BITS, PILANI – DUBAI
Academic City, Dubai
BE (Hons) CS IV Year - 1ST Sem

ARTIFICIAL INTELLIGENCE – EA C461
COMPREHENSIVE EXAMINATION (Closed Book)

Date: 3 Jan 2009
Time: 3 hrs

Max Marks: 80

Answer all questions

1. a. How does Artificial Intelligence (AI) differ from non-AI computer science? 2 mks

b. Consider the following puzzle: You have a sink (with a tap) and two measuring jars, Jar A and Jar B. Jar A holds 10 cups. Jar B holds 8 cups. Neither has any intermediate markings. How can you use these two mugs to measure 4 cups of water into Mug A? How would you construe this puzzle in terms of state space search. What are the states and the operators?
8 mks

2. a. Give three reasons why hill-climbing strategies might fail to find a usable solution.
3 mks

b. At a first glance, the iterative deepening depth first search method seems to involve a substantial amount of repeat work by re-constructing the search tree from scratch each time. Why is this not as bad as it seems? Quantify your justification.

3. Consider the following game tree. The root is a maximizing node, and children are visited left to right.

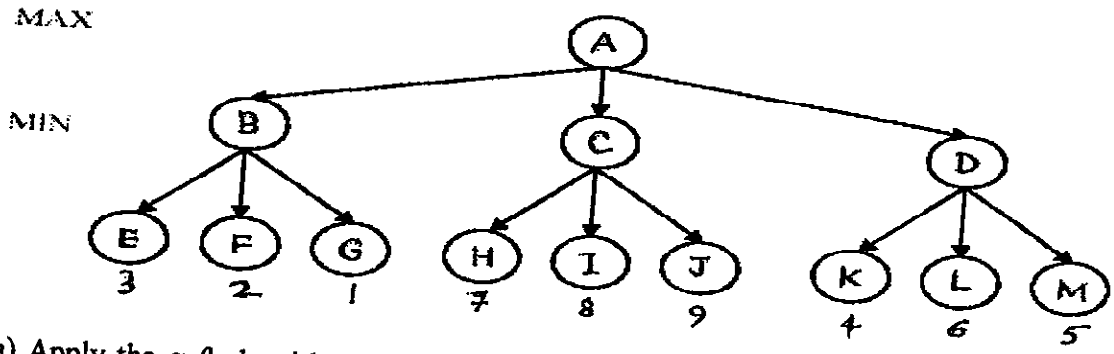

(a) Apply the α-β algorithm to the game tree above. Report the estimated values of the intermediate nodes and indicate the proper move of the maximiser
6 mks

(b) Draw a new game tree by re-ordering the children of each internal node, such that the new game tree is equivalent to the tree above, but alpha-beta pruning will prune as many nodes as possible.
4 mks

4. a. What is the difference between monotonic logic and non-monotonic logic? 2 mks

b. Consider the following statements. 8 mks

All cats like fish, cats eat everything they like, and Ziggy is a cat.

(i) Translate the sentences into FOPL.

(ii) Convert the sentences into conjunctive normal form.

(iii) Answer using resolution if Ziggy eats fish?

5. Consider the following Bayesian Network containing three Boolean random variables:

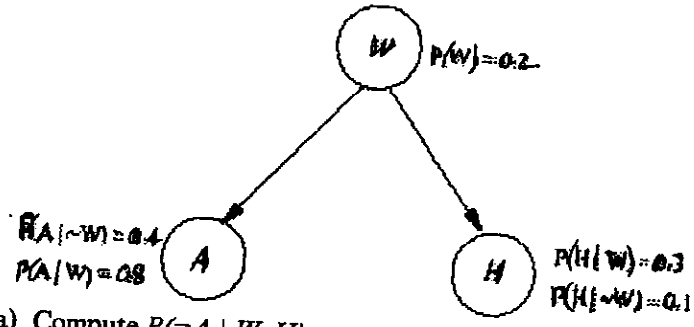

(a) Compute $P(\neg A \mid W, H)$ 2 mks

(b) Compute $P(\neg A, W, H)$ 4 mks

(c) Compute $P(\neg A \mid H)$ 4 mks

6. a. Consider a Perceptron with 3 inputs and one output unit that uses a linear threshold activation function with threshold 0.7, , and initial weights W1=0.2, W2=0.7, W3=0.9.

(i) What is the output (output = 1 if if exceeds threshold, otherwise 0) of the Perceptron given the inputs I1=1, I2=0, I3=1? 1 mk

(ii) What are the weights' values after applying the Perceptron Learning Rule with the above input and desired output 0 (learning rate (η) = 0.2)? 3 mks

(b) A standard 2-layer feed-forward neural network with 10 input units, 5 hidden units and 3 output units contains how many *weights*? (Include biases.). Draw the feedforward network and show your calculations. 6 mks

7. a. What is the Sussman anomaly? How do you solve the Sussman anomaly? 5 mks

b. Write a context free grammar and show a parse tree to correctly parse the sentence
The cat sat on the mat. 5 mks

8. Write short notes on:

a. Candidate Elimination Algorithm 5 mks

b. Information Gain 5 mks

BITS, PILANI – DUBAI
Academic City, Dubai
BE (Hons) CS IV Year - 1st Sem

Course No: EA C461
Course Title: Artificial Intelligence
Test -1 (Closed book)

Date: 2 November 2008 **Time: 50 min** **Max Marks: 25**

Note: Answer all questions.

1. Suppose you are given a full search tree with the constant branching factor b and depth d, b>0, d>0 (i.e. each node except the leaves at depth d has exactly b children). Write a precise expression (do not use O notation) of the following quantities in terms of b and d:

- The number of leaf nodes in the tree
- The number of nodes in the tree

4 mks

2. (a) For the state-space below, the path costs are shown on the links. The start state is S and the goal state is G. All search algorithms are assumed to obey the following rules:

- when a node is expanded in the search tree, the parent of the node is not added as a child
- when multiple nodes are contending, they are placed according to the alphabetical order (i.e., A goes in front of B, etc)

Otherwise, the search algorithms behave as defined in the pseudocode description in class.

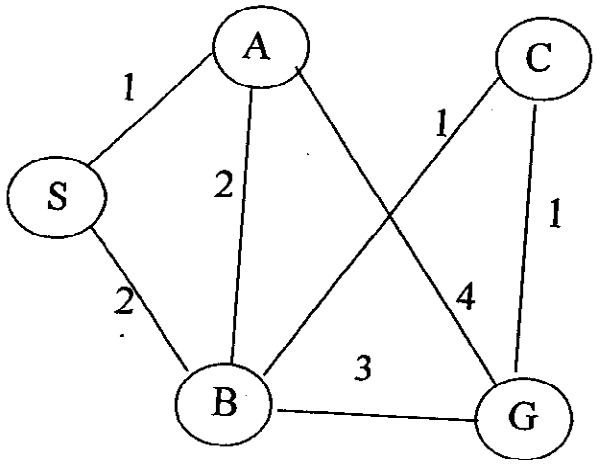

Suppose that the following heuristic is defined for the state space above:
h(S) = 3; h(A) = 2; h(B) = 2; h(G)=h(C) = 0.
Write down the sequence of nodes as they will be expanded by

i) Depth First Iterative Deepening — 2 mks
ii) A* Search — 2 mks
iii) Hill Climbing Search — 2 mks

(b) Say we define the evaluation function for a heuristic search problem (not the one in (a)) as $f(n) = (1 - w)g(n) + wh(n)$

where $g(n)$ is the cost of the best path found from the start state to state n, $h(n)$ is an admissible heuristic function that estimates the cost of a path from n to a goal state, and $0 \leq w \leq 1.0$. What search algorithm do you get when

(i) $w = 0.0$ (ii) $w = 0.5$ (iii) $w = 1.0$ 3 mks

State your reason.

2. Apply the α-β algorithm to the game tree below, where it is the *minimizer's* turn to play. Report the estimated values of the intermediate nodes and indicate the proper move of the minimizer. 4 mks

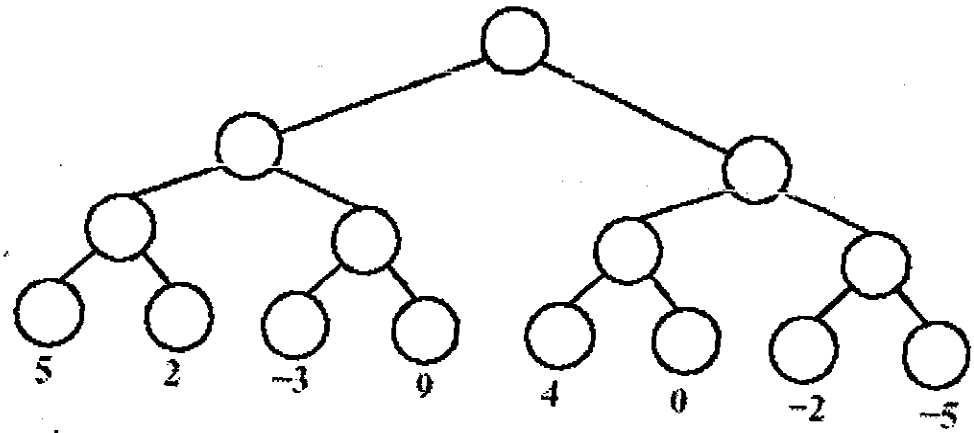

Indicate, by crossing out, *one* (1) unnecessary call to the static board evaluator. Explain why this call to the board evaluator is unnecessary. Give all the steps to get full credit.

3. a. Represent the following English sentences in predicate logic 3 mks

- All employees earning Rs 14000 or more per year pay *taxes*
- If an exam is easy, some students are happy.

(Define your predicates clearly for full credit)

b. Use binary resolution to solve the problem below. Given 5 mks

Given

i. $\forall X [p(X) \rightarrow q(X)]$

ii. $\forall X [p(X) \rightarrow [\exists Y w(Y)]]$

iii. $\forall X \forall Y [[q(X) \wedge w(Y)] \rightarrow s(X)]$

iv. p(mary)

Prove: s(mary)

BITS, PILANI – DUBAI
Academic City, Dubai
BE (Hons) CS IV Year - 1st Sem

Course No: EA C461
Course Title: Artificial Intelligence
Test -2 (Open book)

Date: 14 December 2008 **Time: 50 min** **Max Marks: 20**

Note: Answer all questions.

1. Draw an inheritance hierarchy that satisfies the "is-a" relationship and that includes these objects:
Square, Rectangle, Cube, Box, Shape, Two_Dimensional_Shape, Three Dimensional Shape, Circle, Ellipse, Point, Right Angled Triangle, Equilateral Triangle, Triangle, Sphere, and Cylinder.
3 mks

2. Consider the following cell phone domain. Having a charged battery (CB) enables a cell phone to be operational (PO). Having an operational cell phone, being in the coverage area (CA), and knowing the number to call (KN) enables the placing of a call (PC). Assume there are no other dependencies. All random variables are Boolean.

(a) Draw the Bayesian Network for this domain, including nodes, arcs and the form of each conditional probability table (CPT) at each node (values for probability can be left blank).
5 mks

(b) How many entries are in the full joint probability distribution table for this problem? Why?
2 mks

3. Consider a domain in which there are two attributes, X and Y, each with three possible values: 1, 2, 3. The concept to be learned is X >=Y. (E.g., if X=3 and Y=1, then the output is true.) Draw a decision tree that represents this concept (assume your training set contains all possible combinations of values).
5 mks

(b) Suppose we generate a very large training set from a given decision tree, TREE. Using this training set, we build a new decision tree using the decision tree learning algorithm. Is the new tree guaranteed to be identical to TREE? Briefly explain why or why not.
2 mks

4. A three-input, one-output parity detector outputs a 1 if the number of "1" inputs is even; otherwise it outputs a 0. Is this function linearly separable? Draw a three dimensional picture to support your answer.
3 mks

BITS, PILANI – DUBAI
Academic City, Dubai
BE (Hons) CS IV Year - 1st Sem

Course No: EA C461
Course Title: Artificial intelligence (AI)
Quiz -1

Date: 12 October 2008 **Time: 15 min** **Max Marks: 10**

ID NO:_____ **NAME:**_____

Note: Answer all questions. All questions carry equal marks

1. There are 3 towers (A, B, C), and 2 disks (small one and large one). The purpose of this problem is to move both disks from the tower A to tower C (as illustrated in the figure below), subject to following conditions:

- You can move only one disk at a time
- You cannot put the large disk on top of the small disk.

The possible state can be denoted as follows:
x:(y z) where x is the state number, y is the tower number for the large disk, and z is the tower number for the small disk. For example 3:(A C) implies that the large disk is on tower A and the small disk is on tower C in state 3.

If we use this notation, the following nine states are possible.
1:(A A), 2:(A B), 3:(A C), 4:(B A), 5:(B B), 6:(B C), 7:(C A), 8:(C B), 9:(C C)

a) Which ones are the initial state and goal state? 1 mks

b) Find a solution using depth-first search. Show the search tree and the order in which it is traversed. Assume that cycles are detected and eliminated by never expanding a node containing a state that is repeated on the path back to the root. 4 mks

2. a) Using the following tree, what is the final backed up value using MINIMAX algorithm. Show the correct backed up value at each node in order to get complete marks 2 mks

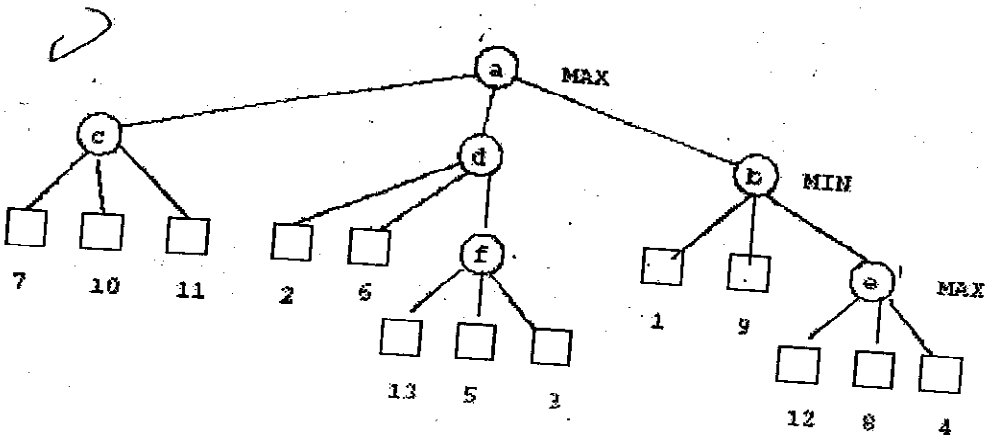

b) Indicate all the nodes (by giving their values below & also crossing on the graph) that are pruned using α-β pruning, assuming the search always goes LEFT to RIGHT. 3 mks

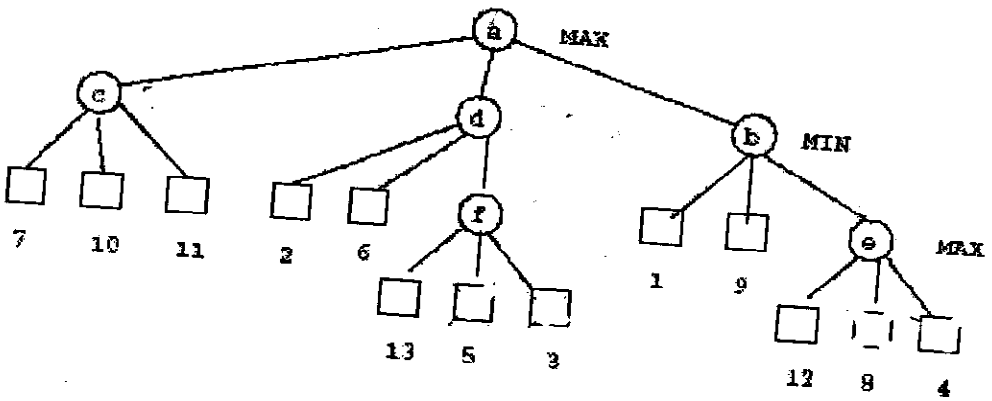

BITS, PILANI – DUBAI
Academic City, Dubai
BE (Hons) CS IV Year - 1st Sem

Course No: EA C461
Course Title: Artificial Intelligence
Quiz -2 A

Date: 29 October 2008 **Time: 15 min** **Max Marks: 10**

ID NO:_____ **NAME:**_____

I. Multiple choice (5 mks)

1. Breadth first search is uniform-cost search with ( )
   a. h(n) = 1
   b. g(n) = h(n)
   c. g(n) = 1
   d. g(n) < h(n)

2. Depth first search (in the worst case) takes ( )
   a. more space than Breadth First Search
   b. less time than Breadth First Search
   c. less space than Breadth First Search
   d. more time than Breadth First Search

3. Assume the following is typed to a newly-started LISP:
   ```
   (setq b '(8 (9)))
   (setq c '(list b c))
   ```
   What does the following return
   ```
   (cons b c)
   ```
   ( )
   a. (8 (9) b c)
   b. ((8 (9)) b c)
   c. (8 (9) (b c))
   d. None of the above

4. Which of the following first-order logic statements closely correspond to the English sentence: "Paul hit somebody with a stone" ( )
   a. $\exists x : hit(Paul, x) \wedge had(Paul, Stone)$
   b. $\exists x : hit(Paul, x) \wedge had(x, Stone)$
   c. $\exists x : hit(Paul, x, Stone)$
   d. $\exists x : hit(x) \wedge had(Paul, Stone)$

5. Use predicate P(x,y) to state that "x can praise y". The translation of the following sentence into predicate logic. "There is no one who can praise everybody" ( )
   a. $\neg\exists x\ \forall y\ P(x,y)$
   b. $\neg\forall x\ \exists y\ P(x,y)$
   c. $\neg\forall x\ \exists y\ P(y,x)$
   d. None of the above

II. Short Answer: (5mks)

Given

~A

~A => ~B Λ ~C

~B => ~D Λ ~E

~C => D V E V F

prove by resolution refutation the fact F (Show all the steps to get full credit)

BITS, PILANI – DUBAI
Academic City, Dubai
BE (Hons) CS IV Year - 1st Sem

Course No: EA C461
Course Title: Artificial Intelligence
Quiz -3 A

Date: 19 November 2008 **Time: 15 min** **Max Marks: 10**

ID NO:_____ **NAME:**_____

I. Multiple choice (4 mks)
Assume that we have evaluated the following lisp-expression:
(setq a '(1 2 3))
What do each of the following return ?

1. (list 'a 1 a)
a. ((1 2 3) 1 (1 2 3)) b. ((123) 1 a) c. (a 1 (1 2 3)) ( ) d. (a 1 a)

2. (cons a '(a b))
a. ((1 2 3) (a b)) b. ((1 2 3) a b) c. (1 2 3 (a b)) ( ) d. (1 2 3 a b)

3. (append (cddr a) (list 1 2))
a. (3 (1 2)) b. (3 1 2) c. ((3) 1 2) ( ) d. ((3) (1 2))

4. '(cons 3 (list a))
a. (cons 3 (a)) b. (cons 3 (list a)) c. (3 (list a)) ( ) d. (3 (a))

II. Short Answers (6 mks)
Assume the following is typed to a newly-started LISP:
(setq a (+ 3 5))
(setq b '(8 (9)))
(setq c '(cons b c))
What does each of the following return ?

5. a _____

6. (- 31 a) _____

7. (cons b c) _____

8. (list b c) _____

9. (append b c) _____

10. (list 'a a) _____