THE UNIVERSITY OF WESTERN AUSTRALIA

Achieve International Excellence

Computer Science and Software Engineering

SEMESTER 1, 2016 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. Dynamic programming

(a) What is the basic principle of *dynamic programming* (DP)? **3 marks**

(b) Describe briefly **two** DP applications in the context of string algorithms. **2 marks**

(c) Show how a DP solution to LCS would work for the strings *011* and *110*. **5 marks**

Q2. Optimisation algorithms

(a) Define the *activity selection* problem, and list **four** plausible greedy rules. **3 marks**

(b) What is the principle behind *iterative improvement algorithms*? **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) Why is the space complexity of *depth-first search* much better than *breadth-first*? **2 marks**

(b) Briefly describe *iterative deepening search*, and describe how it tries to get the best performance features of both breadth-first and depth-first. **4 marks**

(c) What is the principle behind *bidirectional search*? **2 marks**

(d) Describe **two** problem features that can cause problems for bidirectional search. **2 marks**

Q4. Informed search

(a) What is the difference between *informed search* and *uninformed search*? **2 marks**

(b) What does it mean for a heuristic to be *admissible*? **2 marks**

(c) Given admissible heuristics h_1 and h_2, what does it mean if h_1 *dominates* h_2? In what way will A* using h1 out-perform A* using h2? **3 marks**

(d) What is the principle behind *Simplified Memory-bounded A**? **3 marks**

Q5. Game-playing

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

(b) What is an *evaluation function* in the context of game-playing algorithms? Give an example of a linear weighted sum as an evaluation function. **3 marks**

(c) Describe the operation of the *minimax* algorithm. **3 marks**

(d) A game-playing AI usually has to make a move within a certain time limit. How does iterative deepening help with this issue? **2 marks**

Q6. Sequential decision problems

(a) What is the *transition model* of a sequential decision problem (SDP)? **2 marks**

(b) Describe the operation of the *value iteration* algorithm for solving SDPs. **4 marks**

(c) What is *policy loss* in the context of an SDP? How it is related to the maximum error in the utilities? **4 marks**

Q7. Learning agents

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

(b) How is a *decision tree* constructed from a set of examples (i.e. ({attributes}, value) pairs)? **3 marks**

(c) What are the three base cases for the recursive induction algorithm used to construct decision trees? **3 marks**

Q8. Reinforcement learning

(a) What is the difference between *utility learning* and *Q-learning*? **2 marks**

(b) Describe the operational behaviour of *temporal-difference learning*. **3 marks**

(c) What is meant by *exploration* and *exploitation* in the context of learning? **2 marks**

(d) Compare the performance of *adaptive learning* and *temporal-difference learning*. **3 marks**

Q9. Logical agents

(a) Describe and illustrate with an example the main way in which *first-order logic* is more expressive than propositional logic. **3 marks**

(b) What is an *inference system* in the context of logical agents? **2 marks**

(c) What does it mean for an inference system to be *sound* and *complete*? **2 marks**

(d) Describe and illustrate with an example what it means to *unify* two sentences in first-order logic. **3 marks**

Q10. Planning and acting

(a) Describe briefly how a *partial-order planner* works. **4 marks**

(b) What are the two remedies when step S_k clobbers step S_i? **2 marks**

(c) What are the **two** principal ways that planning agents deal with *uncertainty*? **4 marks**

END OF PAPER