		UNIT-IV			
8	a)	Explain planning as state-space search with suitable examples.	L2	CO1 CO3 CO4	5 M
	b)	Discuss the importance of relevant states in regression planning.	L2	CO1 CO3 CO4	5 M
		OR			
9	a)	Explain the working of backward (regression) state-space search with a neat example.	L3	CO1 CO3 CO4	5 M
	b)	Write a detailed note on multi-agent planning. Explain the challenges and solutions.	L2	CO1 CO3 CO4	5 M
		UNIT-V			
10		Explain the theory of learning in AI. Discuss the PAC (Probably Approximately Correct) learning framework.	L3	CO1 CO2 CO4	10 M
		OR			
11		Write a detailed note on statistical learning methods with suitable examples.	L4	CO1 CO2 CO4	10 M

Code: 23ES1305

II B.Tech - I Semester – Regular / Supplementary Examinations
NOVEMBER 2025

ARTIFICIAL INTELLIGENCE
(ARTIFICIAL INTELLIGENCE & MACHINE LEARNING)

Duration: 3 hours Max. Marks: 70

Note: 1. This question paper contains two Parts A and B.
2. Part-A contains 10 short answer questions. Each Question carries 2 Marks.
3. Part-B contains 5 essay questions with an internal choice from each unit. Each Question carries 10 marks.
4. All parts of Question paper must be answered in one place.

BL – Blooms Level CO – Course Outcome

PART – A

		BL	CO
1.a)	Write two components of a problem-solving agent.	L3	CO1
1.b)	List the differences between strong AI and weak AI.	L1	CO1
1.c)	Give one example of a real-life problem that can be modelled as a search problem.	L2	CO1 CO2 CO4
1.d)	Write one difference between BFS and DFS.	L1	CO1 CO2 CO4
1.e)	State one limitation of propositional logic.	L1	CO1 CO2 CO4

1.f)	Give an example of an inference rule.	L2	CO1 CO2 CO4
1.g)	Define a planning graph.	L1	CO1 CO3 CO4
1.h)	What is the main goal of regression planning?	L1	CO1 CO3 CO4
1.i)	What is overfitting in decision trees?	L1	CO1 CO2 CO4
1.j)	Differentiate between positive and negative reinforcement.	L2	CO1 CO2 CO4

PART – B

			BL	CO	Max. Marks
		UNIT-I			
2	a)	Explain the foundations of AI with suitable examples.	L2	CO1	5 M
	b)	Describe the concept of rationality in intelligent agents. Why is it important?	L2	CO1	5 M
		OR			
3	a)	Discuss the structure of an intelligent agent with a neat diagram.	L2	CO1	5 M
	b)	Explain the role of Agent with vacuum world example.	L2	CO1	5 M

		UNIT-II			
4	a)	Define state space search. Explain its role in solving AI problems.	L3	CO1 CO2 CO4	5 M
	b)	Describe the mini-max algorithm with an example of a two-player game.	L2	CO1 CO2 CO4	5 M
		OR			
5	a)	Discuss uninformed vs. informed search strategies with examples.	L2	CO1 CO2 CO4	5 M
	b)	Explain alpha-beta pruning with a suitable example. How does it improve efficiency?	L3	CO1 CO2 CO4	5 M
		UNIT-III			
6	a)	Describe the architecture of a knowledge-based agent with a neat diagram.	L2	CO1 CO2 CO4	5 M
	b)	Explain the working of a simple inference procedure with an example.	L2	CO1 CO2 CO4	5 M
		OR			
7	a)	Compare propositional logic and first-order logic in detail.	L4	CO1 CO2 CO4	5 M
	b)	Apply resolution inference rule with example problems.	L3	CO1 CO2 CO4	5 M

Code: 23ES1305 PVP 23

II B.Tech. – I Sem- Regular Examinations

NOVEMBER 2025

CSE (AI & ML)

ARTIFICIAL INTELLIGENCE

PART –A

SHORT SCHEME

Q.NO	PART - A	Marks
1.a)	Write two components of a problem-solving agent. Any two components of problem-solving agent – 2M	2
1.b)	List the differences between strong AI and weak AI Any Two differences – 2M	2
1.c)	Give one example of a real-life problem that can be modelled as a search problem. Any one example of real-life problem – 2M	2
1.d)	Write one difference between BFS and DFS? Any one difference – 2M	2
1.e)	State one limitation of propositional logic. Any one limitation of propositional logic – 2M	2
1.f)	Give an example of an inference rule. Any one example of inference rule – 2M	2
1.g)	Define a planning graph? planning graph definition – 2M	2
1.h)	What is the main goal of regression planning? Explanation – 2M	2
1.i)	What is overfitting in decision trees. Explanation – 2M	2
1.j)	Differentiate between positive and negative reinforcement Any one difference – 2M	2

PART –B

SHORT SCHEME

UNIT I			
Q.No		**Questions**	**Max Marks**
2	a)	Explain the foundations of AI with suitable examples. Explanation of any four foundations of AI – 5M	5M
	b)	Describe the concept of rationality in intelligent agents. Why is it important. Explanation– 5M	5M
OR			
3	a)	Discuss the structure of an intelligent agent with a neat diagram. Structure and diagram – 5M	5M
	b)	Explain the role of Agent with vacuum world example. Example Explanation – 5M	5M
UNIT II			
4	a)	Define state space search. Explain its role in solving AI problems Any three points -- 5M	5M
	b)	Describe the mini-max algorithm with an example of a two-player game Explanation – 5M	5M
OR			
5	a)	Discuss uninformed vs. informed search strategies with examples. Any three differences -- 5M	5M
	b)	Explain alpha-beta pruning with a suitable example. How does it improve efficiency. Alpha-beta pruning – 4M Any Example-1M	5M
UNIT-III			
6	a)	Describe the architecture of a knowledge-based Agent with a neat diagram. Architecture – 4M Diagram-1M	5M
	b)	Explain the working of a simple inference procedure with an example	5M

		Explanation– 5M	
OR			
7	a)	Compare propositional logic and first order logic in detail. Any three points– 5M	5M
	b)	Apply resolution inference rule with example problems Explanation – 3M Any relevant Example-2M	5M
UNIT-IV			
8	a)	Explain planning as state-space search with suitable examples. Explanation – 5M	5M
	b)	Discuss the importance of relevant states in regression planning. Any 3 points related to regression – 5M	5M
OR			
9	a)	Explain the working of backward (regression) state-space search with a neat example. Explanation with any example – 5M	5M
	b)	Write a detailed note on multi-agent planning. Explain the challenges and solutions. Multi-agent planning – 4M Challenges and Solutions-1M	5M
UNIT - V			
10	Explain the theory of learning in AI. Discuss the PAC (Probably Approximately Correct) learning framework. Explanation – 10 M		10 M
OR			
11	Write a detailed note on statistical learning methods with suitable examples. Explanation – 10M		10 M

Code: 23ES1305

PVP 23

II B.Tech. – I Sem- Regular Examinations
NOVEMBER 2025
CSE (AI & ML)
ARTIFICIAL INTELLIGENCE

PART – A
DETAILED SCHEME

Q.NO	PART - A	Marks
1.a)	Write two components of a problem-solving agent. Any two components of problem-solving agent from the below– 2M • Goal Formulation • Problem formulation • Search • Execution of the solution	2
1.b)	List the differences between strong AI and weak AI Any Two differences – 2M • Definition – strong AI- similar to human-Exact Match • weak AI- solve narrow problems • Any example – Relevant match	2
1.c)	Give one example of a real-life problem that can be modelled as a search problem. Any one example of real-life problem – 2M • GPS Navigation-Finding route in a city- Any example – Relevant match	2
1.d)	Write one difference between BFS and DFS? Any one difference – 2M • BFS- Explores level by level • DFS- Explores deep along a branch	2
1.e)	State one limitation of propositional logic. Any one limitation of propositional logic – 2M • Can not represent relationships between objects -Relevant match	2
1.f)	Give an example of an inference rule. Any one example of inference rule – 2M	2

	• ∀ x King(x) ∧ Greedy(x) ⇒ Evil(x) Inference • King(John) ∧ Greedy(John) ⇒ Evil(John)- Any Example	
1.g)	Define a planning graph? planning graph definition – 2M • A special data structure called a planning graph can be used to give better heuristic estimates.	2
1.h)	What is the main goal of regression planning? Explanation – 2M • In regression search we start at the goal and apply the actions backward until we find a sequence of steps that reaches the initial state.	2
1.i)	What is overfitting in decision trees. Explanation – 2M • Overfitting happens when a model learns random noise instead of the true pattern-relavant match	2
1.j)	Differentiate between positive and negative reinforcement Any one difference – 2M Positive- Encourage by providing a reward Negative- Encourage by removing a punishment	2

PART – B
DETAILED SCHEME

2. a) Explain the foundations of AI with suitable examples. 5M

Ans: Explanation of any four foundations of AI – 5M

- Philosophy
- Mathematics
- Neuroscience
- Economics

(OR)

- Psychology
- Computer Engineering
- Control theory and cybernetics
- Linguistics

2.b) Describe the concept of rationality in intelligent agents. Why is it important. 5M

Ans: Explanation– 5M

The rationality of an agent is measured by its performance measure that evaluates any given sequence of environment states and action sequences, if sequence leads to desirable state then it performed well.

Consider, for example, the vacuum-cleaner agent we might propose to measure performance by

the amount of dirt cleaned up in a single eight-hour shift.

A rational agent can maximize this performance measure by cleaning up the dirt, then dumping it all on the floor, then cleaning it up again, and so on.

A more suitable performance measure would reward the agent for having a clean floor. For example, one point could be awarded for each clean square at each time step (perhaps with a penalty for electricity consumed and noise generated).

A rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

OR

3. a) Discuss the structure of an intelligent agent with a neat diagram. 5M

Structure and diagram – 5M

Ans: The structure of an intelligent agent is a combination of architecture and agent program. It can be viewed as: Agent = Architecture + Agent program

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

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

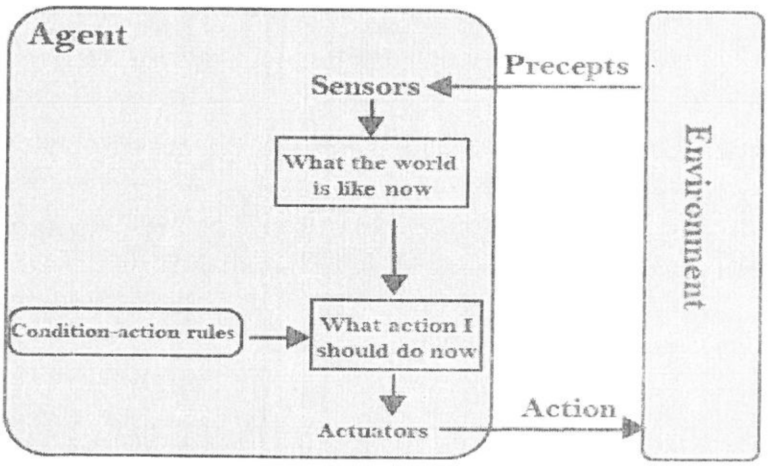

Diagram related to the above one.

3.b) Explain the role of Agent with vacuum world example. 5M

Example Explanation – 5M

Ans: Example—the vacuum-cleaner world shown in Figure below.

One very simple agent function is the following:

Percepts: location/which square it is in and contents, e.g., [A,Dirty]

Actions: Move Left, Right, Suck up dirt, NoOp

Agent's function: look-up table, if the current square is dirty, then suck; otherwise, move to the other square.

Percept sequence	Action
[*A*, *Clean*]	*Right*
[*A*, *Dirty*]	*Suck*
[*B*, *Clean*]	*Left*
[*B*, *Dirty*]	*Suck*
[*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮

UNIT – II

4. a) Define state space search. Explain its role in solving AI problems 5M

Any three points – 5M

Ans:

State space search

• a search tree with the initial state at the root; the branches are actions and the nodes correspond to states in the state space of the problem.

Systematically explores possible solutions (Any three points):

Provides a structured representation
Allows use of efficient search algorithms
Helps in decision-making under constraints
Applicable to many AI tasks

4.b) Describe the mini-max algorithm with an example of a two-player game 5M

Explanation – 5M

Ans:
In artificial intelligence, minimax is a decision-making strategy under game theory, which is used to minimize the losing chances in a game and to maximize the winning chances.

This strategy is also known as 'Minmax,' 'MM,' or 'Saddle point.'

Basically, it is a two-player game strategy where if one wins, the other loose the game. This strategy simulates those games that we play in our day-to-day life.

This strategy simulates those games that we play in our day-to-day life. Like, if two persons are playing chess, the result will be in favor of one player and will unfavor the other one. The person who will make his best try, efforts as well as cleverness, will surely win.

We can easily understand this strategy via game tree- where the nodes represent the states of the game and edges represent the moves made by the players in the game. Players will be two namely:

- MIN: Decrease the chances of MAX to win the game.
- MAX: MAXIMIN his chances of winning the game.

OR

5.a) Discuss uninformed vs. informed search strategies with examples. 5M

Any three differences – 5M

Uninformed Search	Informed Search
1. No heuristic knowledge used	1. Uses heuristic knowledge
2. Blind, unguided search	2. Goal-directed search
3. Explores many unnecessary nodes	3. Explores fewer, relevant nodes
4. Generally slower	4. Generally faster
5. May have lower accuracy	5. Often more accurate and efficient
6. Examples: BFS, DFS, UCS	6. Examples: A*, Greedy Best-First Search
7. Solution quality depends on strategy, not domain knowledge	7. Solution quality improved by good heuristics

5.b) Explain alpha-beta pruning with a suitable example. How does it improve efficiency. 5M

Alpha-beta pruning – 4M

Any Example-1M

Ans:

- Alpha-beta pruning is an advanced version of MINIMAX algorithm. The drawback of minimax strategy is that it explores each node in the tree deeply to provide the best path among all the paths.
- The method used in alpha-beta pruning is that it cutoff the search by exploring less number of nodes. It makes the same moves as a minimax algorithm does, but it prunes the unwanted branches using the pruning technique (discussed in adversarial search).
- Alpha: It is the best highest value, a MAX player can have. It is the lower bound, which represents negative infinity value.
- Beta: It is the best lowest value, a MIN player can have. It is the upper bound which represents positive infinity

Example:

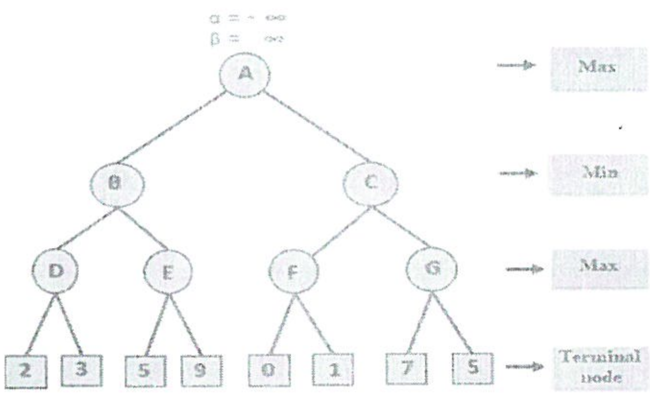

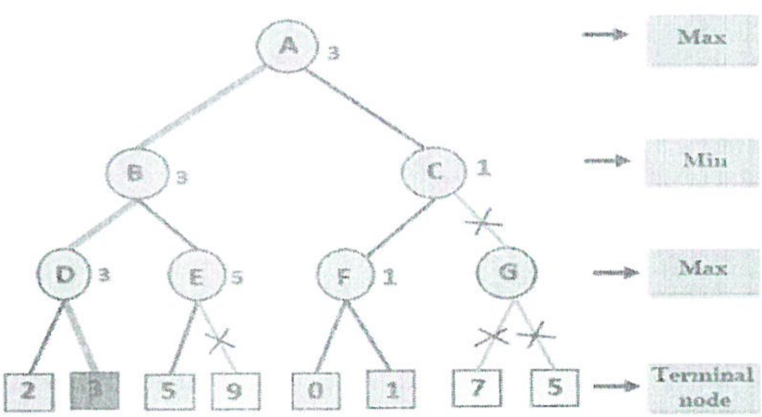

UNIT – III

6.a) Describe the architecture of a knowledge-based Agent with a neat diagram. 5M

Architecture – 4M

Diagram-1M

Ans:

An intelligent agent needs knowledge about the real world for taking decisions and reasoning to act efficiently.

Knowledge-based agents are composed of two main parts:

Knowledge-base and

Inference system.

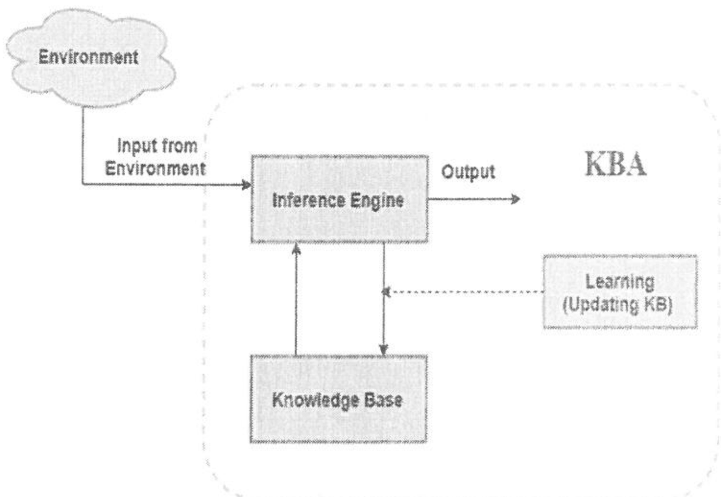

Knowledge base: Knowledge-base is a central component of a knowledge-based agent, it is also known as KB. It is a collection of sentences

Inference system: Inference means deriving new sentences from old. Inference system allows us to add a new sentence to the knowledge base.

6.b) Explain the working of a simple inference procedure with an example. 5M

Explanation– 5M

Ans:

Our first algorithm for inference is a model-checking approach that is a direct implementation of the definition of entailment: enumerate the models, and check that α is true in every model in which KB is true.

$$(\alpha \wedge \beta) \equiv (\beta \wedge \alpha) \quad \text{commutativity of } \wedge$$
$$(\alpha \vee \beta) \equiv (\beta \vee \alpha) \quad \text{commutativity of } \vee$$
$$((\alpha \wedge \beta) \wedge \gamma) \equiv (\alpha \wedge (\beta \wedge \gamma)) \quad \text{associativity of } \wedge$$
$$((\alpha \vee \beta) \vee \gamma) \equiv (\alpha \vee (\beta \vee \gamma)) \quad \text{associativity of } \vee$$
$$\neg(\neg \alpha) \equiv \alpha \quad \text{double-negation elimination}$$
$$(\alpha \Rightarrow \beta) \equiv (\neg \beta \Rightarrow \neg \alpha) \quad \text{contraposition}$$
$$(\alpha \Rightarrow \beta) \equiv (\neg \alpha \vee \beta) \quad \text{implication elimination}$$
$$(\alpha \Leftrightarrow \beta) \equiv ((\alpha \Rightarrow \beta) \wedge (\beta \Rightarrow \alpha)) \quad \text{biconditional elimination}$$
$$\neg(\alpha \wedge \beta) \equiv (\neg \alpha \vee \neg \beta) \quad \text{De Morgan}$$
$$\neg(\alpha \vee \beta) \equiv (\neg \alpha \wedge \neg \beta) \quad \text{De Morgan}$$
$$(\alpha \wedge (\beta \vee \gamma)) \equiv ((\alpha \wedge \beta) \vee (\alpha \wedge \gamma)) \quad \text{distributivity of } \wedge \text{ over } \vee$$
$$(\alpha \vee (\beta \wedge \gamma)) \equiv ((\alpha \vee \beta) \wedge (\alpha \vee \gamma)) \quad \text{distributivity of } \vee \text{ over } \wedge$$

OR

7.a) Compare propositional logic and first order logic in detail. 5M

Any three points– 5M

Propositional Logic	First-Order Logic (FOL)
1. Deals with propositions that are either true or false.	1. Deals with objects, their properties, and relationships.
2. Cannot express internal structure of statements. Each proposition is atomic.	2. Can express complex statements involving variables, relations, and functions.
3. No concepts like quantifiers, variables, or predicates.	3. Uses quantifiers (∀, ∃), variables, predicates, and functions.
4. Limited expressive power.	4. More expressive;
5. Example of statement: "It is raining."	5. Example: ∀x (Human(x) → Mortal(x))

7.b) Apply resolution inference rule with example problems. 5M

Explanation – 3M

Any relevant Example-2M

Ans:

Two clauses, which are assumed to be standardized apart so that they share no variables, can be resolved if they contain complementary literals. Propositional literals are complementary if one is the negation of the other; first-order literals are complementary if one unifies with the negation of the other. Thus, we have

- $C_1: A \vee B$
- $C_2: \neg A \vee C$
- $C_3: \neg C$

Step 1: Negated Goal

We want to prove:

Criminal(West)

So, we start by negating the goal:

¬Criminal(West)

This clause becomes the starting point of the resolution proof.

Step 2: Resolution Spine (Backward Chaining Style)

We now repeatedly resolve with clauses from the knowledge base (KB):

1.¬Criminal(West) resolves wit(¬American(x)∨¬Weapon(y)∨¬Sells(x,y,z)∨¬Hostile(z)∨Criminal(x))

→ gives: ¬American(West) ∨ ¬Weapon(y) ∨ ¬Sells(West, y, z) ∨ ¬Hostile(z).

2. ¬American(West) resolves with

American(West)

→ cancels out, leaving: ¬Weapon(y) ∨ ¬Sells(West, y, z) ∨ ¬Hostile(z).

3. ¬Weapon(y) resolves with

Missile(x)∨Weapon(x)

→ substitution y = M1 gives: ¬Missile(M1) ∨ ¬Sells(West, M1, z) ∨ ¬Hostile(z).

4. Missile(M1) from KB resolves with ¬Missile(M1):

→ result: ¬Sells(West, M1, z) ∨ ¬Hostile(z).

5. ¬Sells(West,M1,z)resolves with

¬Missile(x)∨¬Owns(Nono,x)∨¬Sells(West,x,z)

→ substitution x = M1 gives: ¬Owns(Nono, M1) ∨ ¬Hostile(Nono).

6. Owns(Nono, M1) (from KB) cancels ¬Owns(Nono, M1):

→leaves ¬Hostile(Nono).

7.¬Hostile(Nono) resolves with

¬Enemy(x,America)∨Hostile(x)

→ gives ¬Enemy(Nono, America).

8.Enemy(Nono, America) (from KB) cancels ¬Enemy(Nono, America):

→ Empty clause reached.

Step 3: Conclusion

Since the empty clause is derived, the negated goal ¬Criminal(West) leads to contradiction.

Therefore, Criminal(West) is true.

UNIT – IV

8.a) Explain planning as state-space search with suitable examples. 5M

Explanation – 5M

Ans:

- Forward (Progression) state space search: expand from the initial state toward the goal.
- Backward (Regression) state space search:

In regression search we start at the goal and apply the actions backward until we find a sequence of steps that reaches the initial state. It is called relevant-states search because we only consider actions that are relevant to the goal.

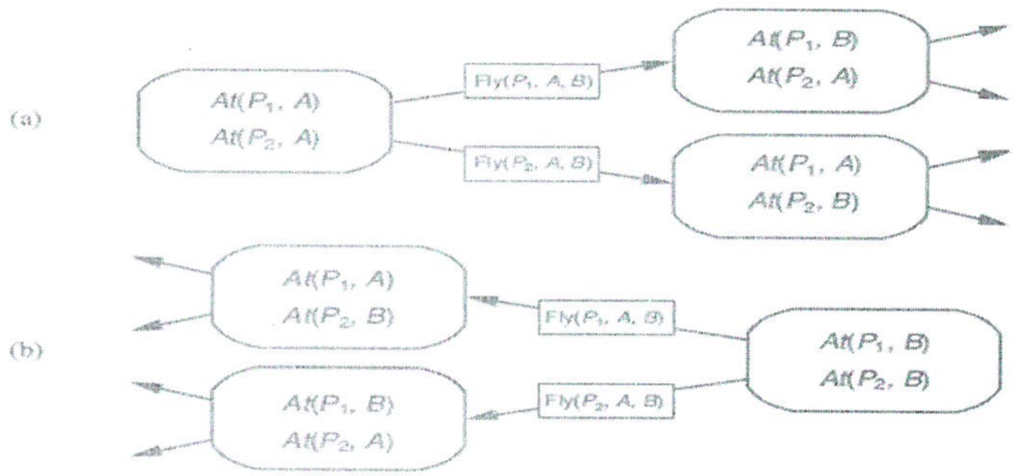

The figure shows forward state transitions (a) and backward regression planning steps (b) for two planes P_1 and P_2 moving between locations A and B using the action Fly(P, A, B).

It illustrates how flying one plane changes its location, while the other plane's location remains the same, demonstrating state updates in progression and goal regression.

8.b) Discuss the importance of relevant states in regression planning 5M

Any 3 points related to regression– 5M

Ans:

Focuses only on states that support achieving the goal: Regression planning starts from the goal and works backward

Reduces the search space significantly:
By including only the states that directly affect the goal conditions, the planner avoids exponential growth of states, making planning much more efficient.

Avoids irrelevant actions:
Only actions that can achieve or maintain the goal conditions are selected, preventing the planner from wasting time evaluating unrelated actions.

Improves planning accuracy:

Relevant states ensure that the planner maintains the needed preconditions and does not violate important goal constraints.

OR

9.a) Explain the working of backward (regression) state-space search with a neat example. 5M

Explanation with any example – 5M

Ans:

Backward or Regression Search starts from the goal state and works backward to find which actions could have produced that goal.

Starting at the goal state.

Finding actions whose effects satisfy the goal.

Determining the preconditions of those actions (what must be true before taking the action).

Repeating the process until reaching the initial state.

Example: Simple Blocks World

Initial State (Init):

On(A, Table) ∧ On(B, Table) ∧ On(C, A) →

- Block A is on the table.
- Block B is on the table.
- Block C is on top of A.

Block(A) ∧ Block(B) ∧ Block(C) → A, B, and C are blocks.

Clear(C) → Nothing is on C (so C can be moved).

Goal (Goal):

On(A, B) ∧ On(B, C) → Stack the blocks so A is on B and B is on C, forming a tower C → B → A.

Actions:

Move(b, x, y): Move block b from x to y.

Preconditions:

- On(b, x) → b is on x.
- Clear(b) → b has nothing on top.
- Clear(y) → y has nothing on top.
- Block(b) ∧ Block(y) → b and y are blocks.
- $b \neq x$, $b \neq y$, $x \neq y$ → cannot move to the same place.

Effects:

- On(b, y) → b is now on y.
- Clear(x) → x becomes clear.
- ¬On(b, x) → b is no longer on x.
- ¬Clear(y) → y is no longer clear.

MoveToTable(b, x): Move block b from x to the table.

Preconditions:

- On(b, x) and Clear(b)
- Block(b)
- $b \neq x$

Effects:

- On(b, Table) → b is on the table.
- Clear(x) → x becomes clear.
- ¬On(b, x) → b is no longer on x.

9.b) Write a detailed note on multi-agent planning. Explain the challenges and solutions. 5M

Multi-agent planning – 4M

Challenges and Solutions-1M

Ans:

When multiple agents exist in an environment, each faces a multiagent planning problem, trying to achieve its own goals—sometimes helping, sometimes hindering others.

Between single-agent and full multiagent planning lies a spectrum:

- **Multieffector planning:** One agent controls several effectors (e.g., typing and speaking simultaneously).
- **Multibody planning:** Effectors are physically separate (e.g., multiple robots) but share information and a common plan.
- **Decentralized planning:** Communication limits prevent full information sharing; each body follows its own subplan with possible communication actions.

Multi-Agent Planning Challenges & Solutions:

Coordination: Actions may clash → use conflict resolution.

Communication: Limited info → distributed planning, protocols.

Partial Observability: Incomplete state → probabilistic/POMDP planning.

Dynamic Environment: Changes occur → replanning/reactive strategies.

Scalability: Many joint actions → hierarchical/abstract planning.

Goal Conflicts: Competing goals → negotiation/shared rewards.

UNIT – V

10. Explain the theory of learning in AI. Discuss the PAC (Probably Approximately Correct) learning framework. 10M

Explanation– 10 M

Ans: Any learning algorithm that returns hypotheses that are probably approximately correct is called a PAC (Probability Approximately Correct) learning algorithm; we can use this approach to provide bounds on the performance of various learning algorithms

- In PAC learning, we aim to show that with enough training examples, any hypothesis consistent with the data is likely to be approximately correct.
- A decision list consists of a series of tests, each of which is a conjunction of literals.
- If a test succeeds when applied to an example description, the decision list specifies the value to be returned.

OR

11. Write a detailed note on statistical learning methods with suitable examples. 10M

Explanation – 10M

Ans:

Statistical Learning Methods: Instantiations of some or all of the random variables describing the domain. The candy is sold in very large bags, of which there are known to be five kinds—again, indistinguishable from the outside:

h_1: 100% cherry,
h_2: 75% cherry + 25% lime,
h_3: 50% cherry + 50% lime,
h_4: 25% cherry + 75% lime,
h_5: 100% lime .

The basic task faced by the agent is to predict the flavor of the next piece of candy. Despite its apparent triviality, this scenario serves to introduce many of the major issues