DESK No.

FAMILY NAME:

GIVEN NAMES:

SIGNATURE:

STUDENT NUMBER:

THE UNIVERSITY OF WESTERN AUSTRALIA

SEMESTER 2, 2018 EXAMINATIONS

CITS3001

Physics, Mathematics & Computing
EMS

Algorithms, Agents and Artificial Intelligence

This paper contains: **6** Pages **(including title page)**

Time Allowed: **2:00** hours

INSTRUCTIONS:

Answer all questions in the answer book provided
Each question is worth 10 marks.
The total for the paper is 50.

THIS IS A CLOSED BOOK EXAMINATION

SUPPLIED STATIONERY	**ALLOWABLE ITEMS**
1 x Answer Booklet 18 Pages	**UWA Approved Calculator with Sticker**

This page has been left intentionally blank

Q1. Optimisation

(a) The *activity selection problem* is, given a set of activities a_i for i = 1 to n, where activity a_i has start time s_i and end time t_i, find the largest set of activities such that none of the activities in the set overlap in time. Describe an algorithm to solve the activity selection problem, and explain why it will always give an optimal solution.

4 marks

(b) Using the geometric version of the travelling salesman problem as an example, describe how a gradient based search can find a good approximate solution to an optimisation problem.

3 marks

(c) With respect to your answer to question (b), explain how *simulated annealing* can help avoid local optima.

3 marks

Q2. Search

(a) Explain what it means for a search strategy to be optimal.

2 mark

(b) Explain what it means for a search strategy to be complete.

2 mark

(c) Consider a variant of the *word-chess* problem where you are given a start word s, a target word t, and a dictionary and you need to find the shortest sequence of words from the dictionary $s_0, s_1, \ldots, s_n$ where $s_0 = s$, $s_n = t$, and s_{i+1} is the word s_i rearranged with a single letter inserted or removed. For example given start word `cat` and end word `dog`, we could have the sequence: `cat`, `at`, `tag`, `goat`, `got`, `go`, `dog`. Describe what you think is the best uninformed search strategy for this problem, and explain why.

3 marks

(d) Explain the A* algorithm in the context of the word-chess variant from question (c), and propose an admissable heuristic that can guide the search.

3 marks

Q3. Game-playing

(a) Carefully describe how the *minimax algorithm* works on the following game tree, assuming nodes are explored left to right.

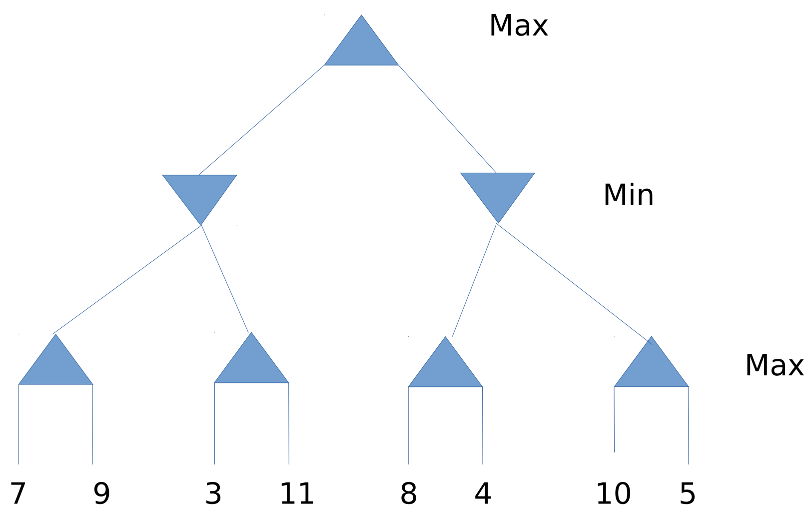

3 marks

(b) Describe how **αβ** search works on the game tree in part (a). How many nodes would be expanded given the presented ordering, and how many nodes would be expanded given an optimal ordering?

3 marks

(c) *Pig* is a simple dice game. A player rolls a six sided die as many times as they like, and the numbers they roll are added to their score, unless they roll a 1, in which case their score is 0, and their game is over. They can choose to stop rolling at any time, and their aim is to maximise their score. Express this game as a sequential decision problem, and describe each part of the Bellman equation in the context of this game.

4 marks

Q4. Learning agents

(a) The CITS3001 project this semester featured the game *Hanabi,* and it is assumed that you are familiar with the game. Suppose that we have observed an agent playing several games, and built a table showing what actions they played, depending on what state the game was in (how many cards had been discarded, whether thay had a playable card, whether someone else had a playable card, how many hints were remaining and how many fuse tokens were left). A small section of the table is below:

Discards	Can play card	Other can play card	Hints remaining	Fuse remaining	Action
20	Yes	Yes	1	2	Play
15	No	Yes	3	1	Hint
23	Yes	Yes	6	3	Play
30	No	No	4	3	Discard
15	Yes	Yes	5	1	Hint
12	Yes	No	3	2	Play
23	No	No	1	3	Play
27	No	No	0	1	Discard
3	No	Yes	8	1	Hint

Describe the process of inducing a decision tree from this data. (You do not have to build the full tree, but you should describe the required steps).

4 marks

(b) Describe the process of *temporal-difference learning.*

3 marks

(c) Describe the process of *Q-learning* and give its advantages and disadvantages relative to temporal difference learning.

3 marks

Q5. Reasoning and Planning

(a) Describe what is meant by a *knowledge base*. How does a knowledge base differ from a database?

2 marks

(b) Suppose we have a knowledge base for an agent exploring the pictured *Wumpus* World.

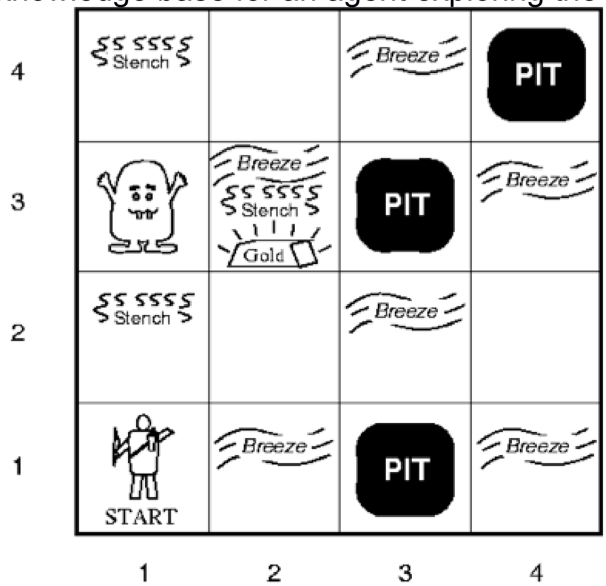

The agent can move *up*, *down, left*, or *right*. If the agent is in a cell next to a *pit* they can feel a *breeze,* and if the agent is in a cell next to the *Wumpus* they can smell a *stench.* Formulate these rules as sentences of first order logic, and describe how the agent's knowledge base evolves as the agent moves from cell (1,1) to (1,2) back to (1,1), and across to (2,1).

5 marks

(c) Explain the key concepts of a partial order planner, using a simple game (like Hanabi) as an example.

3 marks

END OF PAPER