UNIVERSITY OF LONDON

GOLDSMITHS COLLEGE

B. Sc. Examination 2013

Computing

IS53024A Artificial Intelligence

Duration: 2 hours and 15 minutes

Date and time:

There are five questions in this paper. You should answer no more than three questions. Full marks will be awarded for complete answers to a total of three questions. Each question carries 25 marks. The marks for each part of a question are indicated at the end of the part in [.] *brackets.*

There are 75 marks available on this paper.

Electronic calculators must not be programmed prior to the examination. Calculators which display graphics, text or algebraic equations are not allowed.

THIS PAPER MUST NOT BE REMOVED FROM THE EXAMINATION ROOM

Question 1

(a) Explain what is the *first step* that is required in designing a rational software agent. Demonstrate what should be produced specifically in the first step of designing a simple reflex agent Whiteboard-Cleaner. [5]

(b) Define the state space for the Whiteboard-Cleaning world to formulate the Whiteboard Cleaner as a toy problem. Assume there are only two locations for the Whiteboard Cleaner: | Left | Right |. Draw a graph diagram to show the state space for the toy problem. Use links (directed edges) of the graph to denote the following actions: (i) LL: Move to the Left location, (ii) RR: Move to the Right location, (iii) CC: Clean the current location. [7]

(c) What is the main difference between the search processes used for implementing evolutionary genetic learning and symbolic machine learning algorithms? [3]

(d) Give the structure of the simple genetic algorithm. [4]

(e) Explain briefly how each of the three main operators: selection, crossover, and mutation in the simple genetic algorithm works. [6]

Question 2

Consider the instance of searching for a shortest path from vertex A to G of the simple weighted graph below, with its (normalised) adjacency matrix on the right. Assume that the estimated path distances from each vertex to G are, from A:2, from B:1, from C:1, from D:2, from E:2, from F:1, from G:0, from H:1, from I:2, from J:1, from K:1, and from L:2.

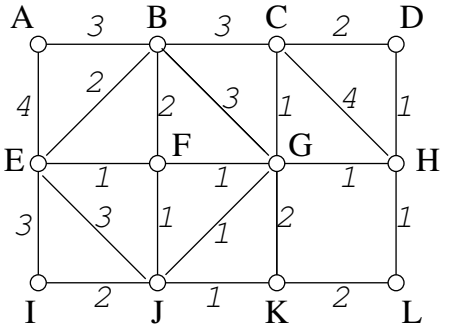

0	3	0	0	4	0	0	0	0	0	0	0
3	0	3	0	2	2	3	0	0	0	0	0
0	3	0	2	0	0	1	4	0	0	0	0
0	0	2	0	0	0	0	1	0	0	0	0
4	2	0	0	0	1	0	0	3	3	0	0
0	2	0	0	1	0	1	0	0	1	0	0
0	3	1	0	0	1	0	1	0	1	2	0
0	0	4	1	0	0	1	0	0	0	0	1
0	0	0	0	3	0	0	0	0	2	0	0
0	0	0	0	3	1	1	0	2	0	1	0
0	0	0	0	0	0	2	0	0	1	0	2
0	0	0	0	0	0	0	1	0	0	2	0

(a) Write down, as specifically as possible, the *problem definition* and the *task environment* for the instance in terms of the (i) state, (ii) action, (iii) transition model, (iv) goal test and (v) path cost. Provide an example for each of the terms to show your understanding. [10]

(b) Draw a series of diagrams to show how a search tree may be expanded step by step applying the Breadth-First Search (Graph-Search version) algorithm for the problem instance. Highlight the characteristics of the algorithm and show all your work. [4]

(c) Demonstrate how the Greedy Best-First search algorithm works. Trace step by step the expansion of the search tree, the evaluation and choice made on each vertex visited in execution of the algorithm. Highlight the characteristics of the algorithm and show all your work. [5]

(d) Demonstrate how the A* search algorithm works. Trace step by step the expansion of the search tree, the evaluation and choice made on each vertex visited in execution of the algorithm. Highlight the characteristics of the algorithm and show all your work. [6]

Question 3

(a) Consider a children game 'Tic-Tac-Toe' for two players X and O. The Tic-Tac-Toe is a pencil-and-paper game where players X and O take turns to fill the spaces in a 3×3 grid that is initially blank. Assume that X goes first, marking an 'X' at one blank place. Next O goes, marking an 'O' at another blank space, and so on. The player who succeeds in placing three respective marks in a horizontal, vertical, or diagonal row wins the game.

i. Draw part of the game tree for the Tic-Tac-Toe including at least three states (i) *X wins*, (ii) *O wins* and (iii) *Draw*, starting from the (X's turn) state X: [5]

X		X
	X	O
O		O

ii. Devise a heuristic function $h(s)$ that computes, for each state s, the difference between the number of X's choices and the number of O's choices for a legal move. Demonstrate how to compute values of the heuristic function using the state (X's turn) as an example. [3]

X		X
	X	O
O		O

iii. Let player X be a single agent. State whether or not its task environment has the following properties, and explain why. [4]

- Static
- Stochastic
- Fully observable
- Episodic

(b) Explain briefly the pruning technique. Decide the missing MinMax values in the game tree below. Use the game tree as an example to demonstrate how the $\alpha\beta$ pruning technique may be applied to avoid exploration of the entire search tree. Explain and demonstrate how elimination can be achieved in a series of diagrams. Show all your work. [13]

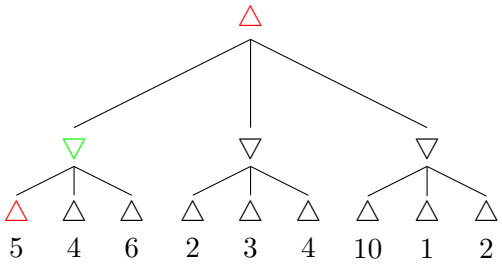

Question 4

(a) What kind of inference systems are used most often to simulate the behaviour of planning agents? Define the notion of planning as a search problem. [5]

(b) Explain whether rule-based systems are inductive or deductive systems. [4]

(c) Consider the following 5 rules describing the consequents of a possible revenge story with variables that can match any given name. Apply a rule-based inference engine operating with the forward chaining algorithm to perform automated reasoning with the given facts. Give the order in which the rules and facts are tried, and illustrate the variable bindings. [16]

```
(RULE 1  (IF   ( ?x2 convinces ?x1 to try to be king ))
         (THEN ( ?x1 has desire to be king )))

(RULE 2  (IF   ( ?x2 wants wealth )
           & ( ?x2 has capacity to influence ?x1 ))
         (THEN ( ?x2 convinces ?x1 to try to be king )))

(RULE 3  (IF   ( ?x1 has desire to be king )
           & ( ?x1 is bad )
           & ( ?x3 is king ))
         (THEN ( ?x1 kills ?x3 )))

(RULE 4  (IF   ( ?x1 is married to ?x2 )
           & ( ?x1 is easily lead ))
         (THEN ( ?x2 has capacity to influence ?x1 )))

(RULE 5  (IF   ( ?x1 kills ?x3 )
           & ( ?x4 is loyal to ?x3 ))
         (THEN ( ?x4 murders ?x1 )))

(FACT 1 ( Tudor is king ))
(FACT 2 ( John is loyal to Tudor ))
(FACT 3 ( Peter is bad ))
(FACT 4 ( Peter is easily lead ))
(FACT 5 ( Peter is married to Macbeth ))
(FACT 6 ( Macbeth wants wealth ))
```

Question 5

(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 the candidate elimination algorithm for version space learning handles positive and negative examples. [5]

(d) A concept description language with 3 attributes is given for symbolic machine learning. Assume that that these attributes can take the following values:

```
     a1              a2              a3
  ---------        ------          ------
 |    |    |      |      |        |      |
 J    M    N      L      B        C      E
```

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

```
1. ( N B C ) +)
2. ( J L E ) -)
3. ( N L C ) +)
4. ( M B C ) -)
```

Give the changes of the boundary sets after learning each example. [12]