COMP3411/9814 Artificial Intelligence

Sample Exam

Instructions

1. Time allowed: 10 minutes reading + 2 hours working
2. This Exam will count for 60% of your Final Mark
3. Marks for each Question are as indicated
4. **All** Questions May be attempted
5. Unless otherwise indicated, numerical answers should be given correct to two decimal places (i.e. within 0.005 of the correct answer)
6. UNSW Approved Calculator may be used
7. No other material may be taken into the examination room
8. **Do not log out until you have saved and submitted your answers**

Question 1 (1 mark)

Which of the following defines a "discrete" task?

A. ○ percept contains all relevant information about the world

B. ○ environment does not change while the agent is thinking

C. ○ finite (or countable) number of possible percepts / actions

D. ○ current state of the world uniquely determines the next

Question 2 (1 mark)

The legged version of Robocup is:

A. ○ simulated, static, discreet, deterministic, known

B. ○ situated, dynamic, continuous, stochastic, unknown

C. ○ situated, static, discreet, deterministic, unknown

D. ○ situated, dynamic, continuous, stochastic, known

Question 3 (1 mark)

Which type of agent uses state-based search to solve complex tasks requiring many individual steps?

A. ◯ reactive agent

B. ◯ model-based agent

C. ◯ planning agent

D. ◯ learning agent

Question 4 (1 mark)

Optimality of a search algorithm answers the question:

A. ◯ Is the algorithm guaranteed to find a solution when there is one?

B. ◯ Does the strategy find the solution that has the lowest path cost of all solutions?

C. ◯ How long does it take to find a solution?

D. ◯ How much memory is needed to perform the search?

Question 5 (1 mark)

Uniform Cost Search, Greedy Search and A*Search can all be viewed as special cases of:

A. ◯ best-First Search

B. ◯ Uninformed Search

C. ◯ Informed Search

D. ◯ Local Search

Question 6 (1 mark)

The search strategy that evaluates nodes by combining the path cost from the start node to node n and the estimated cost of the cheapest path from n to any goal node is:

A. ◯ Greedy Search

B. ◯ Depth First Search

C. ◯ Bidirectional Search

D. ◯ A* Search

Question 7 (1 mark)

The principle "The most likely hypothesis is the simplest one consistent with the data" is called:

A. ◯ Perceptron Learning

B. ◯ Ockham's Razor

C. ◯ Maximum Likelihood

D. ◯ Bayes' Rule

Question 8 (3 marks)

Find a solution for this cryptarithmetic problem:

```
ONE +
ONE
ONE
ONE
---
TEN
```

i. The 3-digit number represented by ONE is:

ii. The 3-digit number represented by TEN is:

Question 9 (7 marks)

Consider the task of finding a path from start state S to goal state G, given the distances and heuristic values in this diagram:

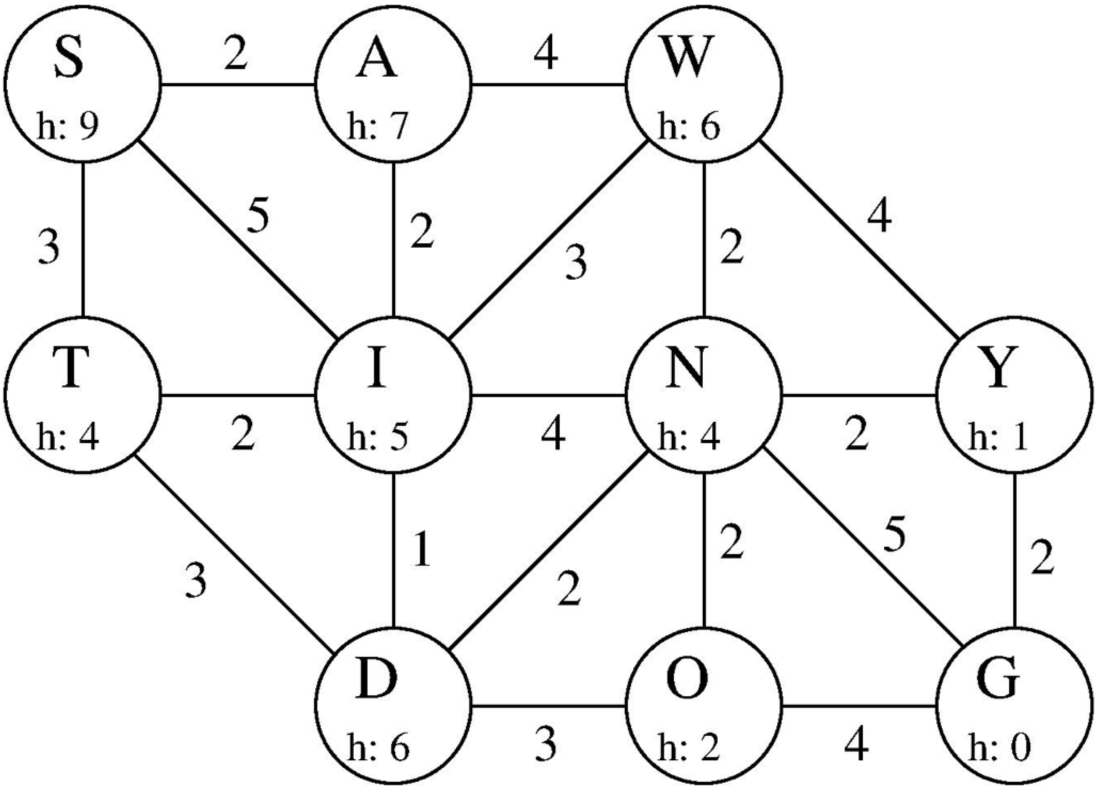

For each of the following strategies, list the order in which the states are expanded. Whenever there is a choice of states, you should select the one that comes first in alphabetical order. In each case, you should skip any states that have previously been expanded, and you should continue the search until the goal node is expanded.

i. Depth First Search

ii. Uniform Cost Search [hint: first compute g() for each state in the graph]

iii. Greedy Search, using the heuristic shown.

iv. A∗ Search, using the heuristic shown.

Question 10 (4 marks)

Consider the arrangement of tiles in the 8-puzzle shown on the left, in comparison to the Goal State shown on the right:

4	2	8
6		3
5	7	1

Initial State

1	2	3
4	5	6
7	8	

Goal State

i. What is the value of the Manhattan Distance Heuristic for this Initial State?

ii. Which of these successor states would be chosen by A* Search as the next state to be expanded?

A. ○

4		8
6	2	3
5	7	1

B. ○

4	2	8
6	3	
5	7	1

C. ○

4	2	8
	6	3
5	7	1

D. ○

4	2	8
6	7	3
5		1

Question 11 (4 marks)

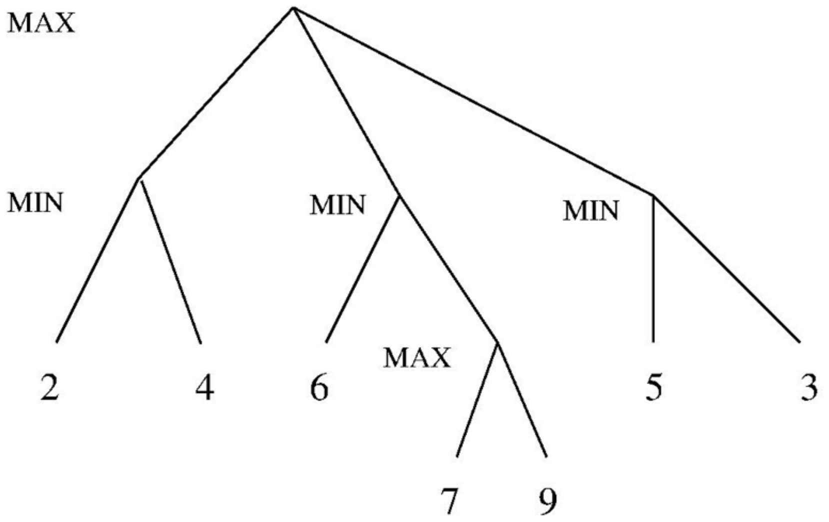

If this game tree is evaluated using alpha-beta search, which leaves would be pruned (i.e. **not** evaluated)? (Select all that apply)

A. ☐ 2

B. ☐ 4

C. ☐ 6

D. ☐ 7

E. ☐ 9

F. ☐ 5

G. ☐ 3

Question 12 (4 marks)

Consider the following probability table concerning these three events:

- rain: it is raining
- light: the light can be switched on
- failure: there is a power failure

	rain		¬ rain	
	light	¬ light	light	¬ light
failure	0.04	0.03	0.02	0.07
¬ failure	0.16	0.02	0.63	0.03

Compute (to two decimal places): Prob(light)

Compute (to two decimal places): Prob(failure | rain ∨ ¬ light)

Question 13 (4 marks)

A marketing company asks a customer to try various kinds of ice cream and report whether they like it (+) or not (−). Each ice cream has a certain flavour (chocolate, vanilla or banana) and contains a certain added ingredient (peanuts, smarties or raisins). The aim is to predict the like attribute (+ or −) from the other two attributes (flavour and ingredient).

flavour	**ingredient**	**like**
chocolate	peanuts	+
vanilla	smarties	+
chocolate	raisons	+
chocolate	smarties	+
vanilla	raisons	−
banana	raisons	−
vanilla	peanuts	−
banana	peanuts	−

(i) What is the information (entropy) gain for the **flavour** attribute?

(ii) What is the information (entropy) gain for the **ingredient** attribute?

Question 14 (4 marks)

The Laplace error estimate for pruning a node in a decision tree is given by

$$E = 1 - \left(\frac{n+1}{N+k}\right)$$

where N is the size of the total population, n is the number of examples in the majority class and k is the number of classes. In the following subtree, we assume that there are just two classes 'yes' and 'no'. The numbers in brackets are the numbers of examples in each class for that node.

```
      [5,9]
     /     \
[4,3]       [1,6]
```

Using minimal error pruning, should the children be pruned or not? Show your calculations.

Question 15 (3 marks)

Training Item	x_1	x_2	Class
a.	0	1	+
b.	2	1	–
c.	2	2	+

Construct by hand a Perceptron network which correctly classifies the above three data items. Your network should have two inputs (x_1, x_2) and one output. Specify the weights w_0, w_1, w_2 of the network.

i. value of w_0 :

ii. value of w_1 :

iii. value of w_2 :

Question 16 (3 marks)

Consider a perceptron whose output is given by $h(w_0 + w_1x_1 + w_2x_2)$ where x_1, x_2 are inputs and $h()$ is the Heaviside (step) function. The current weights are $w_0 =$ 0.5, $w_1 =$ 1, $w_2 =$ -2.

Training Item	x_1	x_2	Class
a.	3	-2	–
b.	2	1	+
c.	4	0	+

Suppose the Perceptron Learning Algorithm is applied to the current weights, using a learning rate of $\eta = 1.0$. We consider only the first step, where **item (a)** is learned. The new values for w_0, w_1 and w_2 after learning **item (a) only** will be:

i. new value of w_0 :

ii. new value of w_1 :

iii. new value of w_2 :

Question 17 (2 marks)

Consider the following multi-layer perceptron, using the threshold activation function, and assume that TRUE is represented by 1; FALSE by 0.

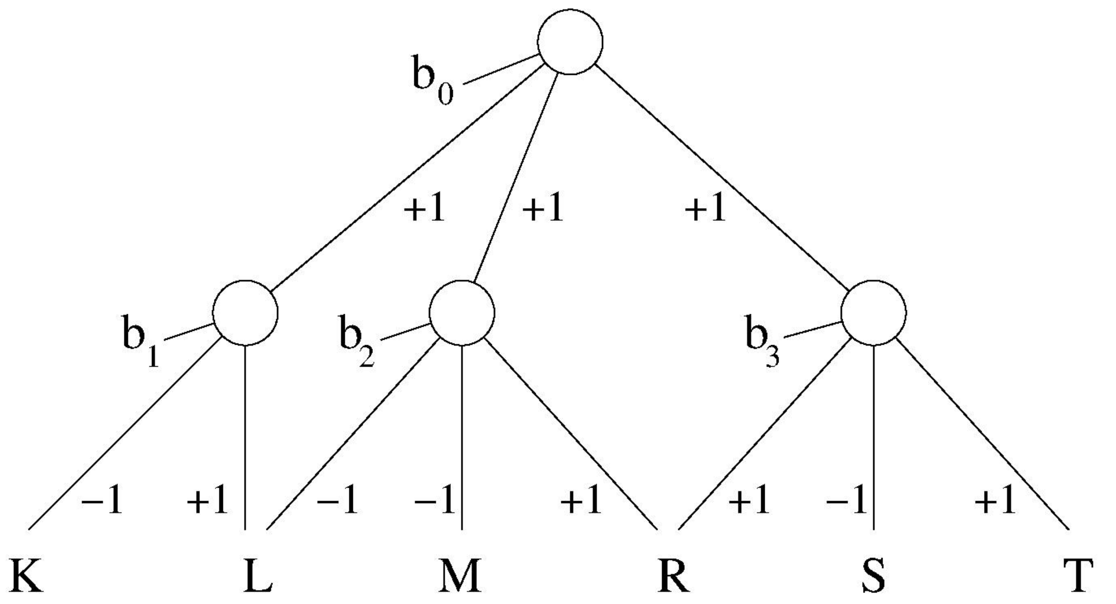

For which values of the biases b_0, b_1, b_2, b_3 would this network compute this logical function ?

(¬ K ∨ L) ∧ (¬ L ∨ ¬ M ∨ R) ∧ (R ∨ ¬ S ∨ T)

i. value of b_0 :

ii. value of b_1 :

iii. value of b_2 :

iv. value of b_3 :

Question 18 (3 marks)

Consider these two sentences on three proposition symbols K, L, M

i. $L \rightarrow K \wedge \neg M$
ii. $\neg M \wedge \neg L \rightarrow K$

Select all models which satisfy both of these sentences: (Select all that apply)

A. ☐ {}

B. ☐ { M }

C. ☐ { L }

D. ☐ { L, M }

E. ☐ { K }

F. ☐ { K, M }

G. ☐ { K, L }

H. ☐ { K, L, M }

Question 19 (2 marks)

Consider the knowledge base:

$$(A \vee B) \wedge (\neg A \vee C \vee D) \wedge (B \vee \neg C) \wedge (\neg B \vee \neg D)$$

Which of the following is a valid resolution from the knowledge base?

A. ○ $\dfrac{A\vee B,\ B\vee\neg C}{A\vee\neg C}$

B. ○ $\dfrac{B\vee\neg C,\ \neg B\vee\neg D}{C\vee D}$

C. ○ $\dfrac{A\vee B,\ \neg A\vee C\vee D}{B\wedge C\wedge D}$

D. ○ $\dfrac{\neg A\vee C\vee D,\ B\vee\neg C}{\neg A\vee B\vee D}$

Question 20 (3 marks)

Translate these logical expressions into plain English:

i. $\neg\forall s\ (\text{Swan}(s) \rightarrow \text{White}(s))$

ii. $\forall c\ \text{Clown}(c) \rightarrow \exists b\ \text{Big}(b) \wedge \text{Red}(b) \wedge \text{Bicycle}(b) \wedge \text{Rides}(c, b)$

Question 21 (7 marks)

Consider an environment with two states $S = \{S_1, S_2\}$ and two actions $A = \{a_1, a_2\}$, where the (deterministic) transitions δ and reward R for each state and action are as follows:

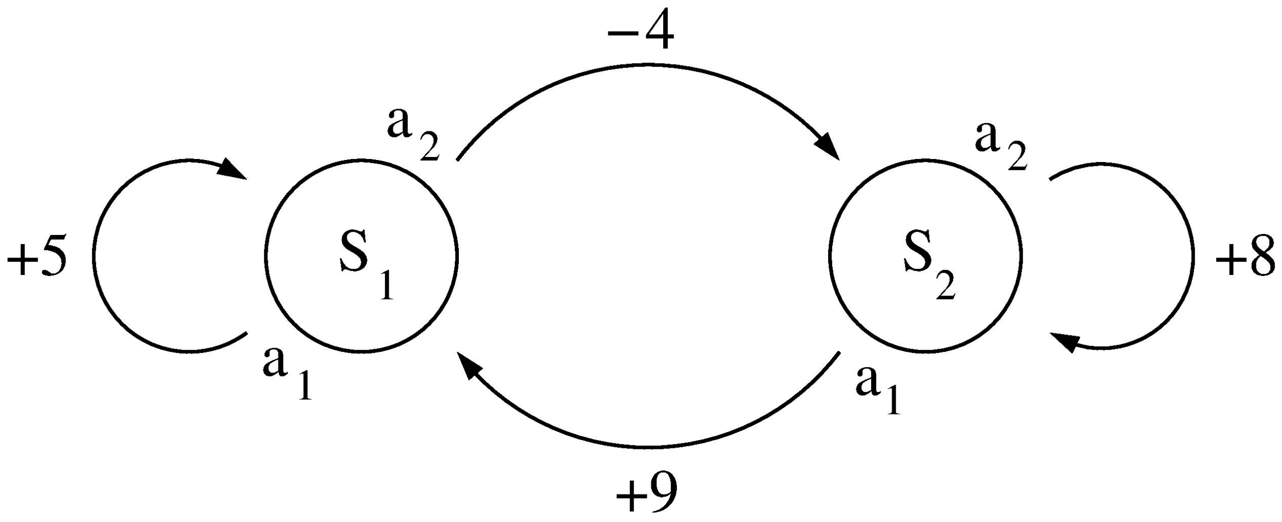

With a discount factor of $\gamma = 0.6$, the optimal policy for this environment is

$$\pi^*(S_1) = a_1, \quad \pi^*(S_2) = a_2.$$

Using this policy, and assuming $\gamma = 0.6$, compute:

i. $Q^*(S_1, a_1)$:

ii. $Q^*(S_1, a_2)$:

iii. $Q^*(S_2, a_1)$:

iv. $Q^*(S_2, a_2)$:

If γ is allowed to vary between 0 and 1, for which range of values of γ is this policy optimal?

v. **Minimum** value of γ :

vi. **Maximum** value of γ :

End of Exam