THE UNIVERSITY OF WESTERN AUSTRALIA

Computer Science and Software Engineering

SEMESTER 2, 2017 EXAMINATIONS

CITS3001
Algorithms, Agents and Artificial Intelligence

FAMILY NAME: _____ GIVEN NAMES: _____

STUDENT ID: | | | | | | | | | SIGNATURE: _____

This Paper Contains: **5** pages **(including title page)**
Time allowed: **2:00 hours (including reading time)**

INSTRUCTIONS:

Answer all questions. Each question is worth 10 marks. The total for the paper is 100.

Most questions require only brief answers: point form answers are fine where appropriate.

PLEASE NOTE

Examination candidates may only bring authorised materials into the examination room. If a supervisor finds, during the examination, that you have unauthorised material, in whatever form, in the vicinity of your desk or on your person, whether in the examination room or the toilets or en route to/from the toilets, the matter will be reported to the head of school and disciplinary action will normally be taken against you. This action may result in your being deprived of any credit for this examination or even, in some cases, for the whole unit. This will apply regardless of whether the material has been used at the time it is found.

Therefore, any candidate who has brought any unauthorised material whatsoever into the examination room should declare it to the supervisor immediately. Candidates who are uncertain whether any material is authorised should ask the supervisor for clarification.

Supervisors Only - Student left at:

This page has been left intentionally blank

Q1. Algorithms

(a) Describe how *dynamic programming* (DP) is used to solve the Longest Common Subsequence Problem. **3 marks**

(b) Describe how the Knuth-Morris-Pratt pattern matching algorithm works. **3 marks**

(c) What is the complexity of MergeSort? Justify your answer. **4 marks**

Q2. Optimisation

(a) Describe an optimal algorithm to solve the Activity Selection Problem. **3 marks**

(b) What is meant by the term *NP-hard*?
Give an example of an NP-hard problem. **3 marks**

(c) Define the *local optima* of a state-space. Describe how *Simulated Annealing* tries to avoid local optima. **4 marks**

Q3. Uninformed search

(a) Given a search tree with branching factor *b*, depth *m*, and least cost solution at depth *d*, what is the time and space complexity of a depth first search? **3 marks**

(b) Is *iterative deepening search* optimal and/or complete?
Explain your answer. **3 marks**

(c) The Eight Puzzle is a 3x3 grid containing eight tiles numbered 1 to 8, and an empty square. Tiles may be slid about the grid to arrange the tiles in order. Explain a bidirectional search in the context of the Eight Puzzle. **4 marks**

Q4. Informed search

(a) Describe, with an example, the PATHMAX modification for the *A** algorithm. What problem does it solve? **3 marks**

(b) Explain, with an example, what it means for one heuristic to dominate another heuristic in the A* algorithm. **3 marks**

(c) Explain how a *Simplified Memory-bounded A** maintains a bound on the memory used in the search. **4 marks**

Q5. Game-playing

(a) Describe any **two** approaches to dealing with *incompleteness* in the context of AI. **2 marks**

(b) How does the *minimax algorithm* deal with intractably large search spaces? **4 marks**

(c) Explain, with an example, how the ordering of nodes affects an **αβ** search. **4 marks**

Q6. Sequential decision problems

(a) Describe, with an example, a *polic*y in a sequential decision problem (SDP)? **2 marks**

(b) What is the difference between value iteration and policy iteration when solving a sequential decision problem? **4 marks**

(c) What are the different components of the Bellman equation? **4 marks**

Q7. Learning agents

(a) Suppose that someone can *drive*, *walk*, or catch the *bus* to work, and each day their choice is influenced by whether the weather is *fine*, whether they are running *late*, and whether the *traffic* is heavy. Given historical data on what choices they made, describe the process of inducing a decision tree to predict how they will travel to work. **6 marks**

(b) What are the **four** basic components of a learning AI agent? **4 marks**

Q8. Reinforcement learning

(a) Explain, with an example, *temporal-difference learning*. **4 marks**

(b) Demonstrate the process of *Adaptive Dynamic Programming* in the following space:

0.1	0.5	0.1

An agent may choose to move left or right, with the values in the squares being the probability that the opposite move is performed. The utility for exiting to the right is 1, and the utility for exiting to the left is -1. **6 marks**

Q9. Logical agents

(a) Use first order logic and the predicates *Student(x), Fails(x), Question(y)* and *Wrong(x, y)* to express the sentence:

A student fails only if they got every question wrong. **5 marks**

(b) Demonstrate the resolution principle using your answer to (a), to formalise the argument:

A student fails only if they got every question wrong.
Brad did not get Question 2 wrong.
Therefore Brad does not fail. **5 marks**

Q10. Planning and acting

(a) Describe the relationship between planning and first order logic. **3 marks**

(b) What is the role of topological sorting in a *partial-order planner*? **3 marks**

(c) Explain, with an example, the concept of *clobbering* in a partial order planner. **4 marks**

END OF PAPER