



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: JUNE 2025	PAPER: THEORY
DURATION: 3 HRS	MARKS: 100

FIRST OPPORTUNITY EXAMINATION QUESTION PAPER	
EXAMINER(S)	Mr NAFTALI INDONGO
MODERATOR	Ms PAULINA SHIFUGULA

THIS QUESTION PAPER CONSISTS OF 8 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: 15 MARKS

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

1. What distinguishes a **reflex agent** from a planning agent?
 - A. Reflex agents always act randomly
 - B. Reflex agents consider the future consequences of their actions
 - C. Reflex agents act based only on the current percept of the environment
 - D. Reflex agents require a goal test

2. In a search problem, the **successor function** defines:
 - A. The set of all possible states
 - B. The transition from one state to another with actions and costs
 - C. The solution plan
 - D. The final goal state

3. Which of the following search strategies uses a LIFO stack for its fringe?
 - A. Breadth-First Search
 - B. Uniform Cost Search
 - C. Iterative Deepening Search
 - D. Depth-First Search

4. Greedy Best-First Search evaluates nodes based on:
 - A. $g(n)$
 - B. $h(n)$
 - C. $f(n) = g(n) + h(n)$
 - D. depth of the node

5. Which of the following is TRUE about an admissible heuristic in A* search?
 - A. It always returns zero
 - B. It never underestimates the true cost to reach the goal
 - C. It guarantees minimal memory usage
 - D. It prefers deeper nodes

6. What is the main advantage of using Iterative Deepening Search (IDS)?
 - A. It is faster than DFS
 - B. It uses less memory than BFS and finds optimal solutions
 - C. It finds the longest path in a graph
 - D. It does not revisit any node

7. Which of the following is NOT a component of a CSP?
 - A. Variables
 - B. Domains
 - C. Heuristics
 - D. Constraints

8. Which filtering technique removes inconsistent values from the domains of future variables after assigning a value?

- A. Constraint Propagation
 - B. Forward Checking
 - C. Genetic Algorithm
 - D. Cutset Conditioning
9. What is the key idea behind the Min-Conflicts heuristic?
- A. Always choose the variable with the largest domain
 - B. Randomly reassign values to all variables
 - C. Select the value that results in the fewest constraint violations
 - D. Avoid changing any variable value
10. What type of games does adversarial search primarily deal with?
- A. Cooperative games
 - B. Randomised games
 - C. Zero-sum, deterministic games
 - D. Solitaire games
11. In the context of MDPs, what does the "Markov property" imply?
- A. Given the present state, the future and the past are independent.
 - B. The future depends on both the past and present
 - C. The system is deterministic
 - D. The policy must be fixed
12. Which of the following statements about the Bellman equation is correct?
- A. It computes the exact reward of each state.
 - B. It defines the transition dynamics of an MDP.
 - C. It recursively relates the value of a state to its expected rewards and future state values.
 - D. It is used only in deterministic environments.
13. Which of the following best describes the core objective of supervised 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.
14. Which of the following best describes a neural network?
- A. A rule-based decision-making algorithm
 - B. A structure that encodes decisions as hard-coded rules
 - C. A tree of decisions with binary splits
 - D. combination of units that compute weighted sums, followed by activations
15. What is the primary purpose of an activation function in a neural network?
- A. To remove non-linearity
 - B. To scale weights
 - C. To reduce the size of the input
 - D. To introduce non-linearity into the network

SECTION B: 15 MARKS

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

1. Depth-First Search (DFS) is always the best choice for finding all possible solutions to a search problem.
2. Local search methods for CSPs require a complete assignment of all variables to start.
3. Heuristic functions are used to evaluate the potential cost of reaching the goal state in a search problem.
4. A CSP involves assigning values to variables such that all the constraints between them are satisfied.
5. Forward checking ensures consistency for all pairs of variables in the problem.
6. In an MDP, the agent has complete knowledge of the entire state of the environment at any given time.
7. MDPs are a powerful tool for modelling and solving decision-making problems in various domains, such as robotics, game playing, and resource allocation.
8. Q-values ($Q^*(s, a)$) are always equal to or greater than the corresponding V-values ($V^*(s)$).
9. The Bellman equation is a key formula used in solving MDPs to determine the optimal value function and policy.
10. Minimax search is used for solving single-agent deterministic search problems.
11. Alpha-Beta pruning can be used to improve the efficiency of the minimax search algorithm in two-player zero-sum games.
12. Markov Decision Processes (MDPs) are deterministic.
13. In machine learning, supervised learning involves labelled training data.
14. Neural networks only perform well when the data is linearly separable.
15. Deep learning models require large datasets and significant computational resources to perform well.

SECTION C: 70 MARKS
STRUCTURED QUESTIONS

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


1. Search Problems

[30 MARKS]

- 1.1. Explain the difference between planning and identification problems. (4)
- 1.2. In your own words, explain what a search problem entails. What are the key components involved? (8)
- 1.3. Considering the properties of search algorithms, which property would be most relevant to add to the following descriptions? (4)

Property	Description
	A measure of how much memory an algorithm needs to run.
	A measure of how long it takes the algorithm to find a solution (typically based on the number of steps involved)
	A guarantee that the algorithm will find the solution with the lowest cost among all possible solutions.
	A guarantee that the algorithm will find a solution if one exists.

- 1.4. Consider the tree shown below. The numbers on the arcs are the arc lengths; the numbers near states B, C and D are the heuristic estimates; all other states have a heuristic estimate of 0. (14)



Assume that the children of a node are expanded alphabetically when the search specifies no order and that the goal is state J. No visited or expanded lists are used. In what order would each type of search expand the states? Write only the sequence of states expanded by each search.

Search Type	List of States
Breadth-First Search	

Depth-First Search	
Uniform Cost Search	
Greedy Best-First Search	
A* Search	

2. Constraint Satisfaction Problems

[20 Marks]

- 2.1. Define the Constraint Satisfaction Problem (CSP) by describing its key components. (10)
- 2.2. Differentiate between explicit and implicit constraints in the CSPs. (4)
- 2.3. You need to schedule guest lectures for 7 historical figures within 4 time slots (1 pm-4 pm). Certain attendee groups want to see specific speakers concurrently (e.g., physics students want to see both Bohr and Newton). Your task is to create a schedule that satisfies these attendance preferences.

The guest lecturers list includes Alan Turing, Ada Lovelace, Niels Bohr, Marie Curie, Socrates, Pythagoras, and Isaac Newton.

Turing has a scheduling constraint. He can only be assigned the 1 pm time slot.

The Course VIII students want to see the physicists: Bohr, Curie and Newton.

The Course XVIII students want to see the mathematicians: Lovelace, Pythagoras and Newton.

The members of the Ancient Greek club want to see the ancient Greeks: Socrates and Pythagoras.

The visiting Wellesley students want to see the female speakers: Lovelace and Curie.

The CME students want to see the British speakers: Turing, Lovelace and Newton.


Finally, you decide that you will be happy if and only if you get to see both Curie and Pythagoras. (Yes, even if you belong to one or more of the groups above.)

Draw a constraint graph using the initials of each pair of lecturers who cannot share a time slot. (6)

Domains for this problem

T	1
L	1,2,3,4
B	1,2,3,4
C	1,2,3,4
S	1,2,3,4
P	1,2,3,4
N	1,2,3,4

Constraint graph for this problem




3. Adversarial Search

[10 Marks]

1.1. What is the main objective of the minimax algorithm in adversarial search? (2)

1.2. Consider the game search tree below. Assume the root node is the max node. The labels on the arcs are the moves. The numbers in the bottom layers are the values of the different outcomes of the game for the max player.



a) What is the value of the game to the max player? (3)
We apply **minimax** recursively:


b) What first move should the max player make? (2)

c) Assuming the max player makes that move, what is the best next move for the min player?

4. Markov Decision Processes

[5 Marks]

Consider the MDP given in the figure below. Assume the discount factor $\gamma = 0.9$. The r-values are rewards, while the numbers next to arrows are the probabilities of outcomes. Note that only the state S_1 has two actions. The other states have only one action for each state.



Write down the numerical value of $V(S_1)$ after the first and the second iterations of the Value Iteration algorithm. (5 Marks)

$$J^0(S_0) = 0; J^0(S_1) = 0; J^0(S_2) = 0; J^0(S_3) = 0;$$

$$J^1(S_1) =$$

$$J^2(S_1) =$$

5. Machine Learning

[5 Marks]

The following table shows a data set and the number of times each point is misclassified during a run of a perceptron algorithm, starting with zero weights. What is the equation of the separating line found by the algorithm as a function of $x_1, x_2,$ and x_3 ? Assume that the learning rate is 1 and the initial weights are all zero.

x_1	x_2	x_3	y	Times misclassified (α_i)
2	3	1	+1	12
2	4	0	+1	0
3	1	1	-1	3
1	1	0	-1	6
1	2	1	-1	11

Use the equation below:

$$\bar{\omega} = \eta \sum_{i=1}^m \alpha_i y^{(i)} \bar{x}^{(i)}$$

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