences.

6. Smart Homes:

- AI-powered devices automate tasks like adjusting temperature, lighting, and security systems.
- Smart home devices learn user preferences and optimize energy consumption for efficiency and convenience.
- AI-driven appliances simplify daily chores and improve the overall living experience.

7. Education:

- Personalized learning platforms use AI to adapt to individual student needs and provide tailored learning experiences.
- AI-powered tools assist with grading, feedback, and administrative tasks, freeing up educators to focus on student interaction.

8. Military and Defense:

- AI assists in decision-making, resource allocation, and strategic planning.
- AI algorithms detect and respond to cyber threats in military networks.
- AI is used in training simulations for military personnel to prepare them for various situations.

9. Agriculture:

- AI-powered robots automate tasks like planting, harvesting, and pest control, optimizing yield and reducing waste.
- AI analyzes data from sensors to monitor crop health and predict optimal growing conditions.

10. Entertainment:

- AI powers recommendation algorithms for movies, music, and other entertainment content.
- AI-generated content, like music and artwork, is also emerging as a new form of entertainment.

A Brief History of AI

The Dawn of Artificial Intelligence (1950s-1960s)

The 1950s , which saw the following advancements , are considered to be the birthplace of AI :

- **1950 :** In 1950 saw the publication of Alan Turing's work ,**"Computing Machinery and Intelligence "** which introduced the Turing Test—a measure of computer intelligence.
- **1956:** A significant turning point in AI research occurs in 1956 when, **John McCarthy** first uses the phrase ***"Artificial Intelligence"*** at the Dartmouth Workshop.
- **1950s–1960s:** The goal of early artificial intelligence (AI) research was to encode human knowledge into computer programs through the use of symbolic reasoning, and logic-based environments.
- **Limited Advancement:** Quick advances are hampered by limited resources and computing-capacity.

- **Early AI systems:** This made an effort to encode human knowledge through the use of logic, and symbolic thinking. The development of early artificial intelligence (AI) systems that, depended on symbolic thinking and logic was hampered by a lack of resources, and processing capacity, which caused the field to advance slowly in the beginning.

AI's Early Achievements and Setbacks (1970s-1980s)

This age has seen notable developments as well as difficulties:

- **1970:** The 1970s witnessed the development of expert systems, which were intended to capture the knowledge of experts in a variety of domains. Data Scientists created rule-based systems that , could use pre-established guidelines to address certain issues.
- **Limitations:** Due to their inability to handle ambiguity and complicated circumstances, these systems had a limited range of applications.
- **The Artificial Intelligence Winter (1970–1980):** A period of inactivity brought on by a lack of funding, and un-met expectations.

Machine Learning and Data-Driven Approaches (1990s)

The 1990s bring a transformative move in AI:

- **1990s:** A worldview move towards machine learning approaches happens.
- **Rise of Machine-Learning:** Calculations learn from information utilizing strategies like neural systems, choice trees, and bolster vector machines.
- **Neural Organize Insurgency:** Propelled by the human brain, neural systems pick up ubiquity for errands like discourse acknowledgment, stock advertise expectation, and motion picture suggestions.
- **Information Powers AI:** Expanded handling control, and information accessibility fuel the development of data driven AI.
- **Unused Areas Rise:** Proposal frameworks, picture acknowledgment and normal dialect handling (NLP) take root.
- **Brilliant Age of AI:** AI frameworks exceed expectations in dis-course acknowledgment, stock determining, and suggestion frameworks.
- **Improved Execution:** Handling control enhancements and information accessibility drive progressions.

The AI Boom: Deep Learning and Neural Networks (2000s-2010s)

The 21st century, witnesses the rise of profound learning, and neural systems :

- **2000s-2010s:** Profound learning a subset of machine learning imitating the human brain's structure and work, came to the cutting edge.

- **Profound Neural Systems:** Multi-layered neural systems exceeded expectations in ranges such as - picture acknowledgment, NLP and gaming.
- **Innovative Progressions:** Profound learning encouraged advance in discourse acknowledgment, NLP, and computer vision.
- **Corporate Speculation:** Tech monsters like ***Facebook, Google, and OpenAI*** made noteworthy commitments to AI inquire about.
- **Counterfeit Neural Systems:** Complex calculations, based on interconnected neurons control profound learning headways.

Generative Pre-trained Transformers: A New Era (GPT Series)

A novel advancement in recent times is the use of Generative Pre-trained Transformers:

- **GPT Series:** Trained on enormous volumes of textual data , these models have rocked the globe.
- **GPT-3:** This model transforms language processing by producing writing that is similar to that of a human being and translating between languages.
- **Learning from Text:** Large volumes of text are absorbed by GPT models, such as - GPT-3, which help them comprehend syntax, context , and comedy.
- **Beyond Translation:** GPT-3 serves as a portable writing assistant by producing essays, poetry , and even language translations.
- **The Upcoming Generation:** This new wave of models , which can write, translate and generate original material as well as provide insightful responses, is exemplified by models such as Bard, ChatGPT, and Bing Copilot.
- **Pushing Boundaries:** These developments have increased the possible applications of AI showcasing its ability in content production, creative projects and language translation.

The State of the Art in AI: Recent advancements

Machine Learning (ML)

Machine Learning is a subset of artificial intelligence (AI) that focuses on building systems that can learn from and make decisions based on data. Instead of being explicitly programmed to perform a task, a machine learning model uses algorithms to identify patterns within data and improve its performance over time without human intervention.

Generative AI

Generative AI refers to a type of artificial intelligence designed to create new content, whether it's text, images, music, or even video. Unlike traditional AI, which typically focuses on analyzing and classifying data, generative AI goes a step further by using patterns it has learned from large datasets to generate new, original outputs. Essentially, it "creates" rather than just "recognizes."

How Generative AI Works

Generative AI works through complex algorithms and deep learning models, often using techniques like neural networks. These networks are trained on vast amounts of data, allowing the AI to understand the underlying structure and patterns within the data.

Reinforcement Learning

Reinforcement Learning (RL) is a branch of machine learning that focuses on how agents can learn to make decisions through trial and error to maximize cumulative rewards. RL allows machines to learn by interacting with an environment and receiving feedback based on their actions. This feedback comes in the form of **rewards or penalties**.

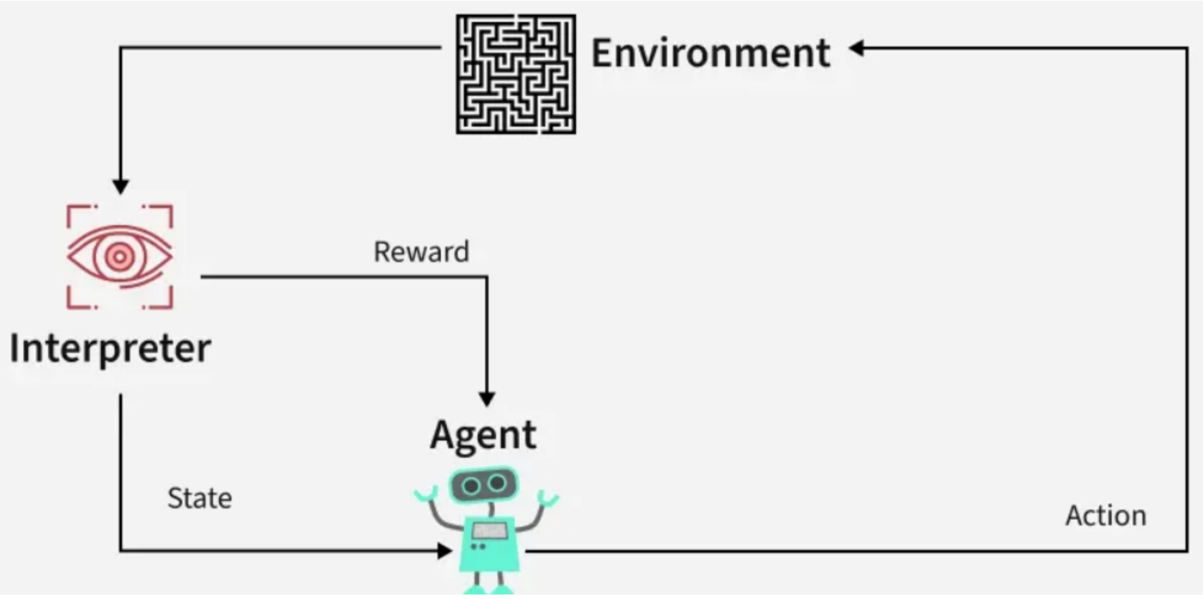

Reinforcement Learning revolves around the idea that an agent (the learner or decision-maker) interacts with an environment to achieve a goal. The agent performs actions and receives feedback to optimize its decision-making over time.

- **Agent**: The decision-maker that performs actions.
- **Environment**: The world or system in which the agent operates.
- **State**: The situation or condition the agent is currently in.
- **Action**: The possible moves or decisions the agent can make.
- **Reward**: The feedback or result from the environment based on the agent’s action.

Modern AI tools and platforms

AI tools are software applications or frameworks that allow users to create, train, deploy, and manage artificial intelligence models. AI platforms provide a full-stack infrastructure (including computing power, storage, and tools) for building AI applications at scale.

OpenAI

OpenAI is known for developing powerful AI models and tools, particularly in the fields of natural language processing, computer vision, and multimodal learning. Its most well-known series of models is GPT, which stands for Generative Pre-trained Transformer. These models are based on the transformer architecture and are trained on vast amounts of internet data using unsupervised learning techniques. The GPT models are designed to predict the next word in a sequence, which enables them to perform a wide range of language tasks such as translation, summarization, question answering, and creative writing. GPT-3 marked a significant leap in scale and capability, while GPT-4 and the multimodal GPT-4o further advanced the technology by enabling real-time understanding and generation of text, images, and audio.

Another significant contribution from OpenAI is Codex, a model trained specifically for programming tasks. It powers GitHub Copilot and helps developers by generating code, suggesting functions, and explaining logic in various programming languages. OpenAI has also introduced DALL·E, a model capable of generating detailed and realistic images from text descriptions, and Whisper, a multilingual automatic speech recognition system. CLIP, another key development, connects vision and language by learning visual concepts through natural language supervision.

OpenAI's approach to training its models involves collecting large, diverse datasets from public and licensed sources, followed by pretraining the models to learn patterns in language or visual data. In addition to standard training, OpenAI uses reinforcement learning from human feedback (RLHF) to fine-tune model behavior. This involves having humans rate and rank model outputs, allowing the AI to better align with human preferences and ethical standards.

TensorFlow

TensorFlow is an open-source machine learning framework developed by the Google Brain team and released in 2015. It is widely used for building and deploying machine learning and deep learning models across various platforms, including desktops, servers, mobile devices, and the web. TensorFlow provides a comprehensive ecosystem of tools, libraries, and community resources that support the development of machine learning applications, from simple experimentation to large-scale production systems. It is written primarily in C++, with user-friendly APIs available in Python, making it accessible to both researchers and developers.

At its core, TensorFlow operates by representing computations as data flow graphs. In these graphs, nodes represent mathematical operations, while the edges represent the multidimensional data arrays, or tensors, that flow between them. This computational graph approach allows for flexible deployment across different hardware environments, including CPUs, GPUs, and TPUs (Tensor Processing Units), which are specialized chips designed by Google for accelerating machine learning tasks. TensorFlow optimizes performance by efficiently managing these computational resources, enabling fast and scalable execution of models.

One of TensorFlow's defining features is its support for automatic differentiation through its component called AutoDiff. This enables the framework to compute gradients automatically, which is essential for training deep neural networks using optimization algorithms like gradient descent. TensorFlow also supports eager execution, a more intuitive programming mode introduced in later versions, where operations are evaluated immediately, rather than building an execution graph beforehand. This allows for easier debugging and interactive model development.

Ethical issues and societal impacts of AI

There are many ethical and societal implications of AI inventions that are worth considering. Some of the most significant ones include:

Bias and discrimination

AI systems can sometimes reflect the biases of the data they are trained on, leading to discriminatory outcomes. There is a risk that AI systems could be biased, either intentionally or unintentionally, in ways that could have negative impacts on certain groups of people. For example, facial recognition systems have been found to be less accurate at identifying people with darker skin tones, which could lead to unequal treatment or disproportionate impact on these groups.

Privacy concerns

AI systems often rely on large amounts of data to function, which raises concerns about data privacy and the potential for this data to be used in ways that individuals do not agree with or that could be harmful to them.

Job displacement

As AI systems become increasingly advanced, there is a risk that they could automate certain jobs, leading to unemployment for some workers. This could have significant societal impacts, including increased inequality and disruption of entire industries.

Autonomy and agency

AI systems may make decisions on their own, leading to questions about accountability and responsibility. If an AI system makes a decision that has negative consequences, it may be difficult to understand why the decision was made and who is responsible.

Security and control

AI systems can be vulnerable to hacking and manipulation, leading to concerns about security and control. As AI systems become more advanced and integrated into society, there is a risk that they could become too powerful and difficult to control. This could lead to unintended consequences and raise questions about who is responsible when things go wrong

It is important for researchers, developers, and policymakers to carefully consider these and other ethical and societal implications as they develop and deploy AI systems.

<table>
<tr><td rowspan="1"></td><td>Pithapur Rajahs Government College(A) Kakinada
Department of Computer Science</td><td rowspan="2">Program: III B.Sc. (AI)
Semester : V</td></tr>
<tr><td rowspan="2">Course 13</td><td>Course Name: ALGORITHMS FOR INTELLIGENT SYSTEMS</td></tr>
<tr><td>Hours Allocated: 3hrs/week</td><td>Credits: 3</td></tr>
</table>

UNIT II

Intelligent Agents: Agents and Environments - What is an agent?**,** Structure of intelligent agents **Rationality:** Rational behavior**,** Bounded rationality, **PEAS Framework:** Performance measure, Environment, Actuators, Sensors**,** Designing task environments.

Environment Types: Fully vs. partially observable, Deterministic vs. stochastic, Episodic vs. sequential, Static vs. dynamic, Discrete vs. continuous

Agent Types: Simple reflex agents, Model-based reflex agents, Goal-based agents, Utility-based agents, Learning agents.

What is an Agent?

In AI, an agent is a software entity that perceives its environment and acts upon it to achieve specific goals. It's essentially a digital assistant or robot that can perform tasks autonomously based on human direction or its own learned capabilities. These agents can range from simple systems with predefined rules to complex ones that utilize machine learning and reasoning to adapt and improve their performance.

Structure of intelligent agents.

So far we have talked about agents by describing behaviour-the action that is performed after any given sequence of precepts. Now, we will have to bite the bullet and talk about how the insides work. The job of A1 is to design the **agent program** that implements the agent function mapping precepts to actions. We assume this program will run on some sort of computing device with physical sensors and actuators-we call this the **architecture.**

agent = architecture +program

Obviously, the program we choose has to be one that is appropriate for the architecture. If the program is going to recommend actions like Walk, the architecture had better have legs. The architecture might be just an ordinary PC, or it might be a robotic car with several onboard computers, cameras, and other sensors. In general, the architecture makes the percepts from the sensors available to the program, runs the program, and feeds the program's action choices to the actuators as they are generated. Most of this book is about designing agent programs, although Chapters 24 and 25 deal directly with the sensors and actuators.

Rationality in AI

Rationality in AI refers to the ability of an artificial agent to make decisions that maximize its performance based on the information it has and the goals it seeks to achieve. In essence, a rational AI system aims to choose the best possible action from a set of alternatives to achieve a specific objective. This involves logical reasoning, learning from experiences, and adapting to new situations.

Types of Rationality

There are two primary types of rationality in AI: bounded rationality and perfect rationality.

1. Bounded Rationality

Bounded rationality recognizes that decision-making capabilities are limited by the information available, cognitive limitations, and time constraints. AI systems operating under bounded rationality use heuristics and approximations to make decisions that are good enough, rather than optimal. This approach is practical in real-world applications where perfect information and infinite computational resources are unavailable.

2. Perfect Rationality

Perfect rationality assumes that an AI system has access to complete information, unlimited computational power, and infinite time to make decisions. While this is an idealized concept, it serves as a benchmark for evaluating the performance of AI systems. Perfectly rational AI would always make the best possible decision in any given situation.

PEAS Framework

In Artificial Intelligence (AI), various types of agents operate to achieve specific goals. The PEAS system is a critical framework used to categorize these agents based on their performance, environment, actuators, and sensors. Understanding the PEAS system is essential for grasping how different AI agents function effectively in diverse environments. Among these agents, Rational Agents are considered the most efficient, consistently choosing the optimal path for maximum efficiency.

PEAS Representation in AI

PEAS is a framework used to specify the structure of an intelligent agent in AI. It breaks down the agent's interaction with the environment into four key components:

Performance Measure: The criteria that define the success of the agent's actions.

Environment: The surroundings or the context in which the agent operates.

Actuators: The mechanisms through which the agent interacts with the environment.

Sensors: The tools the agent uses to perceive its environment.

By defining these elements, PEAS provides a clear outline for designing and evaluating intelligent systems, ensuring they are equipped to perform their tasks effectively.

P: Performance Measure

Performance measure is a quantitative measure that evaluates the outcomes of an agent's actions against a predefined goal. The performance measure is crucial because it guides the agent's decision-making process, ensuring that it acts in a way that maximizes its success.

For example, in a self-driving car, the performance measure could include criteria such as safety (avoiding accidents), efficiency (minimizing travel time), and comfort (ensuring a smooth ride). The car's AI will aim to optimize these factors through its actions.

E: Environment

There are 5 major types of environments:

- Fully Observable & Partially Observable
- Episodic & Sequential
- Static & Dynamic
- Discrete & Continuous
- Deterministic & Stochastic

The environment includes all external factors and conditions that the agent must consider when making decisions. The environment can vary significantly depending on the type of agent and its task.

For instance in the case of a smart thermostat, the environment for a smart thermostat includes the rooms in the house, the outside weather conditions, the heating or cooling system, and the presence of people, all of which the thermostat interacts with to maintain the desired temperature efficiently.

Understanding the environment is critical for designing AI systems because it affects how the agent perceives its surroundings and interacts with them.

A: Actuators

They are responsible for executing the actions decided by the agent based on its perceptions and decisions. In essence, actuators are the "hands and feet" of the agent, enabling it to carry out tasks.

The actuators for a smart thermostat include the heating system, cooling system, and fans, which it controls to adjust the room temperature and maintain the desired comfort level.

The design and choice of actuators are crucial because they directly affect the agent's ability to perform its functions in the environment.

S: Sensors

Sensors collect data from the environment, which is then processed by the agent to make informed decisions. Sensors are the "eyes and ears" of the agent, providing it with the necessary information to act intelligently.

The sensors for a smart thermostat include temperature sensors to measure the current room temperature, humidity sensors to detect moisture levels, and motion sensors to determine if people are present in the house.

The quality and variety of sensors used in an AI system greatly influence its ability to perceive and understand its environment.

Importance of PEAS in AI

The PEAS framework is vital for the design and development of AI system because it provides a structured approach to defining the agent's interaction with its environment. By clearly specifying the performance measure, environment, actuators, and sensors, developers can create AI systems that are more effective and adaptable to their tasks.

Using PEAS helps in:

Defining clear goals: The performance measure ensures that the agent's actions are aligned with the desired outcomes.

Understanding the operational context: Analyzing the environment allows developers to anticipate challenges and design solutions that are robust and effective.

Designing effective interactions: Selecting the right actuators and sensors ensures that the agent can perceive and interact with its environment in a meaningful way.

Exploring Different Types of AI Agents with PEAS Examples

Agent	**Performance Measure**	**Environment**	**Actuator**	**Sensor**
Hospital Management System	Patient's health, Admission process, Payment	Hospital, Doctors, Patients	Prescription, Diagnosis, Scan report	Symptoms, Patient's response
Automated Car Drive	The comfortable trip, Safety, Maximum Distance	Roads, Traffic, Vehicles	Steering wheel, Accelerator, Brake, Mirror	Camera, GPS, Odometer
Subject Tutoring	Maximize scores, Improvement is students	Classroom, Desk, Chair, Board, Staff, Students	Smart displays, Corrections	Eyes, Ears, Notebooks

Agent	Performance Measure	Environment	Actuator	Sensor
Part-picking robot	Percentage of parts in correct bins	Conveyor belt with parts; bins	Jointed arms and hand	Camera, joint angle sensors
Satellite image analysis system	Correct image categorization	Downlink from orbiting satellite	Display categorization of scene	Color pixel arrays

Advantages of PEAS in AI

Structured Design: Provides a clear framework for designing intelligent agents by breaking down their components.

Versatility: Applicable to various AI systems, from simple bots to complex autonomous agents.

Goal-Oriented: Ensures that agents are designed with specific, measurable objectives, leading to better performance.

Systematic Development: Facilitates organized planning and development, making the process more efficient.

Disadvantages of PEAS in AI

Complexity: Can be complex to implement in dynamic environments with many variables.

Over-Simplification: Might oversimplify real-world scenarios, leading to gaps in agent behavior.

Resource-Intensive: Requires significant resources to accurately define and implement each PEAS component.

Limited Adaptability: May struggle to adapt to unexpected changes if not designed with enough flexibility.

Types of Environments in AI

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:

- Fully Observable vs Partially Observable
- Deterministic vs Stochastic
- Static vs Dynamic
- Discrete vs Continuous
- Episodic vs Sequential

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 chess piece 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 time to time.

3. 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.

4. 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.
- **Example:** Consider an example of **Pick and Place robot**, which is used to detect defective parts from the conveyor belts. Here, every time robot(agent) will make the decision on the current part i.e. there is no dependency between current and previous decisions.
- 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.

Types of Agents

1. Simple Reflex Agents

Simple reflex agents act solely based on the **current percept and**, **percept history** (record of past perceptions) is ignored by these agents. Agent function is defined by **condition-action rules**.

A **condition-action rule** maps a state (condition) to an action.

- If the condition is true, the associated action is performed.
- If the condition is false, no action is taken.

Simple reflex agents are effective in environments that are fully observable (where the current percept gives all needed information about the environment). In **partially observable environments**, simple reflex agents may encounter **infinite loops** because they do not consider the history of previous percepts. Infinite loops might be avoided if the agent can randomize its actions, introducing some variability in its behavior.

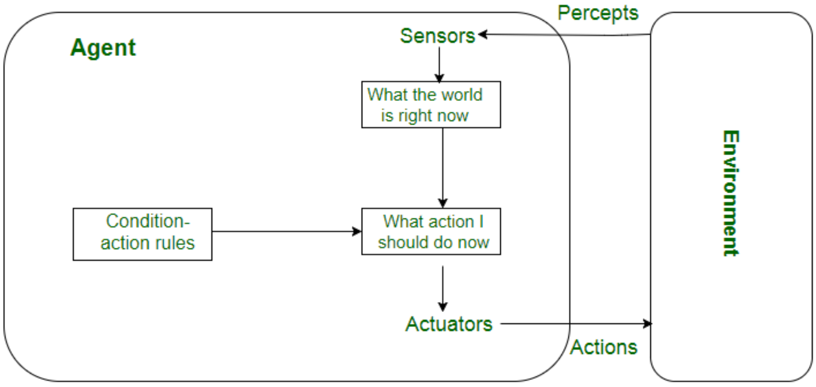

Simple Reflex Agents

2. Model-Based Reflex Agents

Model-based reflex agents finds a rule whose condition matches the current situation or percept. It uses a model of the world to handle situations where the environment is only partially observable.

- The agent tracks its internal state, which is adjusted based on each new percept.
- The internal state depends on the percept history (the history of what the agent has perceived so far).

The agent stores the current state internally, maintaining a structure that represents the parts of the world that cannot be directly seen or perceived. The process of updating the agent's state requires information about:

- How the world evolves independently from the agent?
- How the agent's actions affect the world?

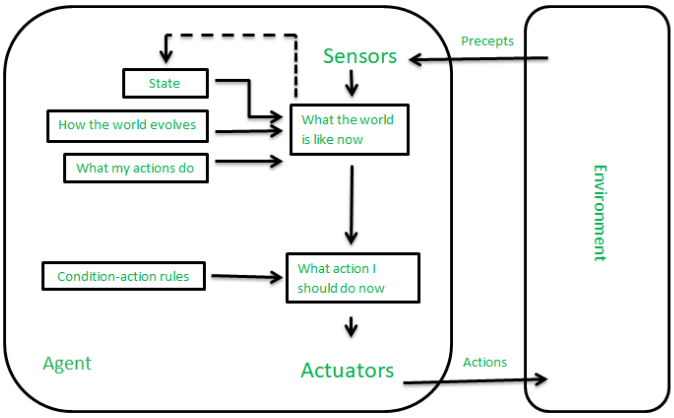

Model-Based Reflex Agents

3. Goal-Based Agents

Goal-based agents make decisions based on their current distance from the goal and every action the agent aims to reduce the distance from goal. They can choose from multiple possibilities, selecting the one that best leads to the goal state.

- Knowledge that supports the agent's decisions is represented explicitly, meaning it's clear and structured. It can also be modified, allowing for adaptability.
- The ability to modify the knowledge makes these agents more flexible in different environments or situations.

Goal-based agents typically require search and planning to determine the best course of action.

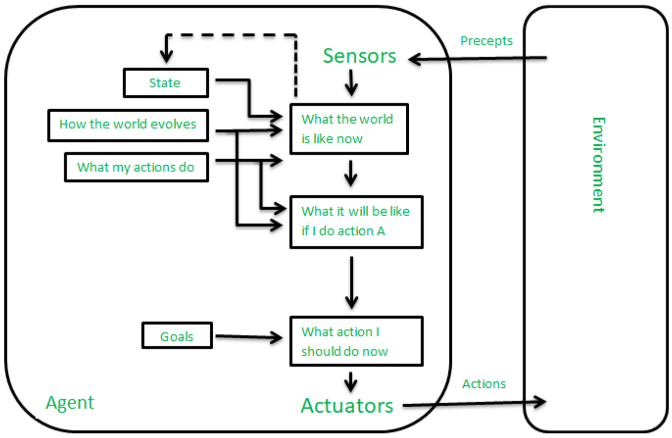

4. Utility-Based Agents

Utility-based agents are designed to make decisions that optimize their performance by evaluating the preferences (or utilities) for each possible state. These agents assess multiple alternatives and choose the one that maximizes their utility, which is a measure of how desirable or "happy" a state is for the agent.

- Achieving the goal is not always sufficient; for example, the agent might prefer a quicker, safer, or cheaper way to reach a destination.
- The utility function is essential for capturing this concept, mapping each state to a real number that reflects the agent's happiness or satisfaction with that state.

Since the world is often uncertain, utility-based agents choose actions that maximize expected utility, ensuring they make the most favorable decision under uncertain conditions.

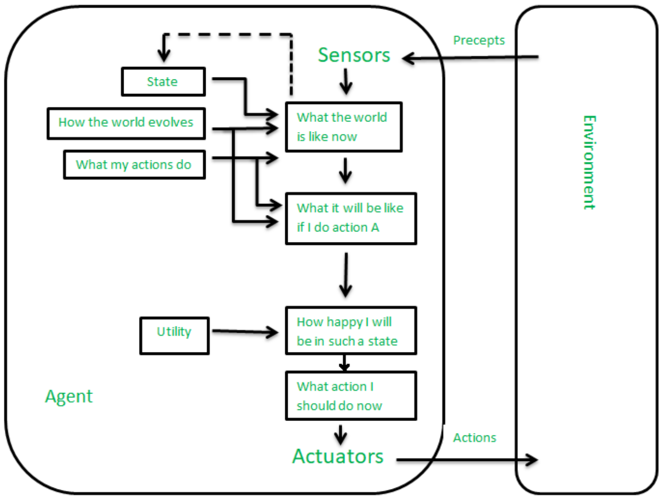

5. Learning Agent

A learning agent in AI is the type of agent that can learn from its past experiences or it has learning capabilities. It starts to act with basic knowledge and then is able to act and adapt automatically through learning. A learning agent has mainly four conceptual components, which are:

1. **Learning element:** It is responsible for making improvements by learning from the environment.
2. **Critic:** The learning element takes feedback from critics which describes how well the agent is doing with respect to a fixed performance standard.
3. **Performance element:** It is responsible for selecting external action.
4. **Problem Generator:** This component is responsible for suggesting actions that will lead to new and informative experiences.