

# CSE 415 Autumn 2022 Assignment 4

Last name: \_\_\_\_\_ First name: \_\_\_\_\_

Due Thursday night November 3 (was Wed. Nov. 2 prior to a correction to the “optimal policy” column of the table in Question 5) via Gradescope at 11:59 PM. You may turn in either of the following types of PDFs: (1) Scans of these pages that include your answers (handwriting is OK, if it’s clear), or (2) Documents you create with the answers, saved as PDFs. When you upload to GradeScope, you’ll be prompted to identify where in your document your answer to each question lies.

Do the following five exercises. These are intended to take 20-25 minutes each if you know how to do them. Each is worth 20 points. If any corrections have to be made to this assignment, these will be posted in ED.

This is an individual-work assignment. Do not collaborate on this assignment.

Prepare your answers in a neat, easy-to-read PDF. Our grading rubric will be set up such that when a question is not easily readable or not correctly tagged or with pages repeated or out of order, then points will be deducted. However, if all answers are clearly presented, in proper order, and tagged correctly when submitted to Gradescope, we will award a 5-point bonus.

If you choose to choose to typeset your answers in Latex using the template file for this document, please put your answers in [blue](#) while leaving the original text black.

---

# 1 Blind Search with the Towers of Hanoi

The 2-disk version of the Towers of Hanoi is a trivial puzzle for humans and machines alike. However, it's a nice and simple context for comparing different algorithms.

Several aspects of this problem will come up again in Assignment 5, and so this problem will not only help you get more familiar with certain details of the search algorithms, but it will provide some insight into the the Towers of Hanoi problem space.

Let us assume that the problem is formulated with the following state representation and operators.

*Initial state:*

Left: 1,2

Middle:

Right:

*Goal state:*

Left:

Middle:

Right: 1,2

Operator  $\phi_0$ : "Move a disk from Left to Middle."

Operator  $\phi_1$ : "Move a disk from Left to Right."

Operator  $\phi_2$ : "Move a disk from Middle to Left."


Operator  $\phi_3$ : "Move a disk from Middle to Right."

Operator  $\phi_4$ : "Move a disk from Right to Left."


Operator  $\phi_5$ : "Move a disk from Right to Middle."

Hand simulate DFS (Depth-First Search, BFS (Breadth-First Search) and IDDFS (Iterative-Deepening Depth-First Search) on this problem, in order to determine the state visitation orderings and compare them.

The problem-space graph for this formulation can be laid out as in the diagrams for sub-questions a-c. In addition, this figure shows how the states in the problem-space correspond to the position of the disks:



Note that the operators always take a specific direction, as shown in this diagram:



Use the copies of the problem space graph below to show the progress of each search algorithm. When doing IDDFS, you'll use a separate copy of the graph for each outer iteration.

For each algorithm stop when the goal node (g) is selected as the current state.


Number the nodes as they are visited (i.e., as they become the current state). The initial state should get numbered 1 when each algorithm starts. However, during IDDFS, it should get multiple numbers since it will be visited multiple times.

In order to get the correct numberings, it is very important to follow the pseudocode in the DFS, and BFS algorithms, and to generate the successors by using the operators in their given order:  $\phi_0, \phi_1, \phi_2, \phi_3, \phi_4, \phi_5$ .


Each time a new node is reached in the search (i.e., put onto the OPEN list), draw an arrow to it from its parent. Note that this does not apply to the initial state, since it has no parent. When doing this for IDDFS, you'll retain (i.e., copy from one diagram to the next) these

arrows from one outer iteration to the next. At this end of IDDFS, you should have an optimal path from the initial state to the goal state, which can be recovered by backtracing from the goal.

- (a) (0 points) Hand-simulate DFS and put the node visitation order on the graph. Note that the answers to this part are done for you as an example.




- (b) (8 points) Hand-simulate BFS and put the node visitation order on the graph. As in the given example, also show the moves using arrows and operator identifiers ( $\phi_0$ , etc.).




- (c) (12 points) Hand-simulate IDDFS and put the node visitation order onto the four graphs below. Use one graph copy for each iteration of IDDFS. The first graph should have only one node (for the initial state) visited. on the graph. The second graph's initial state should have visitation number 2 (since this state is visited again). Note that within one iteration of IDDFS, some node(s) may be reached multiple times along different paths from the initial state. This is OK, and each such repeat visit should be counted as a node visitation and shown in your results. However, there should be no more than one visitation of any given node along the current path between the initial state and current state. That is, the search path must never be allowed to loop back on itself.

First and second iterations:

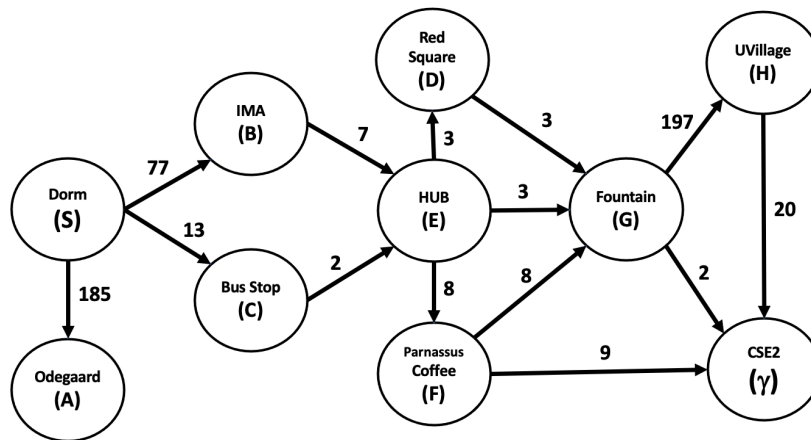


Third and fourth iterations:



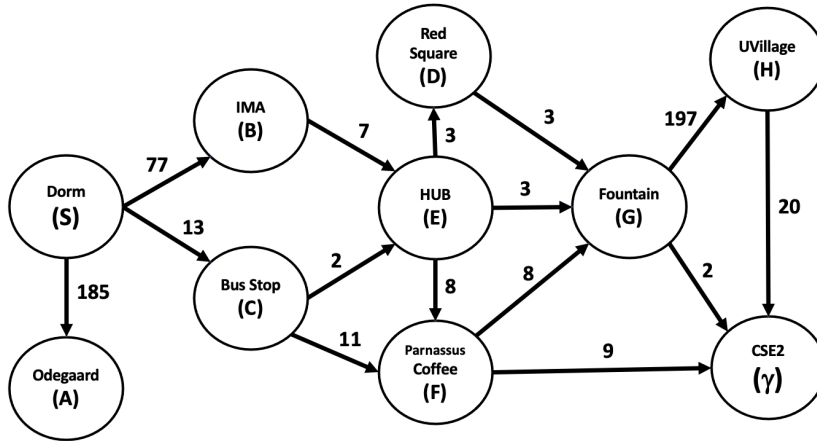

## 2 Heuristic Search

- (a) (5 points) Consider the following statement: *The best heuristic is always the one that gives you the estimate closest to the true cost.* Is this always true? Explain why you agree or disagree with the statement. **Hint:** Consider what “closest” does/doesn’t guarantee and consider how “best” might be determined or defined.



state (s)	$s_0$	A	B	C	D	E	F	G	H	$\gamma$
heuristic $h_1(s)$	13	25	11	5	5	4	10	2	15	0
heuristic $h_2(s)$	14	50	12	6	5	5	8	2	18	0
heuristic $h_3(s)$	14	5	5	7	4	1	7	1	19	0

- (b) (5 points) Which heuristics ( $h_1, h_2, h_3$ ) shown above are admissible?
- (c) (5 points) Which heuristics ( $h_1, h_2, h_3$ ) shown above are consistent?
- (d) (5 points) Which of the 3 heuristics above would you select as the best one to use with A\* search? Why? Refer to both consistency and admissibility in your justification.



state (s)	$s_0$	A	B	C	D	E	F	G	H	$\gamma$
heuristic $h_4(s)$	18	50	11	6	4	4	8	1	18	0

- (e) (5 points) Referring back to the graph (**NOTE:** new edge added – coffee’s important!), determine the path that would be computed by an A\* search, given the heuristics provided above. As you search the space, complete the table below, indicating which nodes are on the OPEN and CLOSED lists, along with their  $f$  values:


Action(s)	OPEN	CLOSED
Starting A* search	$[s_0, 18]$	empty
take $s_0$ off OPEN & add successors to OPEN	<b>NOTE: ADD THE</b> $[\text{successors \& } f \text{ values}]$	$[s_0, 18]$

continue through to goal...

### 3 Adversarial Search

Minimax game-tree search finds the best move under the assumption that both players play rationally to respectively maximize and minimize the utility of the same evaluation function. (It can also do well when those assumptions are relaxed.) The Minimax search, however, typically requires a lot of computation time, especially with a deep and widely branching search tree. Alpha-beta pruning is a method for speeding up Minimax search by skipping any subtrees from the search tree that will not contribute to the outcome. Despite that, Alpha-beta pruning's success is dependent upon the order in which we visit the states. A better order to visit the nodes could result in large difference in performance.

For the following questions, note that the example we are using has a maximizing node at the root, which means our goal is to determine which choice of move at A offers the maximum value.



- (a) (7 points) Apply straight minimax search, depth-first, left-to-right. Show the resulting values at each internal node. Then fill in the table. The total number of nodes visited should include counts for the root node and the leaf nodes.

Number of nodes visited:	
--------------------------	--



- (b) (3 points) Apply alpha-beta pruning, left-to-right, on the given tree. Mark where cutoffs occur in the tree above, and fill out the tables. Note: Use the **final** resulting alpha-beta values for the second table. You can also mark these values in the graph if you prefer. You do **NOT** need to fill in the resulting values at each internal node for this part):

Number of nodes visited:	
Count of cutoffs:	


Node	A	B	C	D	E	F	G	H	I	J
alpha value										
beta value										

- (c) (10 points) Apply alpha-beta pruning, using an auxiliary static evaluation function  $f_2(s)$ . You are still going depth-first in some sense, but use the auxiliary static evaluation function to determine the order for visiting the children nodes. Starting from the root node, sort each of its children from high to low using  $f_2$ . Then visit them following that order. Similarly, do this for all the non-leaf nodes you visit, except that when visiting a minimizing node, sort its children from low to high. You will still look for cutoffs when applicable.

For this problem, Use  $f_2$  as given in the tables. (In practice, we might invest a lot of resources into computing the auxiliary static evaluation function that would result in a good ordering prior to alpha-beta pruning.)

Node $s$ :	A	B	C	D	E	F	G	H	I	J
$f_2(s)$ :	0	1	3	2	2	1	5	6	3	4
Node $s$ :	K	L	M	N	O	P	Q	R	S	T
$f_2(s)$ :	3	2	1	9	8	10	4	5	6	7

Show the resulting values at each internal node (leave it blank if pruned), mark where cutoffs occur on the graph below, and fill out the tables (leave the cell blank if pruned):



Number of nodes visited:	
Count of cutoffs:	


Node	A	B	C	D	E	F	G	H	I	J
alpha value										
beta value										

## 4 Markov Decision Processes

Consider the following problem that is taking place in various instances of a Grid-World MDP. The shaded box is walled off and is not one of the possible states. In all states, the agent has available actions  $\uparrow$ ,  $\downarrow$ ,  $\leftarrow$ ,  $\rightarrow$ . Performing an action that would transition to an invalid state (outside the grid or into a wall) results in the agent remaining in its original state. In states with an arrow coming out, the agent has an *additional* action EXIT. In the event that the EXIT action is taken, the agent receives the labeled reward and ends the game in the terminal state  $T$ . Unless otherwise stated, all other transitions receive no reward, and all transitions are deterministic. (Thus, this MDP is rather different from the Grid-World example shown in lecture, even though they share the same state space.)

For **all** parts of the problem, assume that value iteration begins with all states initialized to zero, i.e.,  $\forall s V_0(s) = 0$ . **Let the discount factor be  $\gamma = 0.5$  for all the following parts.**

Suppose that we are performing Value Iteration on the Grid-World MDP below:



- (a) Fill in the “optimal” values for A and B in the given boxes. (I.e., determine the expected total discounted rewards when starting in A (or B) and following an optimal policy.)


$$V^*(A) : \quad \boxed{\phantom{000000}} \quad V^*(B) : \quad \boxed{\phantom{000000}}$$

- (b) After how many iterations  $k$  will we have  $V_k(s) = V^*(s)$  for all states  $s$ ? If it never occurs, write “never”. Write your answer in the given box.

- (c) Suppose that we wanted to re-design the reward function. For which of the following new reward functions would the optimal policy **remain unchanged**? Let  $R(s, a, s')$  be the original reward function.

- $R_1(s, a, s') = 10 \cdot R(s, a, s')$
- $R_2(s, a, s') = R(s, a, s') + 1$
- $R_3(s, a, s') = R(s, a, s')^2$
- $R_4(s, a, s') = -5$
- None

(d) For the following problem, we add a new state in which we can take the **EXIT** action with a reward of  $+x$ .



(i) For what values of  $x$  is it *guaranteed* that our optimal policy  $\pi^*$  has  $\pi^*(C) = \leftarrow$  (i.e. going left from C)? Write  $\infty$  and  $-\infty$  if there is no upper or lower bound, respectively. Write the upper and lower bounds in each respective box.

< x <

(ii) For what values of  $x$  does value iteration take the **minimum** number of iterations  $k$  to converge to  $V^*$  for all states? Write  $\infty$  and  $-\infty$  if there is no upper or lower bound, respectively. Write the upper and lower bounds in each respective box.

$\leq x \leq$

(iii) Fill the box with value  $k$ , the **minimum** number of iterations until  $V_k$  has converged to  $V^*$  for all states.

## 5 Computing MDP State Values and Q-Values

Recently, Mike has been working on building an intelligent agent to help a friend solve a problem that can be modeled using an MDP. In this environment, there are 3 possible states  $S = \{s_1, s_2, s_3\}$  and at each state the agent always has 2 available actions  $A = \{f, g\}$ . Applying any action  $a$  from any state  $s$  has a (possibly non-zero) probability  $T(s, a, s')$  of moving the agent to one of the *other two* states but will never result in the agent staying at the original state. The rewards for this environment are only dependent on the original state and action taken ( $\forall s' \in S, R(s, a, s') = R(s, a)$ ), not where the agent ended up.

- (a) (2 points) Write down the problem-specific Bellman equations for each of the 3 states ( $V(s) = ?$ ) in this particular MDP. (Use the symbols of the specific states, e.g.,  $s_1$ , and actions, e.g.,  $f$ .)

- (b) (18 points) One fateful day, while Mike was running a VI-based MDP solver on this problem, a mistake in specifying arguments caused the file that recorded the transition probability table  $T(s, a, s')$  to be overwritten with the output solution. Now, Mike has the solution  $V^*(s)$  and optimal policy  $\pi^*(s)$  but has lost the transition probabilities for the problem.

$s$	$a$	$R(s, a)$	$V^*(s)$	$\pi^*(s)$
$s_1$	$f$	-3	-1.1	$f$
$s_1$	$g$	-2.5	-1.1	$f$
$s_2$	$f$	2	2	$g$
$s_2$	$g$	2	2	$g$
$s_3$	$f$	1.41	1.1	$f$
$s_3$	$g$	1.2	1.1	$f$

After looking at the command line history and noting that a discount of  $\gamma = 1$  was specified, Mike muses that it may be possible to recover some parts of the transition probability table  $T(s, a, s')$ . Using the information above, fill in the values in the table below that you recover, rounded to 3 decimal places. HINT: What must transition

probabilities out of a given state add up to? Also, how do you relate  $V^*(s)$  to the transition probabilities, rewards, and other state values?

(Note: the rightmost column of the table was corrected on Nov. 1 from f,g,f,g,f,g to f,f,g,g,f,f.)

$s$	$a$	$T(s, a, s_1)$	$T(s, a, s_2)$	$T(s, a, s_3)$
$s_1$	$f$	0		
$s_1$	$g$	0		
$s_2$	$f$		0	
$s_2$	$g$		0	
$s_3$	$f$			0
$s_3$	$g$			0