UNIVERSITY OF LONDON

GOLDSMITHS COLLEGE

Department of Computing

B. Sc. Examination 2016

IS53024A
Artificial Intelligence

Duration: 2 hours 15 minutes

Date and time:

This paper is in two parts: part A and part B. You should answer ALL questions from part A and TWO questions from part B. Part A carries 40 marks, and each question from part B carries 30 marks. The marks for each part of a question are indicated at the end of the part in [.] brackets.

There are 100 marks available on this paper.

THIS PAPER MUST NOT BE REMOVED FROM THE EXAMINATION ROOM

Part A

Question 1

(a) What is a *rational agent*? [5]

(b) Define: *heuristic function.* [5]

(c) What is the admissibility criterion for heuristic functions? [5]

(d) Why is admissibility important for A^* tree-search? [5]

Question 2

(a) Demonstrate the performance of the forward chaining and backward chaining techniques for rule-based system inference. Use the following rules and facts to prove that the goal (G) is true:

```
(RULE 1 (IF ( C ) and ( K )) (THEN ( E )))
(RULE 2 (IF ( K ) and ( H )) (THEN ( F )))
(RULE 3 (IF ( F ) (THEN ( G )))
(RULE 5 (IF ( A ) (THEN ( H )))
(FACT 1 ( A ))
(FACT 2 ( K ))
```

Show the forward and backward sequences of rule firings using the rules and facts in the order in which they are given. [16]

(b) Explain the relationship between the sequences of firing the rules during forward chaining and backward chaining in this example for rule-based system inference. [4]

Part B

Question 3 This question concerns Algorithm 1

(a) What is the purpose of Algorithm 1? [5]

Algorithm 1

```
procedure WOOD CHIP WORLD(N, E)
    scatter N wood chips randomly in environment E
    place n < N artificial termites randomly in E
    loop
        if termination condition then return
        for each termite do
            MOVE TERMITE(E)
            if current location has a wood chip then
                if termite is not carrying a wood chip then
                    pick wood chip up
                else
                    put wood chip down
```

(b) Suggest a possible termination condition. [5]

(c) Assume that the environment is formed of cells and does not wrap-around. Suppose that Move Termite(E) causes, if possible, a termite to move to a neighbouring cell. Suggest an algorithm for Move Termite(E). [5]

(d) Why does Wood Chip World stipulate that the number of termites should be less than the number of wood chips? [5]

(e) Explain why the number of wood chip piles cannot increase. [5]

(f) How many wood chip piles would you expect there to be if the algorithm were run for a long time? Justify your answer. [5]

Question 4

(a) What is the objective of inductive concept learning? [4]

(b) Explain briefly how the version space learning procedure performs machine learning. [4]

(c) Describe briefly how we represent a version space of hypothetical concept descriptions for symbolic machine learning. [4]

(d) Let the following concept description language with 5 attributes be given for version space learning:

```
  a1         a2         a3           a4           a5
 ----       ----       ----      ---------       ----
|    |     |    |     |    |     |    |    |     |    |
r    q     t    z     c    s     p    g    x     y    n
```

Demonstrate learning with the candidate elimination algorithm using the following positive and negative training examples:

```
1. ( r t c p y ) +)
2. ( q z s g y ) -)
3. ( q t c x y ) +)
4. ( r t c g n ) -)
5. ( q z c x y ) +)
```

Illustrate the changes of the boundary sets after each example. [18]

Question 5

(a) What algorithm would you advocate to solve a unimodal minimisation problem? Justify your answer. [5]

(b) Consider the SIMULATED ANNEALING algorithm (algorithm 2). What is the purpose of of instruction 5: $T \leftarrow \text{schedule}(t)$? [5]

Algorithm 2

1: **procedure** SIMULATED ANNEALING
2:     $current \leftarrow$ start node
3:     $t \leftarrow 0$
4:     **loop**
5:         $T \leftarrow \text{schedule}(t)$
6:         **if** $T = 0$ **then**
7:             return $current$
8:         $next \leftarrow$ a random neighbour of $current$
9:         $\Delta E = f(next) - f(current)$
10:         **if** $u \sim U(0,1) < e^{-\Delta E/T}$ **then**
11:             $current \leftarrow next$
12:         $t \leftarrow t + 1$

(c) Explain, with reference to your answer in part(b), how SIMULATED ANNEALING can solve some multi-modal minimisation problems. [5]

(d) Consider the task of finding a plan for a robot to open a box after moving it to a certain room.

```
(RULE 1  // Robot moves from room ?x to room ?y through door ?d
         (IF   (inroom Robot ?x) (connects ?d ?x ?y))
         (THEN (DELETE (inroom Robot ?x))
               (ADD (inroom Robot ?y)))

(RULE 2  // Open the Box when the Robot and Box1 are in the desired room
         (IF   (inroom Box1 ?x) (inroom Robot ?x) (closed ?b))
         (THEN (DELETE (closed ?b))
               (ADD (open ?b ?x)))

(RULE 3  // Robot pushes ?b from room ?x to room ?y through door ?d
         (IF   (inroom ?b ?x) (inroom Robot ?x) (connects ?d ?x ?y))
 (THEN (DELETE (inroom Robot ?x))
               (DELETE (inroom ?b ?x))
               (ADD (inroom Robot ?y))
               (ADD (inroom ?b ?y)))
```

The initial working memory contains the following facts: (inroom Robot R1) (inroom Box1 R3) (closed Box1) (connects D1 R1 R2) (connects D2 R3 R2) (connects D2 R2 R3).

Demonstrate how planning can be organized using the backward chaining algorithm to prove the goal: (open Box1 R2) and give all variable bindings at each step of the algorithm. Use the most recently added facts immediately for matching the remaining rules. [15]