



PAMIBIA UNIVERSITY
OF SCIENCE AND TECHNOLOGY

FACULTY OF COMPUTING AND INFORMATICS
DEPARTMENT OF SOFTWARE ENGINEERING

QUALIFICATION: BACHELOR OF COMPUTER SCIENCE (SOFTWARE DEVELOPMENT)	
QUALIFICATION CODE: 07BACS	LEVEL: 7
COURSE: ARTIFICIAL INTELLIGENCE	COURSE CODE: ARI711S
DATE: JULY 2025	PAPER: THEORY
DURATION: 3 HRS	MARKS: 100

SUPPLEMENTARY / SECOND OPPORTUNITY EXAMINATION QUESTION PAPER	
EXAMINER(S)	Mr NAFTALI INDONGO
MODERATOR	Ms PAULINA SHIFUGULA

THIS QUESTION PAPER CONSISTS OF 6 PAGES
(Including this front page)

INSTRUCTIONS TO STUDENTS

1. Read all the questions, passages, scenarios, etc., carefully before answering.
2. Answer all the questions.
3. Number each answer clearly and correctly.
4. Write neatly and legibly.
5. Making use of any crib notes may lead to disqualification and disciplinary action.
6. Use the allocated marks as a guideline when answering questions.
7. Looking at other students' work is strictly prohibited.

SECTION A: 10 MARKS

MULTIPLE CHOICE (Select the letter that corresponds to the correct answer.)

1. Which search algorithm is guaranteed to find the least-cost solution if the cost is strictly positive?
 - A. Depth-First Search
 - B. Breadth-First Search
 - C. Uniform Cost Search
 - D. Greedy Search

2. What is the main benefit of Iterative Deepening Search?
 - A. It guarantees minimal memory use with minimal runtime
 - B. It combines DFS's space advantage with BFS's optimality
 - C. It avoids redundant search paths
 - D. It expands the shallowest nodes first

3. In A* Search, what happens when a heuristic is not admissible?
 - A. Search becomes faster
 - B. Search returns no solution
 - C. The optimality of A* is not guaranteed
 - D. It behaves like Uniform Cost Search

4. What is the role of backtracking in CSPs?
 - A. It avoids searching through solution space
 - B. It guarantees arc consistency
 - C. It incrementally builds candidates and abandons partial assignments that cannot lead to a solution
 - D. It randomises the search space

5. The Minimum Remaining Values (MRV) heuristic helps by:
 - A. Maximising the domain size
 - B. Choosing the variable with the fewest legal values
 - C. Choosing the variable with the most constraints
 - D. Selecting variables randomly

6. What best describes a policy in a Markov Decision Process?
 - A. A list of rewards collected during execution
 - B. A mapping from actions to state transitions
 - C. A strategy that maps each state to a specific action
 - D. A probability distribution over terminal states

7. What role does the discount factor (γ) play in MDPs?
 - A. It determines how many actions the agent can take in each state
 - B. It represents the probability of reaching the terminal state
 - C. It penalises illegal actions
 - D. It quantifies how much an agent prefers current rewards over future rewards

8. What is the purpose of policy extraction in solving MDPs?
 - A. To extract and discard poor-performing policies
 - B. To determine the optimal set of rewards based on current actions
 - C. To generate an action plan by selecting the best action from each state based on computed utilities
 - D. To simulate agent behaviour during training

9. Which of the following best describes the core objective of unsupervised learning?
 - A. Discover patterns in unlabelled data.
 - B. Learn by interacting with the environment through trial and error
 - C. Learn a function that maps inputs to known outputs.
 - D. Generate random outputs from input distributions.

10. What is the primary purpose of an activation function in a neural network?
 - A. To introduce non-linearity into the network
 - B. To remove non-linearity
 - C. To scale weights
 - D. To reduce the size of the input

SECTION B: 10 MARKS

TRUE/FALSE (Determine whether the following statements are True or False)

1. Breadth-First Search (BFS) is guaranteed to find the shortest path between two points in a graph.
2. The least constraining value heuristic chooses the value that eliminates the most options for other variables.
3. In cutset conditioning, we instantiate variables to reduce a CSP to a tree structure.
4. In an MDP, the reward function can be defined based on states, actions, or transitions.
5. Q-values ($Q^*(s, a)$) are always equal to or greater than the corresponding V-values ($V^*(s)$).
6. The discount factor in an MDP helps agents prioritise immediate rewards over future ones.
7. Policy iteration may converge to an optimal policy faster than value iteration under certain conditions.
8. Backtracking with filtering and ordering can efficiently solve large N-Queens problems.
9. Inductive reasoning is used in machine learning to draw general conclusions from specific examples.
10. The self-attention mechanism helps neural networks remember distant information in sequences.


SECTION C: 80 MARKS
STRUCTURED QUESTIONS

- Answer all the questions in the provided booklet.
- This section consists of 5 questions.

1. Search Problems

[25 Marks]

- 1.1. In your own words, explain what a search problem entails. What are the key components involved? **(6)**
- 1.2. Consider the graph shown below, where the numbers on the links are link costs and the numbers next to the states are heuristic estimates. Note that the arcs are undirected. Let A be the start state and G be the goal state.



Simulate A* search with a strict expanded list on this graph. At each step, show the path to the state of the node that is being expanded, the length of that path, the total estimated cost of the path (actual + heuristic), and the current value of the expanded list (as a list of states). You are welcome to use scratch paper or the back of the exam pages to simulate the search. However, please transcribe (only) the information requested into the table below: **(10)**

Path to State Expanded	Length of the Path	Total Cost	Estimated	Expanded List (Closed list)
A	0	5		(A)


- 1.3. Is the Heuristic given in 1.2 admissible? Explain. **(2)**
- 1.4. Is the Heuristic given in 1.2 consistent? Explain **(3)**
- 1.5. Did the A* algorithm with a strictly expanded list find the optimal path in 6.1? If it did, find the optimal path, and explain why you would expect that. If it didn't find the

optimal path and give a simple (specific) change of state values of the heuristic that would be sufficient to get the correct behaviour. (4)

2. Constraint Satisfaction Problems

[15 Marks]

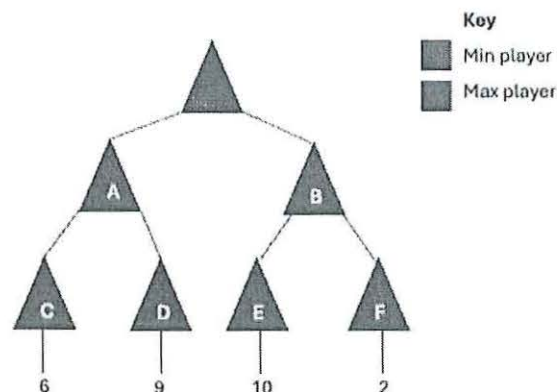
- 2.1. What is the difference between explicit and implicit constraints in CPS? (4)
- 2.2. Explain the difference between unary, binary, and higher-order constraints with suitable examples. (3)
- 2.3. Show the sequence of variable assignments during backtracking with forward checking (BT+FC) in a colouring problem; assuming that the variables are examined in numerical order and the values are assigned in the order shown next to each node. Show the assignments by writing the variable number and the value, e.g., 1R. Show each step to the solution. (8)



3. Adversarial Search

[10 Marks]

- 3.1. Consider a two-player game with the following game tree. Assume the MAX player wants to maximise their score and the MIN player wants to minimise the MAX player's score. Calculate the minimax value for each node, assuming each player plays optimally. (6)



3.2. Explain how alpha-beta pruning works and how it helps improve the efficiency of minimax search. Provide an example where alpha-beta pruning significantly reduces the number of nodes explored. (4)

4. Markov Decision Processes [15 Marks]

4.1. In the context of Markov Decision Processes (MDPs), explain the meaning of the following four quantities and how they relate to each other: (8)

- a) Policy
- b) Utility
- c) Value
- d) Q-Values

4.2. What are the main steps involved in policy iteration? (2)

5. Machine Learning [15 Marks]

Machine algorithms can be categorised based on the type of data they use and the learning approaches they employ. Explain the fundamental differences between supervised learning, unsupervised learning, and reinforcement learning and provide examples of common algorithms or applications associated with them.

Supervision (3)

Learning Goal (3)

Data Interaction (3)

Examples (Provide 2 examples for each in the table below) (6)

Supervised Learning	Unsupervised Learning	Reinforcement Learning

***** End of the Paper *****