CS 188 Fall 2021 | Introduction to Artificial Intelligence | Exam Prep 2

Q1. Searching with Heuristics

Consider the A* searching process on the connected undirected graph, with starting node S and the goal node G. Suppose the cost for each connection edge is **always positive**. We define $h^*(X)$ as the shortest (optimal) distance to G from a node X.

Answer Questions (a), (b) and (c). You may want to solve Questions (a) and (b) at the same time.

(a) Suppose h is an **admissible** heuristic, and we conduct A* **tree search** using heuristic h' and finally find a solution. Let C be the cost of the found path (directed by h', defined in part (a)) from S to G

(i) Choose **one best** answer for each condition below.

1. If $h'(X) = \frac{1}{2}h(X)$ for all Node X, then ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$
2. If $h'(X) = \frac{h(X)+h^*(X)}{2}$ for all Node X, then ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$
3. If $h'(X) = h(X) + h^*(X)$ for all Node X, then ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$
4. If we define the set $K(X)$ for a node X as all its neighbor nodes Y satisfying $h^*(X) > h^*(Y)$, and the following always holds

$$h'(X) \leq \begin{cases} \min_{Y \in K(X)} h'(Y) - h(Y) + h(X) & \text{if } K(X) \neq \emptyset \\ h(X) & \text{if } K(X) = \emptyset \end{cases}$$

then, ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$

5. If K is the same as above, we have

$$h'(X) = \begin{cases} \min_{Y \in K(X)} h(Y) + cost(X,Y) & \text{if } K(X) \neq \emptyset \\ h(X) & \text{if } K(X) = \emptyset \end{cases}$$

where $cost(X,Y)$ is the cost of the edge connecting X and Y,
then, ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$

6. If $h'(X) = \min_{Y \in K(X)+\{X\}} h(Y)$ (K is the same as above), ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$

(ii) In which of the conditions above, h' is still **admissible** and for sure to dominate h? Check all that apply. Remember we say h_1 dominates h_2 when $h_1(X) \geq h_2(X)$ holds for all X. □ 1 □ 2 □ 3 □ 4 □ 5 □ 6

(b) Suppose h is a **consistent** heuristic, and we conduct A* **graph search** using heuristic h' and finally find a solution.

(i) Answer exactly the same questions for each conditions in Question (a)(i).

1. ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$
2. ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$
3. ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$
4. ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$
5. ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$
6. ○ $C = h^*(S)$ ○ $C > h^*(S)$ ○ $C \geq h^*(S)$

(ii) In which of the conditions above, h' is still **consistent** and for sure to dominate h? Check all that apply.
□ 1 □ 2 □ 3 □ 4 □ 5 □ 6

(c) Suppose h is an **admissible** heuristic, and we conduct A* **tree search** using heuristic h' and finally find a solution.

If $\epsilon > 0$, and X_0 is a node in the graph, and h' is a heuristic such that

$$h'(X) = \begin{cases} h(X) & \text{if } X = X_0 \\ h(X) + \epsilon & \text{otherwise} \end{cases}$$

- Alice claims h' can be inadmissible, and hence $\underline{C = h^*(S) \text{ does not always hold}}$.
- Bob instead thinks the node expansion order directed by h' is the same as the heuristic h'', where

$$h''(X) = \begin{cases} h(X) - \epsilon & \text{if } X = X_0 \\ h(X) & \text{if otherwise} \end{cases}$$

  Since h'' is admissible and will lead to $C = h^*(S)$, and so does h'. Hence, $\underline{C = h^*(S) \text{ always holds}}$.

The two conclusions (underlined) apparently contradict with each other, and **only exactly one of them are correct and the other is wrong**. Choose the **best** explanation from below - which student's conclusion is wrong, and why are they wrong?

◯ Alice's conclusion is wrong, because the heuristic h' is always admissible.

◯ Alice's conclusion is wrong, because an inadmissible heuristics does not necessarily always lead to the failure of the optimality when conducting A* tree search.

◯ Alice's conclusion is wrong, because of another reason that is not listed above.

◯ Bob's conclusion is wrong, because the node visiting expansion ordering of h'' during searching might not be the same as h'.

◯ Bob's conclusion is wrong, because the heuristic h'' might lead to an incomplete search, regardless of its optimally property.

◯ Bob's conclusion is wrong, because of another reason that is not listed above.

Q2. Iterative Deepening Search

Pacman is performing search in a maze again! The search graph has a branching factor of b, a solution of depth d, a maximum depth of m, and edge costs that may not be integers. Although he knows breadth first search returns the solution with the smallest depth, it takes up too much space, so he decides to try using iterative deepening. As a reminder, in standard depth-first iterative deepening we start by performing a depth first search terminated at a maximum depth of one. If no solution is found, we start over and perform a depth first search to depth two and so on. This way we obtain the shallowest solution, but use only O(bd) space.

But Pacman decides to use a variant of iterative deepening called **iterative deepening A***, where instead of limiting the depth-first search by depth as in standard iterative deepening search, we can limit the depth-first search by the f value as defined in A* search. As a reminder $f[node] = g[node] + h[node]$ where $g[node]$ is the cost of the path from the start state and $h[node]$ is a heuristic value estimating the cost to the closest goal state.

In this question, all searches are tree searches and **not** graph searches.

(a) Complete the pseudocode outlining how to perform iterative deepening A* by choosing the option from the next page that fills in each of these blanks. Iterative deepening A* should return the solution with the lowest cost when given a consistent heuristic. Note that cutoff is a boolean and new-limit is a number.

```
function ITERATIVE-DEEPENING-TREE-SEARCH(problem)
    start-node ← MAKE-NODE(INITIAL-STATE[problem])
    limit ← f[start-node]
    loop
        fringe ← MAKE-STACK(start-node)
        new-limit ← [ (i) ]
        cutoff ← [ (ii) ]
        while fringe is not empty do
            node ← REMOVE-FRONT(fringe)
            if GOAL-TEST(problem, STATE[node]) then
                return node
            end if
            for child-node in EXPAND(STATE[node], problem) do
                if f[child-node] ≤ limit then
                    fringe ← INSERT(child-node, fringe)
                    new-limit ← [ (iii) ]
                    cutoff ← [ (iv) ]
                else
                    new-limit ← [ (v) ]
                    cutoff ← [ (vi) ]
                end if
            end for
        end while
        if not cutoff then
            return failure
        end if
        limit ← [ (vii) ]
    end loop
end function
```

$\mathbf{A_1}$	$-\infty$	$\mathbf{A_2}$	0	$\mathbf{A_3}$	∞	$\mathbf{A_4}$	*limit*
$\mathbf{B_1}$	True	$\mathbf{B_2}$	False	$\mathbf{B_3}$	*cutoff*	$\mathbf{B_4}$	not *cutoff*
$\mathbf{C_1}$	*new-limit*	$\mathbf{C_2}$	*new-limit* + 1	$\mathbf{C_3}$	*new-limit* + f[*node*]	$\mathbf{C_4}$	*new-limit* + f[*child-node*]
$\mathbf{C_5}$	MIN(*new-limit*, f[*node*])	$\mathbf{C_6}$	MIN(*new-limit*, f[*child-node*])	$\mathbf{C_7}$	MAX(*new-limit*, f[*node*])	$\mathbf{C_8}$	MAX(*new-limit*, f[*child-node*])

(i) ◯ $\mathbf{A_1}$ ◯ $\mathbf{A_2}$ ◯ $\mathbf{A_3}$ ◯ $\mathbf{A_4}$

(ii) ◯ $\mathbf{B_1}$ ◯ $\mathbf{B_2}$ ◯ $\mathbf{B_3}$ ◯ $\mathbf{B_4}$

(iii) ◯ $\mathbf{C_1}$ ◯ $\mathbf{C_2}$ ◯ $\mathbf{C_3}$ ◯ $\mathbf{C_4}$
◯ $\mathbf{C_5}$ ◯ $\mathbf{C_6}$ ◯ $\mathbf{C_7}$ ◯ $\mathbf{C_8}$

(iv) ◯ $\mathbf{B_1}$ ◯ $\mathbf{B_2}$ ◯ $\mathbf{B_3}$ ◯ $\mathbf{B_4}$

(v) ◯ $\mathbf{C_1}$ ◯ $\mathbf{C_2}$ ◯ $\mathbf{C_3}$ ◯ $\mathbf{C_4}$
◯ $\mathbf{C_5}$ ◯ $\mathbf{C_6}$ ◯ $\mathbf{C_7}$ ◯ $\mathbf{C_8}$

(vi) ◯ $\mathbf{B_1}$ ◯ $\mathbf{B_2}$ ◯ $\mathbf{B_3}$ ◯ $\mathbf{B_4}$

(vii) ◯ $\mathbf{C_1}$ ◯ $\mathbf{C_2}$ ◯ $\mathbf{C_3}$ ◯ $\mathbf{C_4}$
◯ $\mathbf{C_5}$ ◯ $\mathbf{C_6}$ ◯ $\mathbf{C_7}$ ◯ $\mathbf{C_8}$

(b) Assuming there are no ties in f value between nodes, which of the following statements about the number of nodes that iterative deepening A* expands is True? If the same node is expanded multiple times, count all of the times that it is expanded. If none of the options are correct, mark None of the above.

◯ The number of times that iterative deepening A* expands a node is greater than or equal to the number of times A* will expand a node.

◯ The number of times that iterative deepening A* expands a node is less than or equal to the number of times A* will expand a node.

◯ We don't know if the number of times iterative deepening A* expands a node is more or less than the number of times A* will expand a node.

◯ None of the above

Q3. CSPs

In this question, you are trying to find a four-digit number satisfying the following conditions:

1. the number is odd,
2. the number only contains the digits 1, 2, 3, 4, and 5,
3. each digit (except the leftmost) is strictly larger than the digit to its left.

(a) CSPs

We will model this as a CSP where the variables are the four digits of our number, and the domains are the five digits we can choose from. The last variable only has 1, 3, and 5 in its domain since the number must be odd. The constraints are defined to reflect the third condition above. Thus before we start executing any algorithms, the domains are

1 2 3 4 5	1 2 3 4 5	1 2 3 4 5	1 3 5

(i) Before assigning anything, enforce arc consistency. Write the values remaining in the domain of each variable after arc consistency is enforced.

(ii) With the domains you wrote in the previous part, which variable will the MRV (Minimum Remaining Value) heuristic choose to assign a value to first? If there is a tie, choose the leftmost variable.

○ The first digit (leftmost)
○ The second digit
○ The third digit
○ The fourth digit (rightmost)

(iii) Now suppose we assign to the leftmost digit first. Assuming we will continue filtering by enforcing arc consistency, which value will LCV (Least Constraining Value) choose to assign to the leftmost digit? **Break ties from large (5) to small (1).**

☐ 1
☐ 2
☐ 3
☐ 4
☐ 5

(iv) Now suppose we are running min-conflicts to try to solve this CSP. If we start with the number 1332, what will our number be after one interation of min-conflicts? Break variable selection ties from left to right, and **break value selection ties from small (1) to large (5)**.

____

(b) The following questions are completely unrelated to the above parts. Assume for these following questions, there are only binary constraints unless otherwise specified.

(i) [*true* or *false*] When enforcing arc consistency in a CSP, the set of values which remain when the algorithm terminates does not depend on the order in which arcs are processed from the queue.

(ii) [*true* or *false*] Once arc consistency is enforced as a pre-processing step, forward checking can be used during backtracking search to maintain arc consistency for all variables.

(iii) In a general CSP with n variables, each taking d possible values, what is the worst case time complexity of enforcing arc consistency using the AC-3 method discussed in class?

○ 0 ○ $O(1)$ ○ $O(nd^2)$ ○ $O(n^2d^3)$ ○ $O(d^n)$ ○ ∞

(iv) In a general CSP with n variables, each taking d possible values, what is the maximum number of times a backtracking search algorithm might have to backtrack (i.e. the number of the times it generates an assignment, partial or complete, that violates the constraints) before finding a solution or concluding that none exists?

○ 0 ○ $O(1)$ ○ $O(nd^2)$ ○ $O(n^2d^3)$ ○ $O(d^n)$ ○ ∞

(v) What is the maximum number of times a backtracking search algorithm might have to backtrack in a general CSP, if it is running arc consistency and applying the MRV and LCV heuristics?

○ 0 ○ $O(1)$ ○ $O(nd^2)$ ○ $O(n^2d^3)$ ○ $O(d^n)$ ○ ∞