KENYA METHODIST UNIVERSITY

END OF TRIMESTER EXAM APRIL 2008

FACULTY	**:**	**SCIENCE AND SOCIAL STUDIES**
DEPARTMENT	**:**	**COMPUTER INFORMATION SCIENCE**
COURSE CODE	**:**	**COMP 440**
COURSE TITLE	**:**	**ARTIFICIAL INTELLIGENCE**

Total Marks (60)

TIME: 2 ½ HOURS

Instructions

Answer all questions in SECTION A and ANY ONE question in SECTION B

SECTION A – Answer all questions (30 marks)

1. Define the following terms
    a. Ideal rational agent.
    b. Problem formulation
    c. State-space search
    d. Skolemization

    (4 marks)
2. Name and define the six capabilities that a computer needs to posses to pass the Total Turing test.

    (6 marks)
3. Why did Turing's test deliberately avoid direct physical interaction between the interrogator and the computer? (1 mark)

4. Explain the difference between goal-based agents and utility-based agents.

    (2 marks)

5. What is the difference between a performance measure and a utility function?

    (2 marks)

6. Describe a search space in which iterative deepening search performs much worse than depth-first search. (2 marks)

7. Consider the MIN-MAX game tree given below. (To answer this question you will have to draw two neat sketches of the tree on your answer sheet).

    a. Perform alpha-beta pruning and illustrate this on your sketch. (3 marks)
    b. Calculate the difference in branching factor before and after pruning.

    (2 marks)

c. The minimax algorithm returns the best move fro MAX under the assumption that MIN plays optimally. What happens when MIN plays suboptimally?

(3 marks)

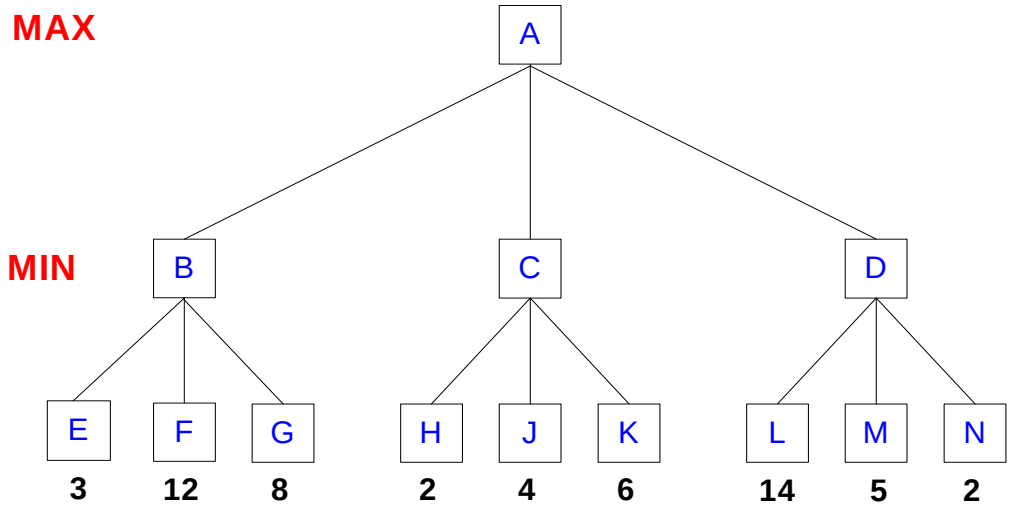

8. Give a predicate calculus sentence such that every world in which it is true contains exactly one object.

(3 marks)

9. Describe how one-point crossover in genetic algorithms works. (2 marks)

SECTION B – Answer ANY ONE question

Question 1 – 30 marks

1. The effectiveness of a search can be measured in at least three ways. Describe these three ways. (6 marks)

2. Suppose that we run a greedy search algorithms with h(n) = -g(n). What sort of search will the greedy search emulate? (2 marks)

3. The traveling salesperson problem (TSP) can be solved using the minimum spanning tree (MST) heuristic, which is used to estimate the cost of completing a tour, given that a partial tour has already been constructed. The MST cost of a set of cities is the smallest sum of the link costs of any tree that connects all the cities.

a. Show how this heuristic can be derived using a relaxed version of the TSP.
b. Show that the MST heuristic dominates straight-line distances.

(8 marks)

4. Consider the map below *(not drawn to scale)*

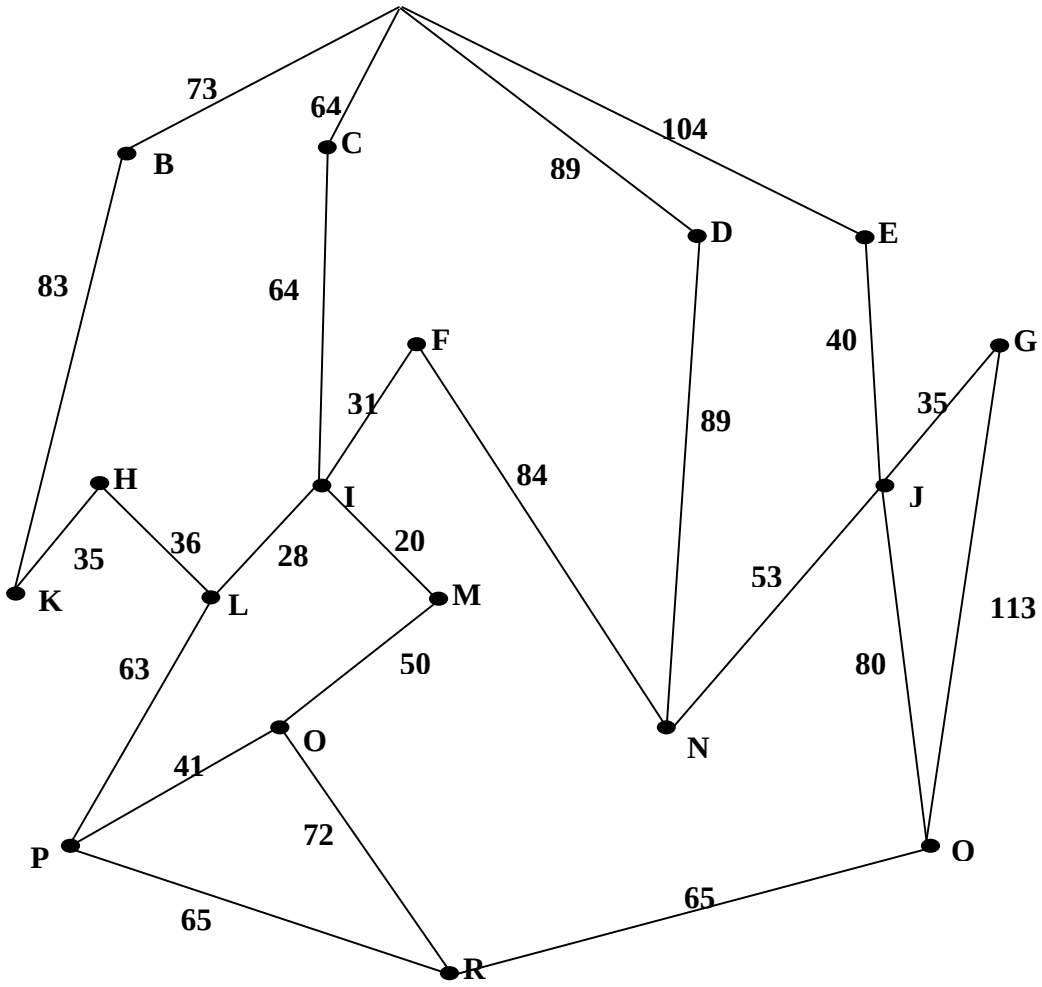

Using the A* algorithm work out a route from A to R, using the following cost functions

g(n) = the distance between each town (shown on map)

h(n) = the straight line distance between any town and town R.

These distances are given in the table below as Straight Line Distance to R

A	240		**G**	165		**M**	100
B	186		**H**	139		**N**	77
C	182		**I**	120		**O**	72
D	163		**J**	130		**P**	65
E	170		**K**	122		**Q**	65
F	150		**L**	104		**R**	0

In your answer provide the following

(i) The search tree that is produced, showing the cost function at each node. (10 marks)

(ii) State the order in which the nodes were expanded and the route that is taken, and give the total cost. (4 marks)

Question 2 – 30 marks

1. Represent the following sentences in first-order logic, using consistent vocabulary (which you must define):
    a. No person likes a professor unless the professor is smart.
    b. Every person who dislikes all vegetarians is smart.
    c. There are no blind drivers. (12 marks)

2. Represent the sentence "All Germans speak the same languages" in predicate calculus. Use *Speaks(x,l)* meaning that person *x* speaks language *l.* (4 marks)

3. Consider the problem of designing a logical agent for the wumpus world using a Boolean circuit-that is, a collection of logic gates connecting the inputs(percept values) to output (action values).
    a. Explain why you would need flip-flops.
    b. Give an order-of-magnitude estimate of how many gates and flip-flops would you need. (6 marks)

4. One of the approaches to knowledge representation and inference is structured objects. Explain where these are used. (5 marks)

5. An artificial neural network based application is used in a bank to decide which applicant gets a loan or not. Is this an expert system? Explain your answer. (3 marks)

Question 3 – 30 marks

1. Draw a decision tree for the problem of deciding whether or not to move forward at a road intersection given that the light has just turned green. (5 marks)
2. We never see the same attribute twice alone one path of a decision tree. Why not? (3 marks)

3. We know that a simple perceptron cannot represent XOR (or generally, the parity function of inputs). Describe what happens to the weights of a four-input, step-function perceptron, beginning with all the weights set to 0.1, as examples of the parity function arrive. (6 marks)

4. A perceptron is used to learn a simple function. It has two regular inputs, X_1 and X_2, and extra fixed input X_0 which always has the value 1. The perceptron's output is such that if summation is greater than 0, then output is 1, else it is 0. Given the following training set with inputs (X_0,X_1,X_2), and target output T:

X_0	X_1	X_2	T
-1	1	1	1
-1	1	0	0
-1	0	1	0
-1	0	0	0

Show the change in the weights of the perceptron for every presentation of a training instance. The initial weights are randomly set as w_1= 0.2 w_2=0.3 and w_0=0.0, the rate of learning η is 0.2. You should iterate over three epochs. Does the perceptron converge?

(10 marks)

5. Construct by hand a neural network that computes the XOR function of two inputs. Make sure you specify what sort of units you are using.

(6 marks)