Indian Institute of Technology (Indian School of Mines) Dhanbad
Department of Computer Science and Engineering
Principles of Artificial Intelligence (CSO303)
Mid Semester Exam. Date: September 15, 2024

Timing: 01:30 PM to 03:30 PM   Monsoon Semester (2023-2024)   Maximum marks: 60

1. (a) Compare atomic, factored, and structured representations. Provide an example for each.

   (b) Agents operate in various types of environments. Explain the differences between a deterministic and a stochastic environment. How does the nature of the environment influence the design of an intelligent agent?

   (c) A financial trading agent operates in stock markets and makes decisions on buying and selling stocks autonomously. Define the PEAS description for this agent. What metrics would indicate that the trading agent is performing well (e.g., profit maximization, risk minimization)? What aspects of the financial environment does the agent need to consider (e.g., market fluctuations, economic news)? What actions can the agent perform to execute trades or manage portfolios? What data sources and sensors would the agent rely on to make decisions (e.g., stock prices, economic indicators)?

(3+3+4)

2. (a) What is a problem-solving agent, and how does it differ from other types of agents? Define the terms state space, initial state, and goal state in the context of the search problem.

   (b) Differentiate between well-defined and ill defined problems. Provide an example of each.

   (c) How does abstraction help in solving complex problems? Provide an example where abstraction can be applied to simplify a real-world problem.

(5+3+2)

3. Suppose you are given the following search space:

State	Next	Cost
P	Q	5
P	R	2
Q	S	6
Q	T	7
R	U	3
R	V	4
S	G	4
T	G	3
U	W	5
V	G	7
W	G	6

   (a) Draw the state space for the search problem, with P as the initial state and G as the goal state. Clearly indicate all paths and their respective costs.

   (b) Assume the initial state is P and the goal state is G. Using the search space from part (a), show how each of the following search strategies would expand the search tree to find a solution:

- Breadth-first search (BFS)
- Depth-first search (DFS)
- Uniform-cost search (UCS)
- Iterative deepening search (IDS)

At each step, show which node is being expanded and maintain a fringe (frontier) table. Also, indicate the final solution found by each algorithm and the cost associated with the path. Compare the results of the algorithms based on time complexity, space complexity, and optimality in the context of this search space. Why do some algorithms find a solution more efficiently than others?

(2+8)

4. (a) Define the time complexity of BFS in terms of the branching factor **b** and the depth **d** of the shallowest solution. How does the branching factor impact the performance of BFS?

   (b) Discuss how Uniform Cost Search (UCS) ensures finding the optimal solution in weighted graphs, particularly focusing on the importance of non-negative edge costs. Additionally, explain how the behavior of UCS would change if the graph contained negative edge costs.

   (c) Why is random-restart more effective in overcoming the limitations of basic hill climbing as opposed to stochastic hill climbing? Under what conditions would random-restart hill climbing be less effective? Provide an example scenario.

(2+3+5)

5. (a) Explain what a heuristic function is in the context of informed search strategies. Discuss the characteristics that make a heuristic admissible and consistent.

   (b) For the following scenario, design two possible heuristics (h1 and h2) for a robot navigating a city grid where it can only move in four directions (up, down, left, right). State which of your heuristics is admissible and consistent, and which one dominates the other with reasons.

   (c) In heuristic-based search algorithms like A*, the choice of the heuristic function directly affects performance. Provide examples of a well-designed heuristic and a poorly designed heuristic for the traveling salesman problem (TSP). What are the specific characteristics of each heuristic that lead to good or poor performance?

(3+3+4)

6. (a) Compare and contrast A* search and Greedy Best-First Search in terms of their use of heuristic functions, efficiency, memory usage, and optimality.

   (b) Sketch a proof of the optimality of A* under an admissible heuristic tree-search paradigm. Why is consistency more critical in graph-based A* search compared to tree-based search?

(4+6)