PRIFYSGOL ABERYSTWYTH UNIVERSITY

SEMESTER TWO 2026

COMPUTER SCIENCE

CS26520 : ARTIFICIAL INTELLIGENCE

Time allowed: **2 HOURS**

Answer all **FOUR** questions

All questions carry equal marks.

YOU MAY BRING THE FOLLOWING INTO THE EXAMINATION:
UP TO 3 SHEETS OF NOTES ON A4 PAPER (WHICH MAY BE DOUBLE-SIDED). THERE ARE NO ADDITIONAL RESTRICTIONS ON CONTENT OR FORMATTING OF THESE NOTES.
NO ELECTRONIC DEVICES.

CALCULATORS ARE NOT PERMITTED.

STUDENTS MAY SUBMIT ANSWERS TO THIS PAPER IN EITHER WELSH OR ENGLISH

Once seated, students may complete the front cover of the answer book(s) and the attendance slip ONLY.

DO NOT TURN THE EXAMINATION QUESTION PAPER OVER UNTIL INSTRUCTED TO DO SO.

1. This question concerns fuzzy logic and fuzzy sets.

a) Briefly discuss one key advantage of fuzzy logic over classical (Boolean) logic. [2 marks]

b) You are modeling a car buyer's preferences using fuzzy sets. The universe of discourse consists of 5 different car models. Three fuzzy sets are defined for the concepts of being *Affordable* (A), *Fuel Efficient* (E), and *Safe* (S).

A (Affordable) = {1.0/Car1, 0.8/Car2, 0.5/Car3, 0.2/Car4, 0.1/Car5}

E (Efficient) = {0.3/Car1, 0.6/Car2, 0.9/Car3, 0.7/Car4, 0.4/Car5}

S (Safe) = {0.4/Car1, 0.7/Car2, 0.8/Car3, 0.9/Car4, 1.0/Car5}

A "desirable" car is defined as one that is **Safe AND (Affordable OR Fuel Efficient)**.

(i) First, calculate the fuzzy set for the concept of being "value-conscious", defined as $V = A \vee E$. [2 marks]

(ii) Now, using your result from part (i), calculate the fuzzy set for a "desirable" car, defined as $D = S \wedge V$. [2 marks]

(iii) Based on your final fuzzy set D, which car model is the most "desirable"? Justify your answer by explaining what the final membership value for that car represents in the context of the problem. [3 marks]

c) You are designing a fuzzy logic controller for a car's autonomous emergency braking system. The system determines the required braking force based on the distance to the vehicle ahead and the relative speed at which the car is closing the gap. The system uses the following rules:

R1: If Distance is Far, then BrakingForce is None

R2: If Distance is Medium, then BrakingForce is Light

R3: If Distance is Close or RelativeSpeed is ClosingFast, then BrakingForce is Hard

The fuzzy sets for Distance (in metres) are:

Close = {1.0/0, 0.8/15, 0.4/30, 0.1/45, 0.0/60}

Medium = {0.0/0, 0.0/15, 0.5/30, 1.0/45, 0.5/60}

Far = {0.0/0, 0.0/15, 0.1/30, 0.3/45, 0.6/60}

The fuzzy sets for RelativeSpeed (in km/h) are:

Stable = {1.0/0, 0.7/10, 0.3/20, 0.0/30, 0.0/40}

ClosingFast = {0.0/0, 0.0/10, 0.4/20, 0.8/30, 1.0/40}

The fuzzy sets for BrakingForce (as a percentage of maximum) are:

None = {1.0/0, 0.6/20, 0.2/40, 0.0/60, 0.0/80, 0.0/100}

Light = {0.0/0, 0.5/20, 1.0/40, 0.5/60, 0.0/80, 0.0/100}

Hard = {0.0/0, 0.0/20, 0.0/40, 0.3/60, 0.7/80, 1.0/100}

Calculate the final BrakingForce (using Mamdani inference and centre-of-gravity defuzzification) for a situation where the car ahead is at a **Distance of 30 metres** and the **RelativeSpeed is 20 km/h**. You can leave your defuzzified value as a fraction. Show your workings. [12 marks]

d) Discuss two challenges in designing the membership functions for a fuzzy logic system. [4 marks]

2. This question concerns classification and clustering.

a) Compare and contrast the k-means clustering algorithm with the k-Nearest Neighbour (k-NN) classification algorithm. Highlight differences in their purpose, process (how the algorithms work), and computational demands. [6 marks]

b) Given the one-dimensional dataset {3, 6, 10, 15, 20, 25}, apply k-means to this data until convergence, using the initial centroids 7 and 22. You should round down centroids if these are not whole numbers (e.g., 6.33 becomes 6). Show your working and give the final clustering produced. [6 marks]

c) Briefly discuss one advantage that Hierarchical Agglomerative Clustering (HAC) has over k-means. [2 marks]

d) You have the following dataset about whether a student passes or fails a module based on certain features. The class 'Pass' indicates the final outcome.

Student	Weekly Study Hours	Assignments Done	Entry Test Score	Pass?
1	7	4	7	Yes
2	2	1	3	No
3	4	3	4	Yes
4	3	1	5	No
5	7	2	5	No

A new student's record is available, and you would like to predict their outcome. The data is:

Student	Weekly Study Hours	Assignments Done	Entry Test Score
6	5	3	6

(i) Decide, using 3-NN and the Euclidean distance measure, whether this student is predicted to pass or fail. (You may leave distances in their square root form, e.g. $\sqrt{6}$.) [8 marks]

(ii) For this data using 3-NN there is not an agreement between all three neighbours. Briefly discuss this issue in general and how it may be addressed. [3 marks]

3. This question concentrates on the area of search.

Figure 1 shows the search tree for a problem with two goal nodes, $G1$ and $G2$. The cost of moving from one node to the next is indicated by a labelled arrow. For example, the cost of moving from node Q to node S is 3. The values for the heuristic measure for each node are as follows:

h(P)=10, h(Q)=8, h(R)=9, h(S)=5, h(T)=6, h(U)=3, h(V)=4, h(W)=2, h(X)=1, h(G1)=0, h(G2)=0.

The lower the value, the closer the node is estimated to be to a goal node.

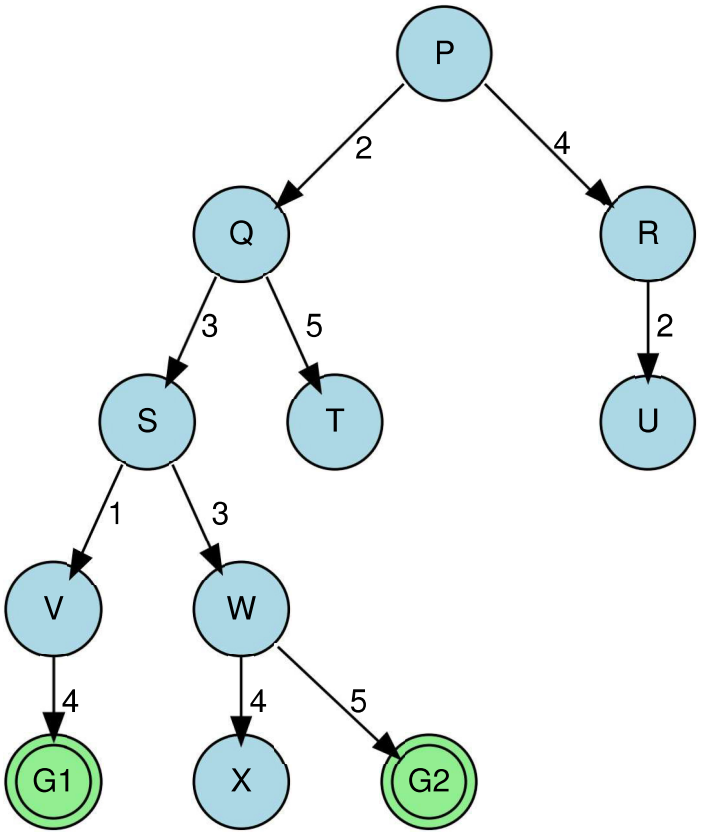


Figure 1: Search tree

a) Assuming that the nodes are processed from left to right (until a goal node is reached), give the order of node expansion for the following algorithms:

(i) Depth-first search [2 marks]

(ii) Breadth-first search [2 marks]

(iii) Depth-limited search (limit = 2) where the root node is level 0. [3 marks]

b) Discuss the primary limitation of depth-limited search as observed in the previous question. How can it be extended into an algorithm that is complete (i.e., guaranteed to find a solution if one exists)? [5 marks]

c) The A* search algorithm is a popular algorithm for search.

(i) State the full evaluation function, $f(n)$, used by the A* search algorithm, explaining what each component of the function represents. [3 marks]

(ii) Using a priority queue, show all the steps taken by the A* algorithm to find a path from the start node P to a goal node. At each step, you must show the contents of the priority queue and state which node is being expanded. Assume that if two nodes have the same $f(n)$ value, the one that is alphabetically first is chosen. Give the final path found. [7 marks]

(iii) Despite its guarantee of optimality, what is the biggest practical limitation of the A* algorithm when applied to very large search spaces, and why? [3 marks]

4. This question is about the area of evolutionary algorithms and swarm systems.

a) Genetic algorithms (GAs) can find solutions to difficult problems. Briefly describe how a GA works and how it can find good solutions. [6 marks]

b) Genetic Algorithms must balance the needs of exploration (searching new and unknown areas of the solution space) and exploitation (refining and improving upon the best solutions already found). Describe the distinct roles that the crossover and mutation operators play in managing this balance. [4 marks]

c) The Number Partitioning Problem involves dividing a set of numbers into two subsets such that the sum of the numbers in each subset is as close as possible. Consider the set of the first 10 positive integers: S = {1, 2, 3, 4, 5, 6, 7, 8, 9, 10}. The goal is to partition S into two disjoint (i.e., non-overlapping) subsets, Subset A and Subset B, such that the sum of the numbers in A is as close as possible to the sum of the numbers in B.

(i) Give a suitable representation of this problem for a GA and justify your choice. [4 marks]

(ii) Using your representation, show how the following partition would be encoded:
Subset A: 1, 4, 5, 8, 9
Subset B: 2, 3, 6, 7, 10 [2 marks]

(iii) Using your representation, explain how mutation could be performed. [2 marks]

(iv) Describe a potential fitness function for this problem. Your function should be designed so that the GA will be guided to find the best solution. Demonstrate that it achieves the best score when an optimal solution is evaluated. [3 marks]

d) Compare and contrast the core operation of PSO and ACO regarding:

- How individuals (particles/ants) communicate information about good solutions.
- How solutions are generated or modified.

[4 marks]

END OF PAPER