SRI VENKATESWARA UNIVERSITY, TIRUPATI

Department of Computer Science

REVISED CURRICULUM FOR MCA PROGRAMME WITH EFFECT FROM 2025–2026

MCA 204A- Artificial Intelligence

UNIT I

Introduction Concept of AI, history, current status, scope, agents, environments, Problem Formulations, Review of tree and graph structures, State space representation, Search graph and Search tree.

UNIT II

Random search, Search with closed and open list, Depth first and Breadth first search, Heuristic search, best first search, A* algorithm, Game Search.

UNIT III

Probability, conditional probability, Constraint Satisfaction, Propositional Logic & Satisfiability, Uncertainty in AI, Bayes Rule, Bayesian Networks- representation, construction and inference, temporal model, hidden Markov model.

UNIT IV

MDP formulation, utility theory, utility functions, value iteration, policy iteration and partially Observable MDPs.

UNIT V

Passive reinforcement learning, direct utility estimation, adaptive dynamic programming, temporal difference learning, active reinforcement learning- Q learning. Introduction to Machine learning, Deep Learning.

Text Books

1. Stuart Russell and Peter Norvig, “Artificial Intelligence: A Modern Approach” , 3rd Edition, Prentice Hall
2. Elaine Rich and Kevin Knight, “Artificial Intelligence”, Tata McGraw Hill

Reference Books

1. Saroj Kaushik, “Artificial Intelligence”, Cengage Learning India, 2011
2. David Poole and Alan Mackworth, “Artificial Intelligence: Foundations for Computational Agents”,Cambridge University Press 2010.

Prof. M. PADMAVATHAMMA
CHAIRPERSON BOS
COMPUTER SCIENCE (COMBINED)
DEPARTMENT OF COMPUTER SCIENCE
S.V. UNIVERSITY, TIRUPATI-517 502

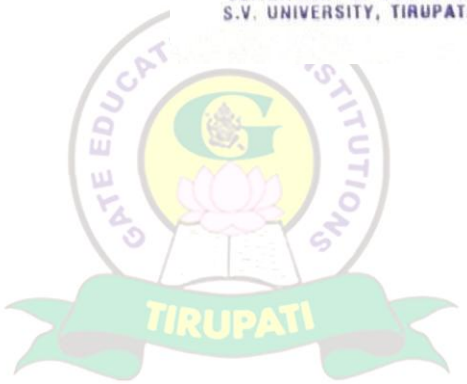

UNIT – I

(Introduction – Concept of AI, History, Current Status, Scope, Agents, Environments, Problem Formulations, Review of Tree and Graph Structures, State Space Representation, Search Graph and Search Tree)

1. Introduction – Concept of Artificial Intelligence

Artificial Intelligence (AI) is a branch of Computer Science that deals with the design and development of intelligent systems capable of performing tasks that normally require human intelligence. These tasks include reasoning, learning, planning, perception, natural language understanding and decision-making.

Artificial Intelligence aims to simulate human cognitive abilities in machines. Traditional computer programs operate based on predefined instructions, whereas AI systems are capable of adapting to new situations, learning from experience and improving performance over time.

AI systems are designed to act rationally. A rational system selects actions that maximize its expected performance measure based on available information and prior knowledge.

Artificial Intelligence combines concepts from multiple disciplines including:

• Computer Science
• Mathematics
• Logic
• Psychology
• Neuroscience
• Engineering

Definitions of Artificial Intelligence

Artificial Intelligence can be defined in several ways:

1. AI is the science and engineering of making intelligent machines.
2. AI is the study of rational agents.
3. AI is the automation of intelligent behavior.
4. AI is the design of systems that can perform tasks requiring human intelligence.

Modern AI primarily follows the rational agent approach, which focuses on building systems that act optimally in a given environment.

Characteristics of Artificial Intelligence

An Artificial Intelligence system typically possesses the following characteristics:

• Ability to learn from experience
• Ability to reason and solve problems
• Ability to handle uncertainty

• Ability to make decisions
• Ability to adapt to new situations
• Ability to understand and process language

Goals of Artificial Intelligence

The main objectives of Artificial Intelligence are:

• To create systems that can learn from data.
• To develop machines that can reason logically.
• To automate complex decision-making tasks.
• To simulate human intelligence in machines.
• To improve efficiency and accuracy in various domains.

Importance of Artificial Intelligence

Artificial Intelligence plays a vital role in modern society. It enhances productivity, reduces human effort and supports intelligent decision-making. AI has transformed various industries including healthcare, finance, education, robotics and transportation.

Conclusion

Artificial Intelligence is a rapidly growing field that focuses on designing intelligent systems capable of reasoning, learning and acting autonomously. It aims to develop rational agents that can solve complex real-world problems efficiently.

2. History of Artificial Intelligence

The history of Artificial Intelligence (AI) dates back to the early development of computing machines. The idea of creating intelligent machines has fascinated scientists and philosophers for centuries. However, Artificial Intelligence as a formal academic discipline began in the mid-20th century.

The development of AI can be divided into several important phases:

Early Foundations (1940s–1950s)

The invention of digital computers during the 1940s laid the foundation for Artificial Intelligence. Researchers began exploring whether machines could perform tasks that require intelligence.

In 1950, Alan Turing proposed a test known as the **Turing Test** to determine whether a machine could exhibit intelligent behavior similar to a human. According to this test, if a human evaluator cannot distinguish between a machine and a human based on conversation, the machine can be considered intelligent.

This idea initiated serious research into machine intelligence.

Dartmouth Conference (1956)

The term “Artificial Intelligence” was officially introduced in 1956 at the Dartmouth Conference organized by John McCarthy. This conference is considered the birth of AI as a scientific field.

Researchers believed that every aspect of learning and intelligence could be described precisely enough to make a machine simulate it. This optimistic view led to significant early research efforts.

Early Research Period (1960s–1970s)

During this period, AI research focused on:

• Problem-solving systems
• Theorem-proving programs
• Symbolic reasoning
• Logic-based approaches

Programs such as the General Problem Solver (GPS) were developed to solve mathematical and logical problems.

However, researchers soon realized that many real-world problems were more complex than expected.

AI Winter (1970s–1980s)

Due to unrealistic expectations and limited computing power, AI research faced criticism and funding cuts. This period is known as the “AI Winter.”

The main reasons for AI Winter were:

• Lack of sufficient computational resources
• Difficulty in handling real-world complexity
• Overpromising results

As a result, research progress slowed down significantly.

Expert Systems Era (1980s)

In the 1980s, AI research revived with the development of Expert Systems. These systems were designed to mimic the decision-making ability of human experts in specific domains.

Expert systems were used in:

• Medical diagnosis
• Financial analysis
• Industrial troubleshooting

They were based on rule-based knowledge representation.

Machine Learning Era (1990s–2000s)

With the availability of more data and improved computing power, AI shifted towards Machine Learning techniques.

Instead of programming machines with explicit rules, researchers developed algorithms that allowed machines to learn patterns from data.

This period saw the development of:

• Decision trees
• Neural networks
• Support Vector Machines
• Statistical learning methods

Modern AI (2010 onwards)

Recent advancements in:

• Deep Learning
• Big Data
• Cloud Computing
• High-performance GPUs have significantly improved AI capabilities.

AI systems now achieve high accuracy in:

• Image recognition
• Speech recognition
• Natural language processing
• Autonomous driving

Modern AI applications include virtual assistants, chatbots, recommendation systems and smart robotics.

Conclusion

The history of Artificial Intelligence shows a journey from early theoretical concepts to powerful real-world applications. Although AI faced challenges during its development, advancements in computing power and data availability have transformed it into one of the most important technological fields today.

3. Current Status of Artificial Intelligence

Artificial Intelligence has emerged as one of the most significant technological advancements of the 21st century. With rapid improvements in computing power, availability of large datasets and advanced algorithms, AI systems have achieved remarkable performance in various domains.

Today, AI is no longer limited to research laboratories. It has become an integral part of daily life and industrial applications.

Technological Advancements

The current status of AI is strongly influenced by the following developments:

• High-performance computing systems
• Graphics Processing Units (GPUs)
• Cloud computing platforms
• Availability of Big Data
• Advanced Machine Learning and Deep Learning algorithms

These advancements have enabled AI systems to process massive amounts of data efficiently and learn complex patterns.

AI in Everyday Applications

Artificial Intelligence is widely used in modern applications such as:

• Virtual Assistants – Voice-based systems that understand and respond to user queries.
• Recommendation Systems – Suggest products, movies or content based on user preferences.
• Facial Recognition Systems – Used in security and authentication.
• Chatbots – Provide automated customer support.
• Autonomous Vehicles – Use AI for navigation and obstacle detection.
• Smart Healthcare Systems – Assist in medical diagnosis and patient monitoring.

These applications demonstrate how AI has become embedded in everyday life.

AI in Industry and Research

Artificial Intelligence is transforming multiple industries:

Healthcare:
AI is used for medical image analysis, disease prediction and personalized treatment planning.

Finance:
AI helps in fraud detection, risk assessment and stock market forecasting.

Manufacturing:
AI supports predictive maintenance and industrial automation.

Agriculture:
AI is used for crop monitoring and yield prediction.

Cyber security:
AI detects suspicious activities and prevents cyber-attacks.

In research, AI is being applied in robotics, natural language processing, computer vision and intelligent decision-making systems.

Strengths of Modern AI

The current generation of AI systems can:

• Recognize speech with high accuracy
• Identify objects in images and videos
• Translate languages automatically
• Analyze large datasets efficiently
• Learn from experience

These capabilities have significantly improved automation and decision-making processes.

Limitations and Challenges

Despite its rapid progress, Artificial Intelligence still faces several challenges:

• Lack of general intelligence (AI systems are task-specific)
• Bias in training data
• Ethical concerns and privacy issues
• Lack of transparency in deep learning models
• High computational requirements

Most current AI systems are examples of Narrow AI, meaning they are designed to perform specific tasks and cannot generalize across domains like humans.

Conclusion

The current status of Artificial Intelligence reflects significant technological advancement and widespread adoption across industries. While AI has achieved impressive success in specific applications, research continues toward developing more advanced and general intelligent systems. AI remains a rapidly evolving field with immense potential for future growth.

4. Scope of Artificial Intelligence

The scope of Artificial Intelligence is vast and continuously expanding. AI is not limited to a single field; instead, it is applied across multiple domains including healthcare, finance, education, robotics, transportation and many more. With advancements in computing power and data availability, the scope of AI continues to grow rapidly.

Artificial Intelligence aims to create intelligent systems capable of solving complex real-world problems efficiently and accurately. The wide applicability of AI demonstrates its importance in modern society.

Scope of AI in Healthcare

Artificial Intelligence plays a crucial role in healthcare by assisting doctors and medical professionals in diagnosis and treatment.

Applications include:

• Medical image analysis (X-rays, MRI scans)
• Disease prediction and early detection
• Drug discovery and research
• Personalized treatment recommendations
• Patient monitoring systems

AI helps improve accuracy and reduce human errors in medical decision-making.

Scope of AI in Finance

In the financial sector, AI is widely used for:

• Fraud detection
• Risk assessment
• Credit scoring
• Stock market prediction
• Automated trading systems

AI systems analyze large volumes of financial data to detect unusual patterns and support decision-making.

Scope of AI in Education

AI enhances the learning experience through:

• Intelligent tutoring systems
• Personalized learning platforms
• Automated grading systems
• Educational chatbots

These systems adapt to individual student needs and improve learning efficiency.

Scope of AI in Robotics

AI enables robots to perform tasks autonomously. It is used in:

• Industrial automation
• Warehouse management
• Service robots
• Space exploration
• Hazardous environment operations

AI-powered robots can make decisions, navigate environments and perform complex tasks.

Scope of AI in Natural Language Processing

AI systems can understand and generate human language. Applications include:

• Language translation
• Speech recognition
• Text summarization
• Chatbots
• Voice assistants

Natural Language Processing allows machines to communicate effectively with humans.

Scope of AI in Computer Vision

Computer Vision enables machines to interpret visual data. Applications include:

• Face recognition
• Object detection
• Surveillance systems
• Medical image analysis
• Autonomous vehicles

AI systems analyze images and videos to extract meaningful information.

Scope of AI in Transportation

Artificial Intelligence is used in:

• Self-driving vehicles
• Traffic management systems
• Route optimization
• Accident prediction systems

AI improves safety and efficiency in transportation systems.

Future Scope of Artificial Intelligence

The future scope of AI includes:

• Development of Artificial General Intelligence (AGI)
• Smart cities
• Advanced healthcare automation
• Human-machine collaboration
• Intelligent decision-support systems

Research is ongoing to develop AI systems capable of reasoning and learning across multiple domains like humans.

Conclusion

The scope of Artificial Intelligence is extensive and continuously evolving. AI has transformed various industries by improving efficiency, accuracy and automation. As technology advances, AI is expected to play an even more significant role in shaping the future of society.

5. Agents

The concept of an agent is fundamental in Artificial Intelligence. Modern AI systems are designed as intelligent agents that interact with their environment to achieve specific goals.

An agent is defined as anything that perceives its environment through sensors and acts upon that environment through actuators.

For example:

- A human agent uses eyes and ears as sensors and hands and legs as actuators.
- A robotic vacuum cleaner uses proximity sensors to detect dirt and wheels to move and clean.
- A software agent uses input data as percepts and produces output actions.

An agent continuously interacts with its environment. It receives percepts (inputs) and produces actions (outputs). The complete history of percepts received by an agent is known as the percept sequence.

Agent Function

The behavior of an agent is described by an agent function.

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

Where:
P* represents the sequence of percepts
A represents the set of possible actions

The agent function maps percept sequences to actions.

In practical implementation:

Agent = Architecture + Program

The architecture refers to the physical device or computational system, and the program implements the logic that defines the agent function.

Rational Agent

A rational agent is one that selects actions that maximize its expected performance measure based on its percept sequence and prior knowledge.

Rationality depends on four main factors:

1. Performance measure
2. Percept sequence
3. Prior knowledge of the environment
4. Available actions

A rational agent does not necessarily know everything but acts optimally based on available information.

Performance Measure

A performance measure evaluates how well an agent performs its task. It is defined by the system designer and depends on the problem.

Examples:

• For a vacuum-cleaning agent, performance may be measured by cleanliness and energy efficiency.
• For a chess-playing agent, performance may be measured by winning the game.

The agent's objective is to maximize its performance measure.

Types of Agents

Artificial Intelligence classifies agents into five main categories:

1. Simple Reflex Agent

A simple reflex agent acts only on the basis of the current percept. It uses condition–action rules (if–then rules).

Example:
If the room is dirty → clean the room.

Limitation:
It works effectively only in fully observable environments.

2. Model-Based Reflex Agent

A model-based agent maintains an internal state to keep track of aspects of the environment that are not directly observable.

It updates its internal state using:

• Previous percepts
• Previous actions

This type of agent is suitable for partially observable environments.

3. Goal-Based Agent

A goal-based agent selects actions that help achieve specific goals.

It requires:

• Goal information
• Search and planning mechanisms

The agent considers future consequences before selecting actions.

4. Utility-Based Agent

A utility-based agent chooses actions based on a utility function that measures how desirable a state is.

Utility allows comparison between different states and helps in decision-making under uncertainty.

The agent selects the action that maximizes expected utility.

5. Learning Agent

A learning agent improves its performance through experience.

A learning agent consists of four components:

1. Performance Element – Selects actions.
2. Learning Element – Improves performance element.
3. Critic – Provides feedback on performance.
4. Problem Generator – Suggests exploratory actions.

Learning agents are widely used in modern AI applications such as recommendation systems and autonomous robots.

Conclusion

Agents form the foundation of Artificial Intelligence systems. The rational agent framework provides a structured approach for designing intelligent systems. Different types of agents are used depending on the nature of the environment and the complexity of the task.

6. Environments

In Artificial Intelligence, an environment refers to everything external to an agent that it interacts with. The environment provides percepts (inputs) to the agent and is affected by the agent's actions.

The nature of the environment plays a crucial role in determining the design of an intelligent agent. Different types of environments require different types of agents and decision-making strategies.

Understanding the characteristics of an environment helps in selecting appropriate AI techniques.

Types of Environments

AI environments are classified based on several properties:

1. Fully Observable vs Partially Observable Environment

Fully Observable Environment

In a fully observable environment, the agent's sensors provide complete information about the current state of the environment.

Example:
A chess game – The agent can see the entire board.

Partially Observable Environment

In a partially observable environment, the agent does not have access to complete information about the environment.

Example:
A self-driving car – It cannot observe all possible obstacles or hidden conditions.

Partially observable environments require agents with memory and internal state (model-based agents).

2. Deterministic vs Stochastic Environment

Deterministic Environment

In a deterministic environment, the next state of the environment is completely determined by the current state and the agent's action.

Example:
A mathematical problem-solving system.

Stochastic Environment

In a stochastic environment, the next state is influenced by randomness or uncertainty.

Example:
Stock market prediction, medical diagnosis.

Stochastic environments require probabilistic reasoning techniques.

3. Episodic vs Sequential Environment

Episodic Environment

In an episodic environment, each decision is independent of previous decisions. The agent's action affects only the current episode.

Example:
Image classification – Each image is classified independently.

Sequential Environment

In a sequential environment, current decisions affect future states.

Example:
Chess game – One move influences future moves.

Sequential environments require planning and long-term reasoning.

4. Static vs Dynamic Environment

Static Environment

In a static environment, the environment does not change while the agent is making a decision.

Example:
Solving a crossword puzzle.

Dynamic Environment

In a dynamic environment, the environment may change while the agent is thinking.

Example:
Self-driving cars in traffic.

Dynamic environments require real-time decision-making.

5. Discrete vs Continuous Environment

Discrete Environment

In a discrete environment, states, percepts and actions are limited and countable.

Example:
Chess – finite number of board positions.

Continuous Environment

In a continuous environment, states and actions vary continuously.

Example:
Robot movement in real-world space.

Continuous environments require advanced mathematical modeling.

Importance of Environment Classification

The classification of environments helps in:

• Selecting appropriate agent type
• Choosing suitable search or learning algorithms
• Designing rational decision-making systems
• Handling uncertainty effectively

Different combinations of these properties define the complexity of the AI problem.

Conclusion

The environment is a critical component in Artificial Intelligence. The characteristics of the environment determine how an agent should perceive, reason and act. Proper understanding of environment types enables the design of efficient and rational AI systems.

7. Problem Formulations

Problem formulation is the process of defining a problem in a structured way so that it can be solved using Artificial Intelligence techniques. In AI, problem-solving agents must clearly define the problem before attempting to find a solution.

A problem is formulated by specifying the initial state, possible actions, transition model, goal state and path cost. Proper problem formulation is essential because it determines how efficiently the problem can be solved.

Components of Problem Formulation

A well-defined problem in AI consists of the following components:

1. Initial State

The initial state describes the starting situation of the agent.

Example:
In the 8-puzzle problem, the initial state is the initial arrangement of tiles.

2. Actions

Actions define the set of possible operations the agent can perform.

Example:
In the 8-puzzle problem, the blank tile can move up, down, left or right.

3. Transition Model

The transition model describes the result of performing an action in a given state.

It defines how one state changes into another state.

Example:
Moving the blank tile to the right changes the configuration of tiles.

4. Goal Test

The goal test determines whether the current state satisfies the goal condition.

Example:
In the 8-puzzle problem, the goal state is the correct arrangement of tiles in order.

5. Path Cost

Path cost measures the total cost of reaching a goal state from the initial state.

Cost may represent:

• Distance
• Time
• Number of moves
• Energy consumption

The objective of the agent is usually to find a solution with minimum path cost.

Problem-Solving Agent

A problem-solving agent follows these steps:

1. Formulate the goal
2. Formulate the problem
3. Search for a solution
4. Execute the solution

The agent explores possible states until it finds a sequence of actions that leads to the goal state.

Example: 8-Puzzle Problem

The 8-puzzle consists of 8 numbered tiles and one blank space arranged in a 3×3 grid.

• Initial State – Random arrangement of tiles
• Actions – Move blank tile
• Goal State – Tiles arranged in order
• Path Cost – Number of moves

This problem is solved using search algorithms such as Breadth-First Search or A* algorithm.

Importance of Problem Formulation

Proper problem formulation helps in:

• Reducing complexity
• Avoiding unnecessary states
• Choosing appropriate search strategies
• Improving efficiency

A poorly defined problem can lead to excessive computation and inefficient solutions.

Conclusion

Problem formulation is a crucial step in Artificial Intelligence. It provides a clear structure for solving problems by defining states, actions, transitions, goals and costs. A well-formulated problem enables the use of efficient search algorithms to find optimal solutions.

8. Review of Tree and Graph Structures

In Artificial Intelligence, many problems are represented using tree and graph structures. These structures help in modeling states and transitions between states during search processes. Understanding trees and graphs is essential for designing search algorithms.

Tree Structure

A tree is a hierarchical data structure consisting of nodes connected by edges. It starts from a special node called the root node and expands into child nodes.

Components of a Tree

• Root Node – The topmost node of the tree.
• Parent Node – A node that has one or more child nodes.
• Child Node – A node derived from another node.
• Leaf Node – A node with no children.
• Depth – The number of edges from the root to a node.
• Height – The maximum depth of the tree.

In AI search problems, each node represents a state, and edges represent actions.

Properties of a Tree

1. A tree has only one path between any two nodes.
2. A tree does not contain cycles.
3. It expands from a single root node.

Example of Tree in AI

In a game-playing problem, each move generates new possible game states. These states are organized in a tree structure known as a game tree.

Similarly, in search problems like the 8-puzzle, possible states are represented as nodes of a search tree.

Graph Structure

A graph is a collection of nodes (vertices) connected by edges. Unlike trees, graphs may contain cycles and multiple paths between nodes.

Graphs are more general than trees.

Components of a Graph

• Vertices (Nodes) – Represent states.
• Edges – Represent connections or transitions.
• Directed Graph – Edges have direction.
• Undirected Graph – Edges do not have direction.

Properties of Graphs

1. A graph may contain cycles.
2. Multiple paths may exist between two nodes.
3. It can be directed or undirected.

Graph Representation in AI

In Artificial Intelligence, many problems are represented as state-space graphs.

Each node represents a state of the problem.
Each edge represents an action that transforms one state into another.

Example:

In route-finding problems, cities are nodes and roads are edges.

Difference Between Tree and Graph

Tree:
• No cycles
• Single path between nodes
• Hierarchical structure

Graph:
• May contain cycles
• Multiple paths possible
• More general structure

Importance in Artificial Intelligence

Tree and graph structures are fundamental in AI because:

• Search algorithms operate on trees and graphs.
• Game-playing strategies use game trees.
• Path-finding problems use graphs.
• State-space representation is based on graph structures.

Understanding these structures is essential before studying search algorithms such as Breadth-First Search and Depth-First Search.

Conclusion

Tree and graph structures provide a systematic way to represent states and transitions in Artificial Intelligence problems. Trees are hierarchical and acyclic, whereas graphs are more general and may contain cycles. These structures form the foundation for search techniques in AI.

9. State Space Representation

State Space Representation is one of the fundamental concepts in Artificial Intelligence. It provides a formal way of describing a problem so that search algorithms can be applied to find a solution.

In AI, a problem is represented as a set of possible states and transitions between those states. The collection of all possible states that can be reached from the initial state is called the state space.

Definition

State Space is defined as the set of all possible states reachable from the initial state by applying valid actions.

Each state represents a unique configuration of the problem at a particular point in time.

Components of State Space Representation

A state space representation consists of:

1. Initial State
2. Set of Actions
3. Transition Function
4. Goal State
5. Path Cost

1. Initial State

The initial state is the starting configuration of the problem.

Example:
In the 8-puzzle problem, the initial arrangement of tiles is the initial state.

2. Set of Actions

Actions define the possible moves that can be performed from a given state.

Example:
In the 8-puzzle problem, the blank tile can move up, down, left or right.

3. Transition Function

The transition function describes the result of performing an action in a particular state.

It defines how one state changes into another state.

4. Goal State

The goal state is the desired final configuration that the agent aims to reach.

The goal test checks whether the current state satisfies the goal condition.

5. Path Cost

Path cost represents the total cost of reaching a state from the initial state.

Cost may represent:

- Distance
- Time
- Number of moves
- Energy

The objective is usually to find a solution with minimum path cost.

Representation as a Graph

State space can be represented using a graph where:

- Nodes represent states
- Edges represent actions

The search process involves exploring this graph to find a path from the initial state to the goal state.

Example: 8-Puzzle State Space

In the 8-puzzle problem:

- Each configuration of tiles is a state.
- Moving the blank tile generates new states.
- The goal is to reach the ordered configuration.

The state space for the 8-puzzle contains thousands of possible states.

Importance of State Space Representation

State space representation is important because:

• It converts real-world problems into formal structures.
• It allows the application of search algorithms.
• It helps in systematic exploration of possible solutions.
• It provides a foundation for heuristic search methods.

Without proper state space representation, solving complex AI problems becomes difficult.

Conclusion

State Space Representation provides a structured framework for modeling AI problems. By defining states, actions and transitions clearly, it becomes possible to apply search algorithms efficiently to reach the goal state. It is a fundamental concept in problem-solving and search techniques in Artificial Intelligence.

10. Search Tree and Search Graph

In Artificial Intelligence, search techniques are used to explore possible states in order to find a path from the initial state to the goal state. During the search process, problems can be represented either as a search tree or as a search graph.

Although both structures represent state exploration, they differ in how states are handled.

Search Tree

A search tree is a tree structure generated during the execution of a search algorithm. It starts from the initial state as the root node and expands by generating child nodes based on possible actions.

Characteristics of Search Tree

• Root node represents the initial state.
• Each child node represents a new state generated by applying an action.
• Nodes may be repeated if the same state is reached through different paths.
• It does not explicitly check for duplicate states.
• No cycles are considered in the tree representation.

Example of Search Tree

In the 8-puzzle problem:

• Initial configuration is the root.
• Each move generates child nodes.
• The tree expands level by level.

If the same configuration appears again through a different path, it is treated as a new node in the search tree.

Advantages of Search Tree

• Simple to implement.
• Easy to understand.

Disadvantages of Search Tree

• May generate repeated states.
• Memory usage increases due to duplication.
• Inefficient for large problems.

Search Graph

A search graph is a more efficient representation that avoids repeated states. It represents the problem as a graph instead of a tree.

In a search graph, each state is represented only once, even if it is reached through multiple paths.

Characteristics of Search Graph

• Nodes represent unique states.
• Edges represent transitions between states.
• Duplicate states are detected and avoided.
• Cycles are properly handled.
• More memory efficient compared to search tree.

Example of Search Graph

In route-finding problems:

• Cities are nodes.
• Roads are edges.
• If a city is visited once, it is not re-expanded unnecessarily.

Search graph representation ensures efficient exploration.

Difference Between Search Tree and Search Graph

Search Tree:

• May contain duplicate states.
• Does not detect cycles.
• Larger memory consumption.
• Simpler implementation.

Search Graph:

• Avoids duplicate states.
• Detects and handles cycles.
• More efficient memory usage.
• Slightly more complex implementation.

Importance in AI

Search Tree and Search Graph representations are essential in AI because:

• Search algorithms like BFS and DFS generate search trees.
• Efficient algorithms use graph search to avoid repetition.
• State-space exploration depends on these structures.
• Heuristic algorithms such as A* often use graph search.

Understanding the difference between search tree and search graph helps in designing efficient problem-solving agents.

Conclusion

Search Tree and Search Graph are two ways of representing the search process in Artificial Intelligence. While search trees may generate duplicate states, search graphs provide a more efficient approach by eliminating repetition and handling cycles properly. Both concepts are fundamental in understanding search algorithms.

UNIT – II

(Random Search, Search with Closed and Open List, Depth First and Breadth First Search, Heuristic Search, Best First Search, A* Algorithm, Game Search).

1. Random Search

Random Search is an uninformed search strategy in which the next state is selected randomly from the available successor states. It does not use any knowledge about the goal state, path cost or heuristic information. Because of this, it belongs to the category of uninformed search algorithms.

In Random Search, there is no systematic order of exploration. The algorithm simply selects one of the possible successor nodes arbitrarily and continues the process until the goal state is found or no more states are available.

Working of Random Search

1. Start from the initial state.
2. Generate all possible successor states.
3. Select one successor randomly.
4. Move to that state.
5. Repeat the process until the goal is found.

The algorithm does not maintain any structured exploration strategy.

Characteristics of Random Search

• It does not use heuristic values.
• It does not guarantee completeness.
• It does not guarantee optimality.
• It may revisit the same state multiple times.
• Performance depends on chance.

Time and Space Complexity

Since exploration is random, time complexity is unpredictable. In worst case, it may explore a large portion of the state space.

Space complexity is low if states are not stored. However, if visited states are stored to avoid repetition, memory usage increases.

Advantages

• Very simple to implement.
• Requires no domain knowledge.
• Useful for comparison purposes.

Disadvantages

• Inefficient for large problems.
• May enter infinite loops.
• Not suitable for structured search problems.

Conclusion

Random Search is the most basic search technique in Artificial Intelligence. Although simple, it is inefficient and impractical for solving real-world AI problems. More advanced search strategies such as BFS, DFS and A* provide better performance.

2. Search with Open List and Closed List

In Artificial Intelligence, graph search algorithms use two important data structures called the **Open List** and the **Closed List**. These lists help manage the exploration of states in a systematic and efficient manner.

The use of Open and Closed Lists distinguishes graph search from simple tree search. It helps in avoiding repeated exploration of the same state and prevents infinite loops in cyclic graphs.

Open List

The Open List contains nodes that have been generated but not yet expanded.

• It represents the frontier of the search.
• It stores states that are waiting to be explored.
• The next node for expansion is selected from the Open List.

The type of data structure used for the Open List depends on the search algorithm:

• In Breadth-First Search, Open List is implemented as a Queue (FIFO).
• In Depth-First Search, Open List is implemented as a Stack (LIFO).
• In Best-First Search and A* Search, Open List is implemented as a Priority Queue.

Closed List

The Closed List contains nodes that have already been expanded.

• It stores visited states.
• It prevents re-expansion of the same node.
• It helps avoid cycles in the state space.

When a node is expanded, it is moved from the Open List to the Closed List.

Working Procedure of Open and Closed Lists

The general graph search procedure using Open and Closed Lists is as follows:

1. Insert the initial state into the Open List.
2. If the Open List is empty, stop (failure).
3. Remove a node from the Open List.
4. If the node is the goal state, stop (success).
5. Otherwise, expand the node to generate successor states.
6. For each successor:
   • If it is not in Open List or Closed List, add it to the Open List.
7. Move the expanded node to the Closed List.
8. Repeat the process.

Importance of Open and Closed Lists

The use of Open and Closed Lists provides several advantages:

• Prevents revisiting the same states.
• Reduces unnecessary computation.
• Avoids infinite loops in cyclic graphs.
• Improves efficiency of search algorithms.
• Enables proper implementation of graph search.

Without a Closed List, the algorithm behaves like a tree search and may repeatedly explore the same state.

Example

Consider a route-finding problem:

• Cities are represented as nodes.
• Roads are represented as edges.

When a city is visited and expanded, it is placed in the Closed List. This prevents the algorithm from revisiting the same city multiple times.

Conclusion

Open and Closed Lists are essential components of graph-based search algorithms in Artificial Intelligence. They ensure systematic exploration of the state space, prevent duplication of states and improve overall efficiency of the search process.

3. Depth First Search (DFS) and Breadth First Search (BFS)

Depth First Search and Breadth First Search are two fundamental uninformed search algorithms used in Artificial Intelligence. Both algorithms systematically explore the state space but differ in their exploration strategies.

Depth First Search (DFS)

Depth First Search is an uninformed search algorithm that explores the deepest node in the search tree first before backtracking.

DFS expands one branch of the tree as far as possible before exploring other branches.

Working of DFS

1. Start with the initial state.
2. Push it onto a Stack (Open List).
3. Remove the top node from the stack.
4. If it is the goal, stop.
5. Otherwise, expand the node.
6. Push its successor nodes onto the stack.
7. Repeat until goal is found or stack becomes empty.

DFS follows the Last-In-First-Out (LIFO) principle.

Characteristics of DFS

• Uses Stack data structure.
• Explores deep into the search tree.
• May get stuck in infinite paths.
• Does not guarantee optimal solution.

Time and Space Complexity of DFS

Let:
b = branching factor
m = maximum depth of tree

Time Complexity = O(b^m)
Space Complexity = O(bm)

Advantages of DFS

• Requires less memory compared to BFS.
• Suitable for deep search problems.
• Simple implementation.

Disadvantages of DFS

• May not find shortest path.
• May get trapped in infinite loops (without cycle checking).
• Not complete in infinite-depth spaces.

Breadth First Search (BFS)

Breadth First Search is an uninformed search algorithm that explores nodes level by level.

BFS expands all nodes at the current depth before moving to the next level.

Working of BFS

1. Start with the initial state.
2. Insert it into a Queue (Open List).
3. Remove the first node from the queue.
4. If it is the goal, stop.
5. Otherwise, expand the node.
6. Add all successor nodes to the end of the queue.
7. Repeat until goal is found or queue becomes empty.

BFS follows the First-In-First-Out (FIFO) principle.

Characteristics of BFS

• Uses Queue data structure.
• Explores nodes level by level.
• Complete (if branching factor is finite).
• Optimal (if all step costs are equal).

Time and Space Complexity of BFS

Let:
b = branching factor
d = depth of shallowest goal

Time Complexity = O(b^d)
Space Complexity = O(b^d)

BFS requires large memory because it stores all nodes at current level.

Difference Between DFS and BFS

Depth First Search:

• Uses Stack
• Explores depth-wise
• Less memory usage
• Not guaranteed optimal

Breadth First Search:

• Uses Queue
• Explores level-wise
• High memory usage
• Guaranteed shortest path (if equal cost)

Conclusion

Depth First Search and Breadth First Search are fundamental uninformed search techniques. DFS is memory efficient but not optimal, while BFS guarantees the shortest solution but requires more memory. The choice of algorithm depends on the nature of the problem and resource constraints.

4. Heuristic Search

Heuristic Search is an informed search strategy that uses additional knowledge about the problem to guide the search process towards the goal more efficiently. Unlike uninformed search methods such as BFS and DFS, heuristic search uses domain-specific information to reduce the number of explored states.

A heuristic is a rule or strategy that helps in making decisions quickly. In Artificial Intelligence, a heuristic function estimates how close a given state is to the goal state.

Heuristic Function

A heuristic function is represented as:

h(n)

Where:

n = current node
h(n) = estimated cost from node n to the goal

The heuristic value helps the algorithm decide which node to explore next.

For a heuristic to be useful, it should:

• Be easy to compute
• Provide a good estimate of distance to goal
• Reduce search space effectively

Informed vs Uninformed Search

Uninformed Search:

• Does not use domain knowledge
• Explores blindly
• Examples: BFS, DFS

Informed Search:

• Uses heuristic function
• Guides search towards goal
• More efficient than uninformed search

Heuristic Search belongs to the informed category.

Example of Heuristic

In a route-finding problem:

• Actual cost = distance traveled so far
• Heuristic value = straight-line distance to destination

The straight-line distance acts as a heuristic to estimate closeness to the goal.

Evaluation Function

Many heuristic search algorithms use an evaluation function:

$f(n) = h(n)$

In some cases, it may combine path cost and heuristic:

$f(n) = g(n) + h(n)$

Where:

$g(n)$ = cost from initial state to node n
$h(n)$ = estimated cost from n to goal

This combination helps in selecting the most promising node.

Properties of a Good Heuristic

A good heuristic should satisfy:

1. Admissibility
   A heuristic is admissible if it never overestimates the true cost to reach the goal.
2. Consistency
   A heuristic is consistent if the estimated cost is always less than or equal to the step cost plus the heuristic value of the next node.

These properties ensure optimal solutions in certain algorithms such as A*.

Advantages of Heuristic Search

• Reduces search space significantly
• Faster than uninformed search
• More suitable for large problems

Disadvantages of Heuristic Search

• Requires domain knowledge
• Poor heuristic may lead to wrong direction
• Designing good heuristic is difficult

Conclusion

Heuristic Search improves efficiency by guiding the search process using domain knowledge. By estimating the distance to the goal, heuristic functions help reduce unnecessary exploration and find solutions faster than uninformed search methods. It forms the foundation for advanced algorithms such as Best First Search and A* Search.

5. Best First Search

Best First Search is an informed search algorithm that expands the most promising node first according to a specified evaluation function. It uses heuristic information to guide the search process toward the goal efficiently.

The main idea behind Best First Search is to always select the node that appears to be closest to the goal based on a heuristic estimate.

Evaluation Function

Best First Search uses an evaluation function:

f(n) = h(n)

Where:

n = current node
h(n) = heuristic estimate of distance from node n to the goal

The algorithm selects the node with the smallest heuristic value.

Working of Best First Search

1. Insert the initial node into the Open List (Priority Queue).
2. If Open List is empty, stop (failure).
3. Remove the node with the lowest heuristic value from the Open List.
4. If the node is the goal state, stop (success).
5. Otherwise, expand the node and generate successor states.
6. Add the successors to the Open List.
7. Move the expanded node to the Closed List.
8. Repeat the process.

The Open List is implemented as a Priority Queue ordered by heuristic value.

Characteristics of Best First Search

• Uses heuristic information.
• Expands the most promising node first.
• Uses Open and Closed Lists.
• May not guarantee optimal solution.

Greedy Best First Search

A common version of Best First Search is Greedy Best First Search.

In this version:

f(n) = h(n)

It ignores the cost already incurred and considers only the estimated distance to the goal.

Greedy approach may find a solution quickly but not necessarily the optimal one.

Time and Space Complexity

Let:

b = branching factor
m = maximum depth

Time Complexity ≈ O(b^m) in worst case
Space Complexity ≈ O(b^m)

It may require significant memory because it stores generated nodes.

Advantages

• Faster than uninformed search in many cases.
• Uses domain knowledge to guide search.
• Reduces unnecessary exploration.

Disadvantages

• Not guaranteed to find optimal solution.
• Performance depends on quality of heuristic.
• May get stuck in local minima.

Example

In a route-finding problem:

• Nodes represent cities.
• h(n) represents straight-line distance to destination.

The algorithm always chooses the city that appears closest to the goal.

Conclusion

Best First Search is an informed search strategy that selects nodes based on heuristic evaluation. It improves efficiency by exploring promising paths first. However, it does not guarantee optimal solutions unless combined with path cost, as in the A* algorithm.

6. A* Algorithm

A* (A-Star) Algorithm is one of the most important and widely used informed search algorithms in Artificial Intelligence. It combines the advantages of Uniform Cost Search and Greedy Best First Search by considering both the path cost and the heuristic estimate.

A* is widely used in path-finding and graph traversal problems because it can find the optimal solution efficiently when a suitable heuristic function is used.

Evaluation Function

A* uses the evaluation function:

$f(n) = g(n) + h(n)$

Where:

n = current node
g(n) = actual cost from initial state to node n
h(n) = estimated cost from node n to goal
f(n) = estimated total cost of solution path through node n

The algorithm selects the node with the smallest f(n) value for expansion.

Working of A* Algorithm

1. Insert the initial node into the Open List (Priority Queue).
2. Set g(start) = 0 and compute f(start).
3. If Open List is empty, stop (failure).
4. Remove the node with the lowest f(n) value.
5. If the node is the goal state, stop (success).
6. Otherwise, expand the node and generate successors.
7. For each successor:
   • Compute g(n)
   • Compute h(n)
   • Compute f(n) = g(n) + h(n)
   • Add it to Open List if not already explored.
8. Move expanded node to Closed List.
9. Repeat until goal is found.

Properties of A* Algorithm

Completeness

A* is complete if the branching factor is finite.

Optimality

A* is optimal if the heuristic function is admissible.

An admissible heuristic never overestimates the true cost to reach the goal.

Admissible and Consistent Heuristics

Admissible Heuristic:

$h(n) \leq$ true cost from n to goal

Consistent Heuristic:

$h(n) \leq cost(n, n') + h(n')$

If heuristic is consistent, A* will not need to re-expand nodes.

Time and Space Complexity

Let:

b = branching factor
d = depth of solution

Time Complexity ≈ O(b^d) in worst case
Space Complexity ≈ O(b^d)

A* requires large memory because it stores all generated nodes.

Advantages

• Finds optimal solution (if heuristic is admissible).
• More efficient than uninformed search.
• Combines path cost and heuristic guidance.

Disadvantages

• High memory consumption.
• Performance depends on quality of heuristic.
• Designing good heuristic is challenging.

Applications of A*

• Route finding (GPS systems)
• Robotics path planning
• Game AI
• Network routing

Conclusion

A* Algorithm is a powerful informed search technique that combines actual cost and heuristic estimate to find optimal solutions efficiently. When a suitable heuristic is used, A* significantly reduces search space and guarantees optimality, making it one of the most widely used algorithms in Artificial Intelligence.

7. Game Search

Game Search is a search strategy used in Artificial Intelligence for decision-making in competitive environments. It is mainly applied in two-player games such as Chess, Tic-Tac-Toe, Checkers and other strategic games.

In game-playing problems, the objective is not only to reach a goal state but also to consider the opponent's moves. Therefore, the search process must account for both maximizing and minimizing strategies.

Game search algorithms are based on adversarial search techniques.

Adversarial Search

Adversarial Search is used when two or more agents compete against each other.

In a typical two-player game:

• One player tries to maximize the score (MAX player).
• The other player tries to minimize the score (MIN player).

The AI agent must choose the best move assuming that the opponent also plays optimally.

Game Tree

Game problems are represented using a Game Tree.

• Root node represents the current game state.
• Child nodes represent possible moves.
• Levels alternate between MAX and MIN players.
• Leaf nodes represent terminal states.

Each leaf node has a utility value that indicates the outcome of the game.

Minimax Algorithm

The Minimax algorithm is the fundamental game search algorithm.

It works under the assumption that:

• MAX player tries to maximize the utility value.
• MIN player tries to minimize the utility value.

Working of Minimax

1. Generate the game tree up to a certain depth.
2. Assign utility values to leaf nodes.
3. Propagate values upward:
   • At MAX level → choose maximum value.
   • At MIN level → choose minimum value.
4. The root node selects the move with the best minimax value.

Properties of Minimax

• Complete if game tree is finite.
• Optimal if opponent plays optimally.
• Time Complexity: O(b^m)

(b = branching factor, m = depth)
• Space Complexity: O(bm)

Minimax can be computationally expensive for large games like chess.

Alpha-Beta Pruning

Alpha-Beta Pruning is an optimization technique for the Minimax algorithm. It reduces the number of nodes evaluated in the game tree.

Alpha represents the best value that MAX can guarantee.
Beta represents the best value that MIN can guarantee.

If at any point:

Alpha ≥ Beta

Further exploration of that branch is stopped (pruned), because it will not affect the final decision.

Advantages of Alpha-Beta Pruning

• Reduces search space significantly.
• Produces the same result as Minimax.
• Improves efficiency without affecting optimality.

Evaluation Function in Game Search

In complex games, it is not feasible to expand the tree fully. Therefore, an evaluation function is used to estimate the desirability of non-terminal states.

Example:

In chess, evaluation may consider:

• Number of pieces
• Board position
• King safety

The evaluation function approximates the utility value.

Applications of Game Search

• Chess-playing programs
• Checkers AI
• Tic-Tac-Toe
• Strategy-based video games

Modern AI game engines combine Minimax, Alpha-Beta Pruning and heuristic evaluation functions.

Conclusion

Game Search is an important application of Artificial Intelligence in competitive environments. Using algorithms such as Minimax and Alpha-Beta Pruning, AI systems can make optimal decisions against intelligent opponents. These techniques form the foundation of modern game-playing AI systems.

UNIT – III

(Probability, Conditional Probability, Constraint Satisfaction , propositional logic and satisfiability, uncertainty in AI, Bayes Theorem, Bayesian Networks-representation, construction and inference, temporal model, Hidden Markov Models)

1. Probability

Probability theory is used in Artificial Intelligence to handle uncertainty. In many real-world problems, complete information about the environment is not available. Therefore, AI systems use probability to represent and reason under uncertainty.

Probability provides a mathematical framework to measure uncertainty and make decisions based on incomplete information.

Basic Concepts of Probability

Random Experiment

A random experiment is an experiment whose outcome cannot be predicted with certainty.

Example:
Tossing a coin, rolling a dice.

Sample Space

The set of all possible outcomes of a random experiment is called the sample space.

Example:
For coin toss:
S = {Head, Tail}

Event

An event is a subset of the sample space.

Example:
Getting Head in a coin toss.

Probability of an Event

If all outcomes are equally likely, the probability of an event A is given by:

P(A) = (Number of favorable outcomes) / (Total number of possible outcomes)

The value of probability lies between 0 and 1.

$0 \leq P(A) \leq 1$

Types of Probability

1. Marginal Probability

Probability of a single event occurring.

Example:
P(A)

2. Joint Probability

Probability of two events occurring together.

$P(A \cap B)$

3. Conditional Probability

Probability of event A occurring given that event B has occurred.

P(A | B)

(Conditional Probability will be discussed in next topic.)

Importance of Probability in AI

Probability is essential in AI because:

- Real-world environments are uncertain.
- Sensors may provide incomplete information.
- Outcomes may not be deterministic.
- Decision-making often involves uncertainty.

Applications include:

- Medical diagnosis
- Spam filtering
- Speech recognition
- Robotics
- Risk assessment

Example in AI

Suppose a medical AI system predicts a disease based on symptoms.

If:

P(Disease) = 0.05
P(Symptom | Disease) = 0.8

The system uses probability theory to estimate the likelihood of disease given symptoms.

Conclusion

Probability provides a systematic way to represent uncertainty in Artificial Intelligence. It enables AI systems to reason and make decisions even when complete information is not available. Probability theory forms the foundation for advanced topics such as Bayesian Networks and Hidden Markov Models.

2. Conditional Probability

Conditional Probability is the probability of an event occurring given that another event has already occurred. In Artificial Intelligence, conditional probability is essential for updating beliefs when new evidence becomes available.

It helps AI systems reason in uncertain environments by adjusting probabilities based on additional information.

Definition

The conditional probability of event A given event B is defined as:

$P(A \mid B) = P(A \cap B) / P(B)$

Where:

$P(A \mid B)$ = Probability of A given B
$P(A \cap B)$ = Joint probability of A and B
$P(B)$ = Probability of B

This formula is valid only if $P(B) \neq 0$.

Interpretation

Conditional probability measures how the likelihood of event A changes when event B is known to have occurred.

In simple terms, it answers the question:

"What is the probability of A, assuming B has already happened?"

Example

Suppose:

P(Disease) = 0.05
P(Symptom | Disease) = 0.8

This means that if a person has the disease, there is an 80% chance of showing the symptom.

Conditional probability is widely used in medical diagnosis systems and expert systems.

Multiplication Rule

From the definition of conditional probability:

$P(A \cap B) = P(A \mid B) \times P(B)$

Similarly:

$P(A \cap B) = P(B \mid A) \times P(A)$

This rule helps compute joint probabilities.

Independence of Events

Two events A and B are independent if:

$P(A \mid B) = P(A)$

This means the occurrence of B does not affect the probability of A.

If events are independent:

$P(A \cap B) = P(A) \times P(B)$

Chain Rule of Probability

For multiple events:

$P(A, B, C) = P(A) \times P(B \mid A) \times P(C \mid A, B)$

This rule is important in Bayesian Networks.

Importance in AI

Conditional probability is used in:

• Bayesian inference
• Medical diagnosis systems

• Spam filtering
• Speech recognition
• Risk prediction
• Machine learning algorithms

It forms the foundation for Bayes Rule and probabilistic graphical models.

Conclusion

Conditional Probability allows AI systems to update their beliefs when new evidence is observed. It provides a mathematical basis for reasoning under uncertainty and is essential for advanced probabilistic models such as Bayesian Networks and Hidden Markov Models.

3. Constraint Satisfaction

Constraint Satisfaction is an important problem-solving framework in Artificial Intelligence. A Constraint Satisfaction Problem (CSP) consists of a set of variables, domains for each variable, and constraints that specify allowable combinations of values.

The objective of a CSP is to find an assignment of values to variables such that all constraints are satisfied.

Definition of Constraint Satisfaction Problem (CSP)

A Constraint Satisfaction Problem is defined by three components:

1. Variables
2. Domains
3. Constraints

1. Variables

Variables represent the unknowns that need to be assigned values.

Example:
In map coloring problem, each region is a variable.

2. Domains

The domain of a variable is the set of possible values that can be assigned to it.

Example:
If three colors are available (Red, Green, Blue), then domain of each region is:

D = {Red, Green, Blue}

3. Constraints

Constraints define the rules that restrict the possible values the variables can take.

Example:
Adjacent regions must not have the same color.

Constraints can be:

• Unary Constraint – Involves a single variable.
• Binary Constraint – Involves two variables.
• Higher-order Constraint – Involves more than two variables.

Representation of CSP

CSPs are usually represented using a Constraint Graph:

• Nodes represent variables.
• Edges represent constraints between variables.

This graphical representation helps visualize relationships among variables.

Example: Map Coloring Problem

Problem:

Color regions of a map such that no two adjacent regions have the same color.

Components:

Variables: Regions
Domain: {Red, Green, Blue}
Constraints: Neighboring regions must have different colors

The goal is to assign colors satisfying all constraints.

Methods for Solving CSP

1. Backtracking Search

A depth-first search approach that assigns values to variables one by one and backtracks when a constraint is violated.

2. Forward Checking

After assigning a value to a variable, forward checking removes inconsistent values from neighboring variables' domains.

3. Constraint Propagation

Techniques like Arc Consistency reduce domain sizes before search.

Arc Consistency ensures that for every value of one variable, there exists a consistent value in the connected variable.

Properties of CSP

• CSP problems are discrete and finite.
• Solutions must satisfy all constraints.
• No objective function is required (unless optimization version).

Applications of CSP

• Map coloring
• Scheduling problems
• Timetable generation
• Sudoku solving
• Resource allocation

Importance in AI

Constraint Satisfaction provides a powerful framework for solving combinatorial problems efficiently. It reduces the search space by enforcing constraints early in the process.

Conclusion

Constraint Satisfaction Problems are widely used in Artificial Intelligence to model problems involving multiple variables and restrictions. By defining variables, domains and constraints clearly, AI systems can efficiently search for valid solutions using techniques like backtracking and constraint propagation.

4. Propositional Logic and Satisfiability

Propositional Logic is a formal system used in Artificial Intelligence to represent knowledge and perform logical reasoning. It is also called propositional calculus. In this logic, statements are represented as propositions that can either be true or false.

Propositional Logic provides a foundation for knowledge representation and reasoning in AI systems.

Propositions

A proposition is a declarative statement that has a definite truth value (either true or false).

Examples:

• “It is raining.”
• “2 + 2 = 4.”

Each proposition is usually represented by symbols such as:

P, Q, R

Logical Connectives

Propositions can be combined using logical connectives.

1. Negation (¬)

Represents NOT.

If P is true, then ¬P is false.

2. Conjunction (∧)

Represents AND.

P ∧ Q is true only if both P and Q are true.

3. Disjunction (∨)

Represents OR.

P ∨ Q is true if at least one of P or Q is true.

4. Implication (→)

Represents IF–THEN.

P → Q is false only when P is true and Q is false.

5. Biconditional (↔)

Represents IF AND ONLY IF.

P ↔ Q is true when both have same truth value.

Truth Table

Truth tables are used to determine the truth value of compound propositions.

Example for conjunction:

P | Q | P ∧ Q
T | T | T
T | F | F
F | T | F
F | F | F

Truth tables help evaluate logical expressions.

Knowledge Representation in Propositional Logic

In AI, knowledge is represented as logical sentences.

Example:

P → Q
P
Therefore Q

This is called Modus Ponens, a valid inference rule.

Satisfiability

A propositional formula is said to be satisfiable if there exists at least one assignment of truth values that makes the formula true.

If no assignment makes it true, it is unsatisfiable.

If it is true for all assignments, it is called a tautology.

Example

Formula:

P ∨ ¬P

This is always true, so it is a tautology.

Formula:

P ∧ ¬P

This is always false, so it is unsatisfiable.

SAT Problem

The Satisfiability (SAT) problem is the problem of determining whether a propositional logic formula is satisfiable.

SAT is an important problem in computer science and is known to be NP-complete.

AI systems use SAT solvers in:

• Planning problems
• Verification systems
• Circuit design
• Knowledge reasoning

Importance in AI

Propositional Logic provides:

• Formal reasoning framework
• Rule-based decision making
• Knowledge representation
• Logical inference

It forms the basis for more advanced logics such as First Order Logic.

Conclusion

Propositional Logic is a fundamental tool in Artificial Intelligence for representing and reasoning about knowledge. The concept of satisfiability helps determine whether logical statements can be true under certain assignments. These concepts are essential for automated reasoning systems.

5. Uncertainty in Artificial Intelligence

In many real-world situations, Artificial Intelligence systems must operate under uncertainty. Uncertainty arises when complete, accurate or reliable information about the environment is not available. Therefore, AI systems must reason and make decisions even when the information is incomplete or ambiguous.

Handling uncertainty is one of the most important challenges in Artificial Intelligence.

Sources of Uncertainty

Uncertainty in AI may arise due to the following reasons:

1. Incomplete Information
   The agent may not have access to all relevant data.
2. Noisy Data
   Sensor readings may contain errors.
3. Ambiguity
   Information may have multiple interpretations.
4. Changing Environment
   The environment may be dynamic and unpredictable.
5. Inherent Randomness
   Some processes are probabilistic by nature.

Need for Handling Uncertainty

In real-world applications such as:

• Medical diagnosis
• Weather forecasting
• Stock market prediction
• Robotics
• Speech recognition

Decisions must be made despite uncertainty.

Classical logic (true/false reasoning) is insufficient in such cases. Therefore, probabilistic reasoning is used.

Approaches to Handle Uncertainty

1. Probability Theory

Probability is the most widely used approach to represent uncertainty. It assigns a numerical value between 0 and 1 to express the likelihood of an event.

2. Bayesian Reasoning

Bayesian reasoning updates beliefs when new evidence is observed.

It is based on Bayes Rule and conditional probability.

3. Fuzzy Logic

Fuzzy logic allows partial truth values between 0 and 1.

Example:

Instead of saying “Temperature is hot” (true/false), fuzzy logic allows degrees such as 0.7 hot.

4. Dempster-Shafer Theory

Represents belief functions without requiring precise probabilities.

Degrees of Belief

In probabilistic reasoning, uncertainty is represented as degrees of belief rather than absolute truth.

Example:

If P(Rain) = 0.8, it means there is 80% belief that it will rain.

Decision Making Under Uncertainty

AI systems often use Expected Utility to make decisions under uncertainty.

Expected Utility = Probability × Utility

The rational agent chooses the action that maximizes expected utility.

Importance in AI

Handling uncertainty is essential because:

• Real-world environments are rarely deterministic.
• Information may be incomplete or noisy.
• Intelligent decision-making requires probabilistic reasoning.

Without uncertainty handling, AI systems cannot function effectively in practical scenarios.

Conclusion

Uncertainty is an unavoidable aspect of real-world Artificial Intelligence applications. By using probability theory, Bayesian reasoning and other mathematical frameworks, AI systems can reason effectively and make rational decisions even in uncertain environments.

6. Bayes Rule

Bayes Rule (or Bayes Theorem) is a fundamental concept in probabilistic reasoning and Artificial Intelligence. It provides a mathematical method for updating the probability of a hypothesis when new evidence is observed.

In AI, Bayes Rule is widely used for reasoning under uncertainty and belief updating.

Bayes Theorem Formula

Bayes Rule is given by:

[
P(H \mid E) = \frac{P(E \mid H) \cdot P(H)}{P(E)}
]

Where:

P(H | E) = Posterior Probability (Probability of hypothesis given evidence)
P(E | H) = Likelihood (Probability of evidence given hypothesis)

P(H) = Prior Probability (Initial belief about hypothesis)
P(E) = Evidence Probability

Interpretation of Terms

1. Prior Probability (P(H))

Represents initial belief about the hypothesis before observing evidence.

Example:
Probability that a person has a disease before any test result.

2. Likelihood (P(E | H))

Probability of observing the evidence assuming the hypothesis is true.

Example:
Probability of a positive test result if the person actually has the disease.

3. Posterior Probability (P(H | E))

Updated probability of the hypothesis after considering evidence.

4. Evidence (P(E))

Overall probability of the evidence occurring.

It can be computed using:

[
P(E) = P(E \mid H)P(H) + P(E \mid \neg H)P(\neg H)
]

Example: Medical Diagnosis

Suppose:

P(Disease) = 0.01
P(Positive Test | Disease) = 0.99
P(Positive Test | No Disease) = 0.05

Using Bayes Rule, we calculate:

P(Disease | Positive Test)

This updates belief about having the disease after observing a positive test result.

Importance in Artificial Intelligence

Bayes Rule is used in:

• Medical diagnosis systems
• Spam filtering
• Fault detection systems
• Machine learning classification
• Decision support systems

It allows AI systems to update beliefs when new data is available.

Bayesian Inference

Bayesian inference is the process of using Bayes Rule repeatedly to update probabilities as more evidence is observed.

It forms the foundation for Bayesian Networks and probabilistic graphical models.

Advantages of Bayes Rule

• Provides systematic belief updating.
• Handles uncertainty effectively.
• Incorporates prior knowledge.

Limitations

• Requires accurate prior probabilities.
• Computation may become complex for large systems.

Conclusion

Bayes Rule is a powerful mathematical tool for reasoning under uncertainty in Artificial Intelligence. It allows systems to update beliefs based on new evidence, making it essential for probabilistic reasoning and intelligent decision-making.

7. Bayesian Networks – Representation, Construction and Inference

A Bayesian Network (also called a Belief Network) is a probabilistic graphical model that represents a set of random variables and their conditional dependencies using a directed acyclic graph (DAG).

Bayesian Networks are widely used in Artificial Intelligence for reasoning under uncertainty.

7.1 Representation of Bayesian Networks

A Bayesian Network consists of two main components:

1. Directed Acyclic Graph (DAG)
2. Conditional Probability Tables (CPT)

1. Directed Acyclic Graph (DAG)

• Each node represents a random variable.
• Directed edges represent probabilistic dependencies.
• The graph does not contain cycles.

If there is an edge from node A to node B, then A is a parent of B, and B depends on A.

Example:

Disease → Fever

This means fever depends on disease.

2. Conditional Probability Table (CPT)

Each node has a Conditional Probability Table that defines the probability of the node given its parent nodes.

Example:

P(Fever | Disease)

If a node has no parents, it has a prior probability.

Joint Probability Representation

A Bayesian Network represents the full joint probability distribution as:

$P(X_1, X_2, ..., X_n) = \prod P(X_i \mid Parents(X_i))$

This factorization reduces complexity significantly compared to full joint distribution tables.

7.2 Construction of Bayesian Networks

Constructing a Bayesian Network involves:

1. Identifying relevant variables.
2. Determining dependencies among variables.
3. Drawing directed edges to represent dependencies.

4. Assigning conditional probability values to each node.

The structure of the network is often determined using domain knowledge.

Steps in Construction

1. Choose variables relevant to the problem.
2. Arrange variables in a logical causal order.
3. Connect variables with directed edges based on dependency.
4. Specify Conditional Probability Tables (CPTs).

7.3 Inference in Bayesian Networks

Inference is the process of computing the probability of certain variables given observed evidence.

Example:

Compute P(Disease | Fever)

Inference methods include:

• Exact Inference (Enumeration, Variable Elimination)
• Approximate Inference (Sampling methods)

Types of Inference

1. Diagnostic Inference
   From effect to cause.
2. Predictive Inference
   From cause to effect.
3. Intercausal Inference
   Between causes of the same effect.

Advantages of Bayesian Networks

• Compact representation of joint probability.
• Handles uncertainty effectively.
• Supports belief updating.
• Suitable for complex real-world problems.

Applications

• Medical diagnosis
• Fault detection
• Risk analysis

• Decision support systems
• Machine learning

Conclusion

Bayesian Networks provide a powerful framework for representing probabilistic relationships among variables. By combining graphical representation with probability theory, they enable efficient reasoning and inference under uncertainty in Artificial Intelligence systems.

8. Temporal Model

A Temporal Model is a probabilistic model used to represent systems that evolve over time. In many real-world Artificial Intelligence applications, the state of a system changes dynamically. Therefore, AI systems must model how variables change with time.

Temporal models extend probabilistic reasoning to sequential and time-dependent processes.

Need for Temporal Models

In static models such as simple Bayesian Networks, variables represent relationships at a single point in time. However, many real-world problems involve sequences of events.

Examples:

• Weather prediction
• Stock market analysis
• Speech recognition
• Robot localization
• Medical monitoring

In such cases, the system state changes over time, and past states influence future states.

Time-Indexed Variables

In temporal models, variables are represented with time indices.

Example:

$X_1, X_2, X_3 \ldots$

Where:

X_t represents the state of variable X at time t.

This representation allows modeling transitions between time steps.

Markov Assumption

Temporal models often rely on the Markov assumption.

The Markov assumption states that:

The future state depends only on the present state and not on the past states.

Mathematically:

$P(X_{t+1} \mid X_t, X_{t-1}, \ldots, X_0) = P(X_{t+1} \mid X_t)$

This simplifies modeling of sequential processes.

First-Order Markov Model

In a first-order Markov model:

The current state depends only on the immediately previous state.

This assumption reduces computational complexity.

Transition Model

A temporal model consists of:

1. Transition Model
   Describes how the state evolves from time t to t+1.
2. Sensor Model
   Describes how observations relate to the hidden state.

Example: Weather Model

Suppose:

State = {Rain, No Rain}

Transition probabilities:

$P(Rain_{t+1} \mid Rain_t)$
$P(Rain_{t+1} \mid No\ Rain_t)$

This defines how weather changes over time.

Applications of Temporal Models

• Speech recognition systems
• Robot tracking systems
• Stock price prediction
• Medical diagnosis over time
• Fault monitoring systems

Temporal models are foundational for Hidden Markov Models (HMM).

Importance in AI

Temporal models help AI systems:

• Predict future states
• Track hidden variables
• Update beliefs over time
• Handle sequential data

They are essential for modeling dynamic environments.

Conclusion

Temporal Models extend probabilistic reasoning to time-dependent systems. By incorporating time-indexed variables and transition probabilities, AI systems can model evolving environments effectively. These models serve as the foundation for advanced sequential models such as Hidden Markov Models.

9. Hidden Markov Model (HMM)

A Hidden Markov Model (HMM) is a statistical model used to represent systems that follow a Markov process with hidden states. It is widely used in Artificial Intelligence for modeling sequential and time-dependent data.

In an HMM, the actual state of the system is not directly observable (hidden), but we can observe outputs that depend on those hidden states.

Basic Idea of HMM

In many real-world systems:

• The true state of the system cannot be directly observed.
• Only indirect observations are available.

Example:

In speech recognition, we cannot directly observe the spoken word's internal phoneme states. Instead, we observe sound signals.

Thus:

Hidden States → Internal states of system
Observations → Visible outputs generated by hidden states

Components of Hidden Markov Model

An HMM consists of the following components:

1. States (S)
2. Observations (O)
3. Transition Probability Matrix (A)
4. Emission Probability Matrix (B)
5. Initial State Probability (π)

1. States (S)

Set of hidden states:

$S = \{S_1, S_2, \ldots, S_n\}$

These states are not directly observable.

2. Observations (O)

Set of possible observable outputs:

$O = \{O_1, O_2, \ldots, O_m\}$

Observations depend on hidden states.

3. Transition Probability (A)

Defines probability of moving from one state to another.

$A = P(S_{t+1} \mid S_t)$

Example:

$P(Rain_{t+1} \mid Rain_t)$

4. Emission Probability (B)

Defines probability of observing a certain output given a hidden state.

$B = P(O_t \mid S_t)$

Example:

P(Umbrella | Rain)

5. Initial Probability (π)

Probability distribution of initial states.

$\pi = P(S_1)$

Three Fundamental Problems of HMM

1. Evaluation Problem

Given model parameters and observation sequence, compute probability of observation sequence.

Solved using: Forward Algorithm.

2. Decoding Problem

Given observations, determine the most likely sequence of hidden states.

Solved using: Viterbi Algorithm.

3. Learning Problem

Given observation sequences, learn model parameters (A, B, π).

Solved using: Baum-Welch Algorithm.

Assumptions of HMM

1. Markov Property
   Current state depends only on previous state.
2. Output Independence
   Observation depends only on current state.

Applications of HMM

• Speech recognition
• Part-of-speech tagging in NLP

- Handwriting recognition
- Bioinformatics
- Weather prediction
- Activity recognition

Importance in Artificial Intelligence

Hidden Markov Models are powerful tools for modeling sequential data where internal states are hidden. They provide a mathematical framework for tracking, prediction and inference in dynamic systems.

HMMs are widely used in Natural Language Processing and speech systems.

Conclusion

Hidden Markov Model is a probabilistic model for sequential data with hidden states. By combining transition probabilities, emission probabilities and initial state distributions, HMM enables effective modeling of time-dependent systems in Artificial Intelligence.

UNIT – IV

(Markov Decision Process (MDP), Utility Theory, utility functions, Value Iteration, Policy Iteration, Partially Observable Markov Decision Process (POMDP))

1. Markov Decision Process (MDP)

A Markov Decision Process (MDP) is a mathematical framework used for modeling decision-making in environments where outcomes are partly random and partly under the control of an agent.

MDPs are widely used in Artificial Intelligence for solving sequential decision-making problems under uncertainty.

Need for MDP

In many AI problems:

• The agent interacts with environment repeatedly.
• Actions influence future states.
• Outcomes may be uncertain.
• Decisions must maximize long-term reward.

MDP provides a formal structure for such problems.

Components of MDP

An MDP is defined by the tuple:

(S, A, P, R, γ)

Where:

S → Set of states
A → Set of actions
P → Transition probability
R → Reward function
γ → Discount factor

1. States (S)

A finite set of states representing all possible situations of the environment.

Example:

In a grid world, each cell is a state.

2. Actions (A)

Set of actions available to the agent.

Example:

Up, Down, Left, Right in grid world.

3. Transition Probability (P)

Defines probability of moving from one state to another after taking an action.

$P(s' \mid s, a)$

Probability of reaching state s' from state s after action a.

4. Reward Function (R)

Specifies immediate reward received after taking action a in state s.

$R(s, a)$

Rewards guide the agent toward desirable behavior.

5. Discount Factor (γ)

$0 \leq \gamma \leq 1$

Discount factor determines importance of future rewards.

If γ is close to 1 → future rewards are important.
If γ is close to 0 → immediate rewards are preferred.

Markov Property

MDP satisfies the Markov property:

The next state depends only on the current state and action, not on previous states.

Mathematically:

$P(S_{t+1} \mid S_t, S_{t-1}, ..., A_t) = P(S_{t+1} \mid S_t, A_t)$

This simplifies decision modeling.

Objective of MDP

The goal of the agent is to find an optimal policy π that maximizes expected cumulative reward.

Policy π(s) → action chosen in state s.

Applications of MDP

• Robotics
• Autonomous vehicles
• Resource allocation
• Game playing
• Reinforcement learning

Conclusion

Markov Decision Process provides a powerful mathematical framework for sequential decision-making under uncertainty. By defining states, actions, transition probabilities and rewards, MDP helps agents determine optimal strategies that maximize long-term rewards.

2. Utility Theory

Utility Theory is a framework used in Artificial Intelligence to model rational decision-making under uncertainty. It provides a way to quantify preferences among different outcomes and helps an agent choose actions that maximize overall satisfaction or benefit.

In uncertain environments, outcomes are not guaranteed. Therefore, instead of simply maximizing reward, agents use utility values to evaluate the desirability of possible outcomes.

Need for Utility Theory

In many real-world problems:

• Outcomes involve uncertainty.
• Multiple possible consequences exist.
• Trade-offs must be considered.
• Decisions involve risk.

Utility theory helps the agent choose the most preferred outcome based on expected utility.

3.Utility Function

A utility function assigns a numerical value to each possible outcome.

U(s) = Utility of state s

Higher utility values indicate more desirable states.

Unlike simple reward functions, utility functions represent preferences and risk attitudes.

Expected Utility

In uncertain environments, the agent cannot predict the exact outcome. Instead, it calculates the Expected Utility.

Expected Utility (EU) = Σ P(s) × U(s)

Where:

P(s) = Probability of outcome s
U(s) = Utility of outcome s

The rational agent chooses the action with the highest expected utility.

Rational Preference

Utility theory assumes that rational agents:

• Have consistent preferences.
• Prefer higher utility over lower utility.
• Choose actions that maximize expected utility.

This is called the Principle of Maximum Expected Utility.

Risk Attitudes

Utility theory also accounts for risk preferences.

1. Risk-Neutral
   Focuses only on expected value.
2. Risk-Averse
   Prefers safer outcomes with lower uncertainty.
3. Risk-Seeking
   Prefers risky outcomes with potentially higher rewards.

Utility functions can be designed to reflect these behaviors.

Example

Suppose an agent has two options:

Option A:
Certain reward of 50.

Option B:
50% chance of reward 100,
50% chance of reward 0.

Expected value of B = 50.

However, depending on risk preference, the agent may prefer Option A or B.

Utility theory helps model such decisions.

Relationship Between Utility and MDP

In Markov Decision Processes:

• Reward represents immediate gain.
• Utility represents long-term value of states.

Utility theory helps determine optimal policies in MDP.

Applications of Utility Theory

• Decision-making systems
• Economic models
• Game theory
• Reinforcement learning
• Autonomous systems

Conclusion

Utility Theory provides a systematic approach for rational decision-making under uncertainty. By maximizing expected utility, AI agents can select actions that best align with their preferences and long-term objectives.

4. Value Iteration

Value Iteration is a dynamic programming algorithm used to compute the optimal policy in a Markov Decision Process (MDP). It is used to determine the optimal value function and the best action to take in each state.

Value Iteration is based on the Bellman Optimality Equation and repeatedly updates the value of each state until convergence.

Objective of Value Iteration

The goal of Value Iteration is to compute the optimal value function V*(s), which gives the maximum expected cumulative reward achievable from state s.

Once the optimal value function is obtained, the optimal policy can be derived.

Bellman Optimality Equation

The Bellman equation for Value Iteration is:

$V(s) = \max_a [R(s, a) + \gamma \Sigma P(s' | s, a) V(s')]$

Where:

V(s) = Value of state s
R(s, a) = Immediate reward
γ = Discount factor
P(s' | s, a) = Transition probability
V(s') = Value of next state

This equation updates the value of each state based on expected future rewards.

Steps in Value Iteration

1. Initialize V(s) = 0 for all states.
2. For each state s, update its value using Bellman equation.
3. Repeat updates for all states.
4. Continue until values converge (change becomes very small).
5. Extract optimal policy from final value function.

Policy Extraction

After convergence, the optimal policy π*(s) is:

$\pi^*(s) = \text{argmax}_a [R(s, a) + \gamma \Sigma P(s' | s, a) V(s')]$

This gives the best action for each state.

Convergence Property

Value Iteration converges to the optimal value function if:

• State space is finite.
• Discount factor $\gamma < 1$.

Advantages

• Guaranteed convergence.
• Computes optimal policy.
• Conceptually simple.

Disadvantages

• Computationally expensive for large state spaces.
• Requires full knowledge of transition probabilities.

Applications

• Robotics navigation
• Game playing
• Reinforcement learning
• Resource allocation problems

Conclusion

Value Iteration is an important algorithm for solving Markov Decision Processes. By repeatedly applying the Bellman equation, it computes the optimal value function and determines the best policy for decision-making under uncertainty.

5. Policy Iteration

Policy Iteration is a dynamic programming algorithm used to compute the optimal policy in a Markov Decision Process (MDP). Unlike Value Iteration, which updates value functions directly, Policy Iteration alternates between evaluating a policy and improving it.

Policy Iteration consists of two main steps:

1. Policy Evaluation
2. Policy Improvement

These two steps are repeated until the policy becomes stable (no further changes occur).

Objective of Policy Iteration

The goal of Policy Iteration is to find the optimal policy π^* that maximizes the expected cumulative reward.

Step 1: Policy Evaluation

In this step, the value function $V\pi(s)$ is computed for a given policy π.

The Bellman equation for policy evaluation is:

$V\pi(s) = R(s, \pi(s)) + \gamma \Sigma P(s' | s, \pi(s)) V\pi(s')$

Where:

$V\pi(s)$ = Value of state s under policy π
$R(s, \pi(s))$ = Immediate reward
γ = Discount factor
$P(s' | s, \pi(s))$ = Transition probability

This equation is applied repeatedly until the value function converges.

Step 2: Policy Improvement

After computing the value function, the policy is updated.

For each state s:

$\pi\_new(s) = \text{argmax}_a [R(s, a) + \gamma \Sigma P(s' | s, a) V\pi(s')]$

If $\pi\_new = \pi$ (no change), the policy is optimal.

Otherwise, replace π with π_new and repeat the process.

Algorithm Outline

1. Initialize a random policy π.
2. Perform policy evaluation.
3. Perform policy improvement.
4. If policy changes, repeat from step 2.
5. If policy does not change, stop.

Difference Between Value Iteration and Policy Iteration

Value Iteration:

• Updates value function directly.
• Extracts policy after convergence.

Policy Iteration:

• Alternates between policy evaluation and improvement.
• Often converges in fewer iterations.

Advantages

• Guaranteed to find optimal policy.
• Often faster convergence than Value Iteration.
• Clear separation between evaluation and improvement.

Disadvantages

• Policy evaluation step may be computationally expensive.
• Requires full knowledge of transition model.

Applications

• Reinforcement learning
• Robotics control
• Decision-making systems
• Game strategies

Conclusion

Policy Iteration is an efficient method for solving Markov Decision Processes. By iteratively evaluating and improving policies, it converges to the optimal decision strategy that maximizes expected long-term reward.

6. Partially Observable Markov Decision Process (POMDP)

A Partially Observable Markov Decision Process (POMDP) is an extension of the Markov Decision Process (MDP) in which the agent does not have complete knowledge of the current state of the environment.

In a standard MDP, the agent fully observes the current state before taking an action. However, in many real-world problems, the state is not directly observable. Instead, the agent receives observations that provide partial information about the true state.

Need for POMDP

In real-world scenarios:

• Sensors may provide incomplete information.
• Observations may be noisy.
• True state may be hidden.

Examples:

• Robot navigation with imperfect sensors
• Medical diagnosis based on symptoms
• Autonomous driving in fog
• Speech recognition

In such situations, POMDP provides a mathematical framework for decision-making under partial observability.

Components of POMDP

A POMDP is defined as:

(S, A, P, R, O, Z, γ)

Where:

S → Set of states
A → Set of actions
P → State transition probability
R → Reward function
O → Set of observations
Z → Observation probability function
γ → Discount factor

1. States (S)

True states of the environment (hidden from the agent).

2. Actions (A)

Actions available to the agent.

3. Transition Probability (P)

P(s' | s, a)

Probability of moving to state s' from state s after action a.

4. Reward Function (R)

Immediate reward obtained after performing action a in state s.

5. Observations (O)

Possible observations received by the agent.

6. Observation Probability (Z)

Z(o | s', a)

Probability of observing o after taking action a and reaching state s'.

Belief State

Since the agent cannot observe the true state directly, it maintains a belief state.

Belief state is a probability distribution over all possible states.

b(s) = Probability that system is in state s.

The agent updates its belief using Bayes Rule after receiving new observations.

Difference Between MDP and POMDP

MDP:

• State fully observable.
• Decision based on current state.

POMDP:

• State partially observable.
• Decision based on belief state.

Objective of POMDP

The goal is to find an optimal policy that maximizes expected cumulative reward over belief states.

The policy maps belief states to actions.

Applications of POMDP

• Robotics navigation
• Autonomous vehicles
• Medical diagnosis
• Dialogue systems
• Surveillance systems

Challenges of POMDP

• Computationally complex.
• Large belief state space.
• Difficult to compute exact optimal solutions.

Approximate methods are often used.

Conclusion

Partially Observable Markov Decision Process extends MDP to handle situations where the true state is not fully observable. By maintaining belief states and updating them using probabilistic reasoning, POMDP enables intelligent decision-making in uncertain and partially observable environments.

UNIT – V

(Passive Reinforcement Learning, Direct Utility Estimation, Adaptive Dynamic Programming, Temporal Difference Learning, Active Reinforcement Learning – Q Learning, Introduction to Machine Learning, Deep Learning)

1. Passive Reinforcement Learning

Passive Reinforcement Learning is a type of reinforcement learning in which the agent follows a fixed policy and tries to evaluate that policy. The agent does not attempt to improve or modify the policy. Instead, it learns how good the current policy is.

In other words, the agent's behavior is fixed, and the goal is to estimate the utility (value) of each state under that policy.

Objective of Passive RL

The objective is to learn the utility function U(s) for each state while following a given policy π.

Unlike active reinforcement learning, the agent does not search for better actions. It simply evaluates the performance of the fixed policy.

Characteristics of Passive RL

• Policy is fixed.
• Agent does not explore new actions.
• Focus is on policy evaluation.
• Learns state utilities.
• Suitable for model evaluation.

Working of Passive RL

1. The agent follows a predefined policy π.
2. It interacts with the environment.
3. It observes rewards and state transitions.
4. It estimates utility of states based on observed experience.

Over time, the agent learns which states are more valuable under the fixed policy.

Methods Used in Passive RL

Passive Reinforcement Learning can be implemented using:

1. Direct Utility Estimation
2. Adaptive Dynamic Programming
3. Temporal Difference Learning

These methods differ in how they estimate state utilities.

Example

Consider a robot that always follows a fixed route in a grid world.

It does not change its route but learns:

• Which states give higher rewards
• Which states are dangerous
• Long-term expected reward from each state

Thus, it evaluates the quality of its fixed strategy.

Applications

• Policy evaluation in reinforcement learning
• Performance analysis of fixed strategies
• Simulation-based evaluation

Conclusion

Passive Reinforcement Learning focuses on evaluating a fixed policy by estimating the utility of states. It does not attempt to improve the policy but provides a foundation for understanding policy performance before moving to active reinforcement learning.

2. Direct Utility Estimation

Direct Utility Estimation is a method used in Passive Reinforcement Learning to estimate the utility (value) of states based on actual experience. It is one of the simplest approaches for policy evaluation.

In this method, the utility of a state is estimated by averaging the total rewards received after visiting that state.

Basic Idea

When the agent follows a fixed policy and starts from a particular state, it eventually receives a sequence of rewards.

The total discounted reward obtained from that state is called the return.

Return (G) = $R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots$

Where:

R = Reward
γ = Discount factor

Direct Utility Estimation calculates the utility of a state as the average of returns observed from that state.

Utility Update Formula

If a state s has been visited n times and returns observed are:

$G_1, G_2, G_3, \ldots, G_n$

Then:

$U(s) = (G_1 + G_2 + \ldots + G_n) / n$

This is simply the sample mean of returns.

Working Procedure

1. Follow the fixed policy π.
2. For each visit to state s, record total return G.
3. Update utility estimate as average of observed returns.
4. Repeat until estimates stabilize.

Characteristics

• Simple to implement.
• Does not require knowledge of transition probabilities.
• Uses actual experience.
• Works only for fixed policy (passive learning).

Advantages

• Easy to understand and implement.
• Model-free approach.
• No need for environment model.

Disadvantages

• Requires many episodes for accurate estimates.
• Slow convergence.
• High variance in estimates.
• Does not use partial information from intermediate states.

Example

Suppose state A is visited 3 times and returns are:

10, 8, 12

Then:

U(A) = (10 + 8 + 12) / 3 = 10

This becomes estimated utility of state A.

Comparison with Other Methods

Direct Utility Estimation:

• Uses complete return after episode ends.
• Does not update values step-by-step.

Other methods like Temporal Difference Learning update utilities during each step.

Conclusion

Direct Utility Estimation is a basic model-free method for estimating state utilities in passive reinforcement learning. By averaging observed returns, it evaluates the effectiveness of a fixed policy. Although simple, it may require many experiences to converge to accurate utility values.

3. Adaptive Dynamic Programming (ADP)

Adaptive Dynamic Programming (ADP) is a method used in Passive Reinforcement Learning to estimate the utility of states by learning a model of the environment and then applying dynamic programming techniques.

Unlike Direct Utility Estimation, which directly averages returns, ADP first estimates the transition probabilities and reward function from experience, and then computes utilities using these estimated models.

Basic Idea of ADP

In Adaptive Dynamic Programming:

1. The agent interacts with the environment.
2. It observes transitions and rewards.
3. It learns an approximate model of the environment.
4. It applies dynamic programming (like Value Iteration) to compute utilities.

Thus, ADP is a model-based learning method.

Steps in Adaptive Dynamic Programming

1. Initialize transition model and reward estimates.
2. Follow fixed policy π.
3. Record transitions (s, a, s').
4. Update estimates of transition probabilities.
5. Update reward estimates.
6. Apply dynamic programming to compute utility values.

Transition Model Estimation

Transition probability is estimated using observed frequencies:

$$P(s' \mid s, a) = \frac{\text{Number of times transition } (s, a \rightarrow s') \text{ occurred}}{\text{Total number of times action a taken in state s}}$$

Reward Estimation

Reward for a state can be estimated as average observed reward:

$R(s)$ = Average reward received in state s

Utility Computation

After estimating model parameters, Bellman equation is used:

$$U(s) = R(s) + \gamma \Sigma P(s' \mid s, \pi(s)) U(s')$$

This updates utilities based on learned model.

Characteristics of ADP

• Model-based method.
• Uses estimated transition probabilities.
• More efficient than direct utility estimation.
• Uses full dynamic programming updates.

Advantages

• Faster convergence compared to direct estimation.
• Uses partial information efficiently.
• More accurate utility computation.

Disadvantages

• Requires storage of transition model.
• More computationally complex.
• Needs sufficient experience to estimate probabilities.

Comparison with Direct Utility Estimation

Direct Utility Estimation:

• Model-free
• Uses sample averages

Adaptive Dynamic Programming:

• Model-based
• Learns transition model
• Uses Bellman updates

Applications

• Robotics
• Control systems
• Resource management
• Simulation-based learning

Conclusion

Adaptive Dynamic Programming is a model-based approach in passive reinforcement learning that learns the environment model and uses dynamic programming techniques to estimate state utilities. It improves efficiency and convergence compared to simple averaging methods.

4. Temporal Difference Learning (TD Learning)

Temporal Difference (TD) Learning is a reinforcement learning method that combines ideas from Direct Utility Estimation and Dynamic Programming. It is used to estimate the utility (value) of states by learning directly from experience without requiring a model of the environment.

TD Learning updates the value of a state based on the difference between successive value estimates. This difference is called the Temporal Difference Error.

Basic Idea of TD Learning

Unlike Direct Utility Estimation, which waits until the end of an episode to compute total return, TD Learning updates utility values step-by-step after each transition.

It uses the current reward and estimated utility of the next state to update the current state's utility.

TD Update Rule

The TD(0) update rule is:

$U(s) \leftarrow U(s) + \alpha [R + \gamma U(s') - U(s)]$

Where:

$U(s)$ = Current estimate of utility of state s
α = Learning rate ($0 < \alpha \leq 1$)
R = Immediate reward
γ = Discount factor
$U(s')$ = Utility of next state

Temporal Difference Error

TD Error = $R + \gamma U(s') - U(s)$

This measures the difference between:

- Observed reward + expected future value
- Current estimate of utility

The utility is adjusted in the direction of reducing this error.

Working Procedure

1. Initialize U(s) arbitrarily.
2. Follow fixed policy π.
3. Observe transition ($s \rightarrow s'$) and reward R.
4. Update U(s) using TD formula.
5. Move to next state and repeat.

Characteristics

- Model-free method.
- Updates values incrementally.
- Does not wait for full episode to end.
- Faster learning compared to direct estimation.

Advantages

- Learns online (step-by-step).
- Requires less memory.
- Efficient and practical.
- Does not need transition probabilities.

Disadvantages

• May converge slowly if learning rate is small.
• Sensitive to learning rate parameter.
• Approximate method.

Comparison with Other Methods

Direct Utility Estimation:

• Uses complete return.
• High variance.

Adaptive Dynamic Programming:

• Model-based.
• Requires transition model.

Temporal Difference Learning:

• Model-free.
• Updates using immediate feedback.

Applications

• Reinforcement learning systems
• Game AI
• Robotics control
• Prediction problems

Conclusion

Temporal Difference Learning is a powerful model-free reinforcement learning method that updates state utilities using the difference between successive predictions. By combining ideas from Monte Carlo methods and dynamic programming, TD Learning provides efficient and practical value estimation.

5. Active Reinforcement Learning – Q-Learning

Active Reinforcement Learning refers to the process in which the agent not only evaluates a fixed policy but also learns the optimal policy by exploring different actions. The agent actively chooses actions to maximize cumulative reward.

Q-Learning is one of the most important model-free algorithms in Active Reinforcement Learning. It allows the agent to learn the optimal action-selection policy without knowing the transition model of the environment.

Difference Between Passive and Active RL

Passive Reinforcement Learning:

• Policy is fixed.
• Agent evaluates policy.
• No action optimization.

Active Reinforcement Learning:

• Agent chooses actions.
• Learns optimal policy.
• Balances exploration and exploitation.

Q-Learning belongs to Active RL.

Basic Idea of Q-Learning

Instead of learning utility of states U(s), Q-Learning learns Q-values.

Q(s, a) represents the quality of taking action a in state s.

The goal is to learn the optimal Q-function:

$Q^*(s, a)$

Once Q-values are learned, optimal policy is:

$\pi^*(s) = \text{argmax}_a\, Q(s, a)$

Q-Learning Update Rule

The Q-value update formula is:

$$Q(s, a) \leftarrow Q(s, a) + \alpha\, [\, R + \gamma\, \max_{a'} Q(s', a') - Q(s, a)\,]$$

Where:

Q(s, a) = Current estimate
α = Learning rate
R = Immediate reward
γ = Discount factor
s' = Next state
$\max_{a'} Q(s', a')$ = Maximum future Q-value

Explanation of Formula

The update is based on:

• Immediate reward (R)
• Best possible future reward
• Current Q estimate

The term:

$R + \gamma \max_{a'} Q(s', a')$

is called the Target Value.

The difference between target and current Q is Temporal Difference Error.

Working Procedure

1. Initialize Q(s, a) arbitrarily.
2. Observe current state s.
3. Choose action a (using ε-greedy strategy).
4. Execute action and observe reward R and next state s'.
5. Update Q(s, a) using formula.
6. Move to state s'.
7. Repeat until convergence.

Exploration vs Exploitation

Q-Learning uses ε-greedy strategy:

• With probability ε → choose random action (exploration).
• With probability 1 − ε → choose best action (exploitation).

This ensures proper learning.

Properties of Q-Learning

• Model-free method.
• Off-policy learning algorithm.
• Converges to optimal policy if learning rate decreases properly.
• Does not require transition probabilities.

Advantages

• Learns optimal policy directly.
• Simple and powerful algorithm.
• Widely used in practice.

Disadvantages

• Slow convergence in large state spaces.
• Requires exploration.
• Memory intensive for large problems.

Applications

• Game playing (Atari games)
• Robotics navigation
• Autonomous driving
• Resource allocation
• Reinforcement learning systems

Conclusion

Q-Learning is a powerful Active Reinforcement Learning algorithm that learns optimal action-value functions through interaction with the environment. By continuously updating Q-values using rewards and future estimates, it enables agents to discover optimal decision-making strategies.

6. Introduction to Machine Learning

Machine Learning (ML) is a branch of Artificial Intelligence that enables computers to learn from data and improve performance without being explicitly programmed.

Instead of writing rules manually, Machine Learning algorithms learn patterns from data and make predictions or decisions automatically.

Definition

Machine Learning is defined as:

"A computer program is said to learn from experience E with respect to some task T and performance measure P if its performance at task T improves with experience E."

Need for Machine Learning

Traditional programming requires explicit instructions for every task. However, in many real-world problems:

• Data is large and complex.
• Patterns are difficult to define manually.
• Systems must adapt over time.

Machine Learning provides automated pattern recognition and prediction.

Types of Machine Learning

1. Supervised Learning

In supervised learning, the model is trained using labeled data.

Input → Output pairs are given.

Examples:

• Classification (Spam detection)
• Regression (House price prediction)

2. Unsupervised Learning

In unsupervised learning, data is unlabeled. The model identifies hidden patterns.

Examples:

• Clustering
• Dimensionality reduction

3. Reinforcement Learning

The agent learns through interaction with environment using rewards and penalties.

Example:

• Game playing
• Robotics

Basic Components of Machine Learning

1. Dataset
2. Features (Input variables)
3. Model
4. Loss Function
5. Optimization Algorithm

Machine Learning Process

1. Data Collection
2. Data Preprocessing
3. Model Selection
4. Training
5. Evaluation
6. Deployment

Applications of Machine Learning

• Image recognition
• Speech recognition
• Medical diagnosis
• Fraud detection
• Recommendation systems
• Natural Language Processing

Difference Between AI and ML

Artificial Intelligence:

• Broader concept
• Focuses on intelligent systems

Machine Learning:

• Subset of AI
• Focuses on learning from data

Advantages of Machine Learning

• Handles large datasets
• Improves automatically with experience
• Detects complex patterns

Limitations

• Requires large amount of data
• Can be computationally expensive
• Risk of overfitting

Conclusion

Machine Learning is a powerful subfield of Artificial Intelligence that enables systems to learn from data and make intelligent decisions. It plays a crucial role in modern AI applications and forms the foundation for advanced techniques such as Deep Learning.

7. Introduction to Deep Learning

Deep Learning is a specialized branch of Machine Learning that uses Artificial Neural Networks with multiple hidden layers to model complex patterns in large datasets. It is inspired by the structure and functioning of the human brain.

Deep Learning has become one of the most powerful techniques in Artificial Intelligence, especially for tasks involving images, speech and natural language.

Definition

Deep Learning is defined as:

"A subset of Machine Learning that uses multi-layered neural networks to learn hierarchical representations of data."

The term “deep” refers to the number of hidden layers in the neural network.

Need for Deep Learning

Traditional Machine Learning methods require manual feature extraction. However, in complex tasks:

• Data is high-dimensional (images, audio, video).
• Feature engineering is difficult.
• Large datasets are available.

Deep Learning automatically learns features directly from raw data.

Artificial Neural Networks (ANN)

Deep Learning is based on Artificial Neural Networks.

An ANN consists of:

1. Input Layer
2. Hidden Layers
3. Output Layer

Each layer contains neurons connected by weights.

Working of Neural Network

1. Input data is fed into input layer.
2. Data passes through hidden layers.
3. Each neuron computes weighted sum of inputs.

4. Activation function is applied.
5. Output is generated.

The network learns by adjusting weights using backpropagation.

Activation Functions

Common activation functions:

• Sigmoid
• ReLU (Rectified Linear Unit)
• Tanh
• Softmax

Activation functions introduce non-linearity into the model.

Types of Deep Learning Models

1. Convolutional Neural Networks (CNN)

Used for image processing and computer vision.

Applications:

• Face recognition
• Object detection

2. Recurrent Neural Networks (RNN)

Used for sequential data.

Applications:

• Speech recognition
• Language translation

3. Deep Neural Networks (DNN)

General multi-layer neural networks.

Advantages of Deep Learning

• Automatically extracts features.
• High accuracy for complex problems.
• Works well with large datasets.

Limitations

• Requires large computational resources.
• Needs large training data.
• Training is time-consuming.

Applications of Deep Learning

• Image recognition
• Speech recognition
• Natural Language Processing
• Autonomous vehicles
• Medical image analysis
• Chatbots

Difference Between Machine Learning and Deep Learning

Machine Learning:

• May require manual feature extraction.
• Works well with smaller datasets.

Deep Learning:

• Automatic feature learning.
• Requires large data and computation.

Conclusion

Deep Learning is an advanced Machine Learning technique that uses multi-layer neural networks to learn complex patterns in data. It has revolutionized fields such as computer vision, speech recognition and natural language processing. Deep Learning represents one of the most significant advancements in modern Artificial Intelligence.

GATE EDUCATIONAL INSTITUTIONS
TIRUPATI