

Name:

Student ID:

CSE 573 Winter 2021 HW1

1/28/2021

100 points

Instructions:


- 1) The homework should be done individually. Don't forget to write your name.
- 2) We highly recommend typing your homework, but writing and scanning, or annotating a PDF are also acceptable.
- 3) Keep your answers brief but provide enough explanations and details to let us know that you have understood the topic.
- 4) The assignment is due on Feb 10.
- 5) You should upload your assignments through gradescope.

Topics:	Points
Search	20
Heuristics for Informed Search	15
Alpha-Beta pruning	15
Expectimax	25
Value Function	15
Markov Decision Processes (MDPs)	10

Problem 1. Search [20 points]


In the following state graph the agent wants to move from the start state (a) to the goal state (G). The agent can move in the direction indicated by the edges.

At every node, the agent observes a list of next possible destinations (children of the current node) in alphabetical order and adds them to the fringe, in that order.



Node	H1	H2
a	9	17
b	3	19
c	14	14
d	10	12
e	7	10
f	8	8

In the following questions, number the nodes in your search tree corresponding to the expansion order (see below for an example) -- break ties alphabetically.



- What is the path returned by Depth First Search (DFS) (ignoring costs on edges)? Show your search tree. (3 points)
- What is the path returned by Breadth-first search (BFS) (ignoring costs on edges)? Show the expanded search tree. (3 points)
- What is the path returned by Uniform Cost Search (UCS) given the costs of the edges in the figure? Show the expanded search tree. (3 points)

- D) Are the functions H_1 and H_2 admissible heuristics for the graph? Why or why not? If not, specify a node that has an inadmissible heuristic value and provide the interval of values that would make the heuristic admissible at that node. (3 points)
- E) What is the path returned by greedy search using heuristic H_2 ? Show the expanded tree. (4 points)
- F) What is the path returned by A* search using heuristic H_1 ? Show the expanded tree. (4 points)


Problem 2. Heuristics for Informed Search [15 Points]

A delivery robot is moving in an $n \times m$ maze. A simple version of the maze is shown in the figure. The robot is programmed to deliver multiple parcels to their destinations. Each parcel starts at some node in the maze and has its own delivery destination. The initial position of the parcels is shown in the figure, and the number on each parcel is its target destination. At every step the robot can take one of the following actions:

- Move: Move in one of these directions: {Up, Right, Down, Left}
- Pick: Pick-up a parcel in a location
- Drop: Put down a parcel at a location.

The cost of each move action is 1, and the costs of Pick and Drop are zero. The robot starts at square number 1 and can move through dotted lines - but the solid lines represent walls. The robot wants to deliver all parcels to their destinations with minimal cost.

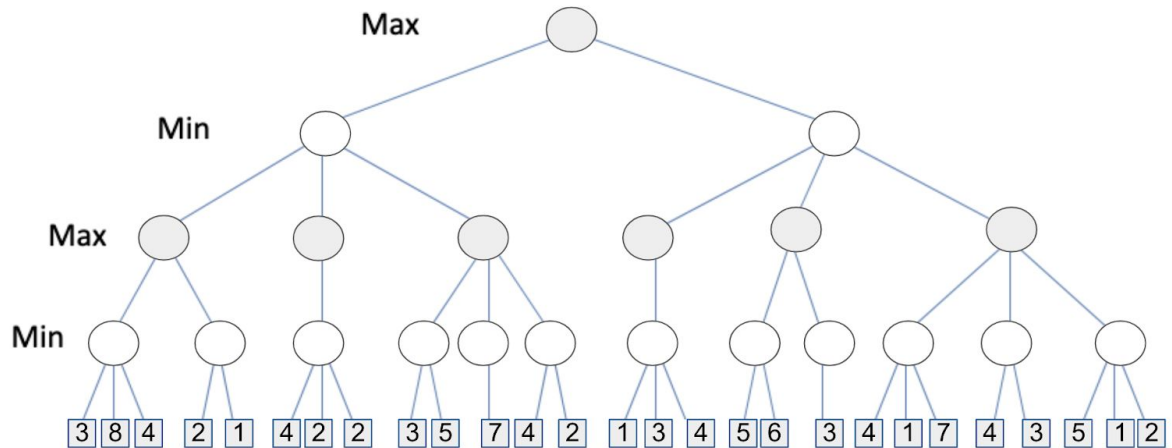
Note that multiple parcels can be placed at the same square. You might come up with different ways of addressing these questions; For each part, there might be several possible heuristic functions. You will get credit if your heuristic follows requested requirements -- clearly write down the assumptions and conditions in which your solution is plausible.



- A) If the robot can carry only one parcel at a time, define an admissible heuristic function for searching the space. Explain in plain english why the heuristic is admissible. Is your heuristic consistent? Why? Make sure your heuristic is not $h(x) = 0$. (5 points)
- B) If the robot can carry multiple parcels at a time, is the function $h(x)$ = "count of packages that are not delivered" admissible and consistent? Why or why not? (4 points)
- C) If the robot can carry multiple parcels at a time, define two admissible heuristic functions. Explain in plain english why each heuristic is admissible. Are your heuristics consistent? Why? Make sure your heuristic is not $h(x) = 0$, and it is not a function of actual cost because it is not practical to compute the actual cost in a general case. *Hint: the heuristic function can be a function of carried parcels and un-carried parcels.* (6 points)

Problem 3. Alpha-Beta pruning [15 points]


Below is the tree showing the states in a 2-player game played by two rational agents. This tree shows the 2-level expansion of decisions, and the values at the leaves are the utility values at those states.



- A) Show the values of every intermediate node after performing the minimax algorithm. (5 points)
- B) Use the Alpha-Beta pruning algorithm to determine the branches that need to be cut. (5 points)
- C) Change the value of one leaf node so as to maximize the number of leaves Alpha-Beta needs to explore in the resulting tree. (5 points)

Problem 4 - Expectimax [25 points]

- A) All the expectimax search tree shown below, fill in the values for interior nodes. Assume the uniform distribution for the expectation nodes (circular nodes). (2 Points)




- B) In the tree above, suppose the opponent chooses the left branch with probability p instead of with equal probability. Find the range of values for p that will change the optimal decision for the root max node, relative to part A. (5 Points)

Small PacMan Maze: In the map below, the pacman and the ghost can move to any adjacent squares (the action space is $\{U, D, R, L\}$). The pacman and the ghost can move in any direction if there is no wall on the way and the solid lines represent the wall. The movement of the ghost is random. Pacman starts first and alternates turns with the ghost. The game starts in the shown figure, and it ends in either one of the two cases. Note that pacman and the ghost cannot move in the direction of a wall (for example, in the start state shown, pacman cannot move 'down' or 'left'; the only valid actions are 'up' or 'right').

- Winning: Pacman (maximizer) eats all the dots.
- Losing: The ghost catches the pacman.

Pacman will receive the score of +1 for eating a dot and a penalty score of -2 if it is caught by the ghost. The final score of the pacman is calculated as the number of dots eaten by the pacman and the penalty if the ghost caught the pacman.




- C) Draw the expectimax search tree for the first four total turns (pacman moves first, then the ghost, then pacman again, and finally the ghost). The ghost moves horizontally with probability of q and it moves vertically with probability of $1-q$. Use the letter 'P' for pacman and 'G' for ghost when drawing game states. On the tree please distinguish the max-nodes, terminal nodes and the expectation nodes. Note that each of the pacman or the ghost's movements counts as the turn. You don't need to expand the whole tree; you can calculate the values of non-terminal states at depth 4. (10 points)
- D) Calculate the expected score of Pacman if it plays optimally for $q=1/3$. Hint: the tree can be infinitely large, but you don't need to expand the whole tree in order to compute this expected score -- Pacman will get a score of +2 for eating both dots and winning the game. (6 points)
- E) Explain in plain English what Pacman's strategy should be (4 points).

Problem 5. Evaluation Function [15 Points]

The Game of Hex was invented by [John Nash](#) in the 1940s. Hex consists of a rhombus game map divided into $n * n$ hexagons. Each player in a 2-player game has a marker (Blue and red).

At each round a player can place a marker on an unmarked hexagon, and players alternate turns. The goal for the players is to link their opposite sides of the board in an unbroken chain. Whoever connects their sides first wins and receives +1 point and the opponent receives -1. It has been proven that a draw is impossible in Hex, hence there's always a winner.

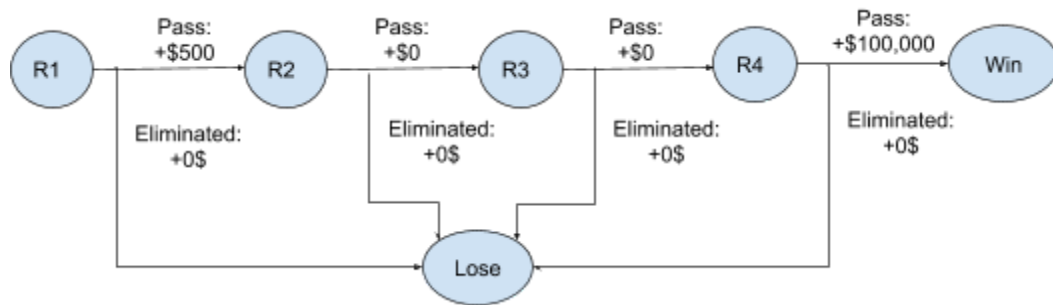


- A) Players can play optimally using a minimax algorithm. Why is expanding the whole game tree not practical? What factors about a state would a good heuristic need to consider? What are the things that we need to consider when we design the evaluation function so we can evaluate different stages of this game? (5 points)
- B) Define an evaluation function that approximates the value of each state. (5 points)
- C) If the size of the game is $n * n$ and each time the agent considers the next three moves (agents' move, minimizers' response, agents' subsequent move). What is the Big- θ time cost of the initial action? (5 points)

Problem 6. MDP [10 Points]

You are controlling a robot walking across a narrow bridge. The goal is to walk straight at every step to get to the reward at the end of the bridge. However, the controller is noisy, so the robot does not always follow your instructions. You control the robot in rounds, one step at a time. At each step, the robot will either step forward, or fall. The robot receives a small \$500 reward after the first successful step (R1), to entice it onto the bridge. In each subsequent step, there is no reward, except at the end of the bridge after the final round (R4), where you receive a \$100,000 reward.

The probability of the robot following your instructions, thereby passing each round, is as follows: $P(R1=pass)=4/5$, $P(R2=pass)=3/4$, $P(R3=pass)=1/2$, $P(R4=pass)=1/5$.



6.1. Model this problem as an MDP. This will require you to fully specify both the transition function, and the reward function. [6 pts]

6.2. What is the value of each state when $\gamma = 1$? [4 pts]