IITJ, III Year (CSE), II Semester
End Semester Examination, 2016(with solutions)
CS323:Artificial Intelligence

Duration: 3 Hours Max. Marks 100

- Attempt all questions.
- Every answer should be detailed, and supported with sufficient logic for justification.

---

1. Consider the 3-puzzle problem, which is a simpler version of the 8-puzzle where the board is 2×2 and there are three tiles, numbered 1, 2, and 3, and *blank*. There are four operators, which move the *blank* tile **up**, **down**, **left**, and **right**. The **start** and **goal** states are given in figure 1. Show, how the path to the goal can be found using:

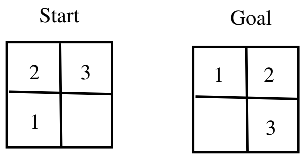

Figure 1: State Space search.

   (a) Breath-first search [5]

   (b) Depth-first search [5]

   (c) A^* search with the heuristic being sum of the number of moves and the number of misplaced tiles. [10]

   Assume that there is no possibility to remember states that have been visited earlier. Also, use the given operators in the given order unless the search method defines otherwise. Label each visited node with a number indicating the order in which they are visited. If a search method does not find a solution, explain why this happened.

Ans. These are shown as figure 2 for breadth first search, as figure 3 for depth first search, and as figure 4 for A*-search.

2. Explain what algorithms or heuristics are suitable for solving constraint satisfaction problems under the following situations. Justify your answers.

   (a) The problem is so tightly constrained that it is highly unlikely that solutions exist. [5]

   Ans. Should cover: Propagate and use strategy of FC (forward checking). Also, use some heuristics, like MVH (minimum value heuristics), degree heuristics with largest number of constraints first, local search, minimum conflicts. Along with formal representation of above as well data structure for management of constraints, branch and bound.

   (b) The domain sizes vary significantly: some variables have very large domains (over 1,000 values) and some have very small domains (with fewer than 10 values). [5]

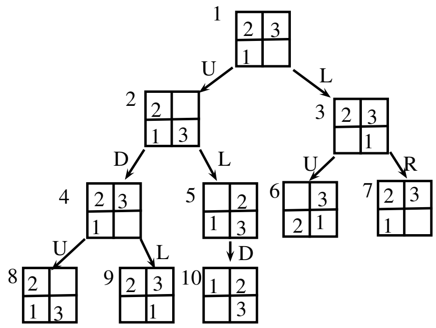

Figure 2: Breadth-first search.

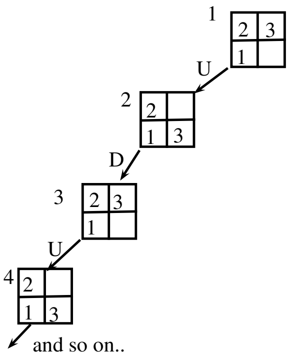

Figure 3: Depth first search.

Ans. Should cover: Synthesizing of constraints, path inconsistency, node inconsistency, augmented networks, constraint propagation, preclusion, branch merging, search rearrangement, propagation of constraints, least constraints value heuristics, formal methods for representation, as well data structures required.

(c) The set of variables and set of domains are handled by a computer say, M. Each constraint is handled by a networked computer, say N. Traffic in the networks is slow. To check a particular constraint, computer M sends a message to computer N through the network, which in turn will send a message back to indicate whether the constraint is satisfied or violated. [10]

Ans. Should cover: Generate and test, backtracking algorithm, propagate and use strategy of FC (forward checking). Formal methods of generate and test, data structures required.

3. Two firms Alpha and Beta serve the same market. They have constant average costs of $2 per unit. The firms can choose either a high price ($10) or a low price ($5) per unit for their product. When both firms set a high price, total demand is $10,000$ units, which is split evenly between the two firms. When both set a low price, total demand is $18,000$ units, which is again split evenly. If one firm sets a low price and the other a high price, the low priced firm sells $15,000$ units, while the high priced firm only $2,000$ units. Analyze the pricing decisions of the two firms as a non-cooperative game.

(a) In the scenarios mentioned above, form representation, construct the pay-off matrix, where the elements of each cell of the matrix are the two firms' profits. [10]

Ans. a) Payoff alpha-high = 10000/2*(10-2) =$40,000.

Payoff beta-high = 10000/2*(10-2) =$40,000.

b) Payoff alpha-low = 18000/2*(5-2) =$27,000.

Payoff beta-low = 18000/2*(5-2) =$27,000.

c) Payoff alpha-low = 15000*(5-2) =$45,000.

Payoff beta-high = 2000*(10-2) =$16,000.

d) Payoff beta-low = 15000*(5-2) =$45,000.

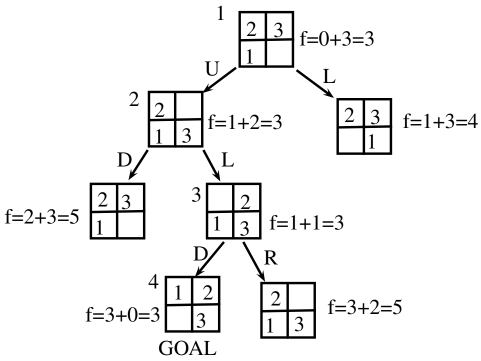

Figure 4: A* search.

Payoff alpha-high = 2000*(10-2) =$16,000.

The payoff matrix is shown in table 1.

Table 1: Payoff Matrix.

		Beta	Beta
		High	Low
Alpha	High	(40k, 40k)	(16k, 45k)
Alpha	Low	(45k, 16k)	(27k, 27k)

(b) Derive the equilibrium set of strategies. [5]

Ans. The equilibrium state is shown in figure 5. Say, alpha choses the max and beta as min. Hence, at root, alpha will choose path A, to have it from the worst, i.e., lowest 16k. Having completed the path A, the beta will choose the path D (note that it is not a cooperative game). Similar will be the situation, if beta player is maximizing and alpha is minimizing. Hence the search path is AD, and equilibrium state is $(27k, 27k)$ dollars.

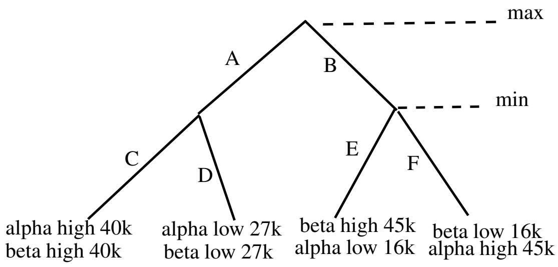

Figure 5: game-tree.

4. Apply alpha-beta search (from left to right) to the game tree in figure 6. Show the backed-up value of each node. Mark with an X any branches that are not searched.

   (a) Identify these as alpha / beta cut-offs, mark the best moves with an arrow from root node, [5]

   (b) compute the time complexity of search with worst case branching factor of 3 and height of tree as h. [5]

   Ans. See the lecture notes.

Ans. The cutoffs are shown in figure 7.

5. Assume that you have three operators:

   O_1 : Precondition: a; effect: $\neg a \wedge b$

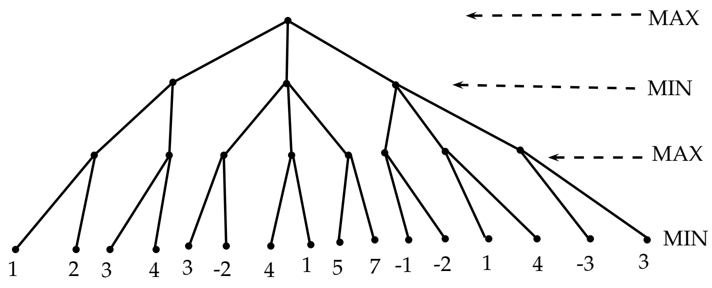

Figure 6: game-tree.

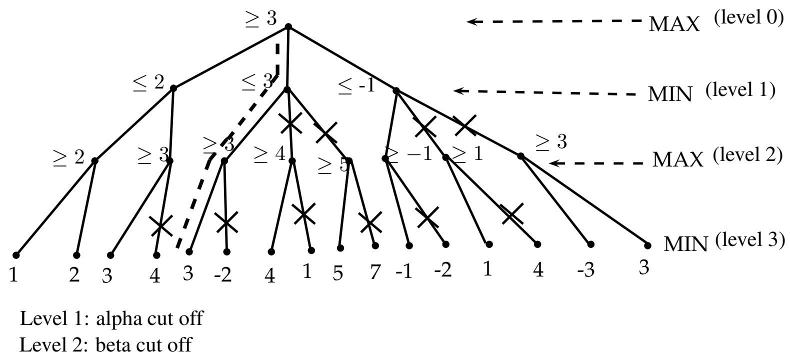

Figure 7: game-tree with alpha-beta cuts.

O_2 : Precondition: $a \wedge c$; effect: $\neg a \wedge b \wedge \neg c$

O_3 : Precondition: $b \wedge c$; effect: $\neg c \wedge d$

Show the first three layers (proposition, action and proposition) of the **graph plan** when initial state is $a \wedge c$ (a and c both are true). Include the mutual exclusions and justify each of them. [10]

Ans. The figure 8 shows the graph plan as solution of this planning problem.

On the first level only the current state ($a \wedge c$) is shown. On the second level (action level 1), all possible actions that are possible to perform are drawn. This also produces the set of possible literals for the third level. Note that the preconditions of O_3 are not satisfied and thus the operator is not in the graph at the first action level. The square show the option of not performing any action for each literal present on previous level.

Now we have the graph, and we notice that some actions cannot be performed at the same level, these are mutually exclusive. The $a, \neg a$ is logically inconsistent hence mutex, also $c, \neg c$ is also logically inconsistent hence mutex, and b, c is inconsistent effect (a). Hence mutex.

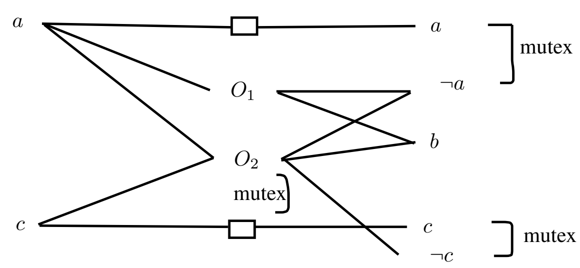

Figure 8: Graph Plan.

6. Represent the following statements using semantic networks:

(a) "John tells his students a lot of useful things." [5]

(b) "Andrea tells John's students an enormous number of useful things." [5]

Suppose you wanted to build an AI system that was able to work out "who tells John's students the greatest number of useful things." How could you do that? [5]

Ans. Standard semantic network representation.

7. Is the following formula a default rule? Justify for yes/no. If it is no, what else is required to make it a default rule. [2]

$$\frac{bird(X) : flies(X)}{flies(X)}$$

Ans. The clauses are required to be ground clauses.

Compute the default extensions of following theories $T = (M, D)$, assuming the closed world representation.

(a) $M = \{a\}, D = \{\frac{:\neg c}{d}, \frac{:\neg d}{e}\}$ [4]

Ans. For the formula to be $\frac{A:B}{C}$, A is null (ϕ, empty set) in all the cases. And, $M \vdash \phi$ holds true in all these cases. The two extensions are $\{a, d\}$ and $\{a, e, d\}$.

(b) $M = \{p \wedge q\}, D = \{\frac{:\neg q}{b}, \frac{:\neg p}{q}\}$ [4]

Ans. There is no extension in this.