Name:
Student ID:

Name:
Student ID:

CSE 473 Autumn 2019 HW1

10/10/2019

100 points + 15 extra points

Instructions:

1) The homework can be done individually or in a group of two people. Mention the name of your partner in the submission, and you can both submit the same file.
2) We highly recommend typing your homework, but writing and scanning also works.
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 Oct 23rd.

Topics:	Points
Search	35
Heuristics for Informed Search	15+5
Alpha-Beta pruning	15
Expectimax	25
Value Function	10 +10

Problem 1. Search [35 points]

In the following state graph the agent wants to move from the start state (A) to the goal state (G).

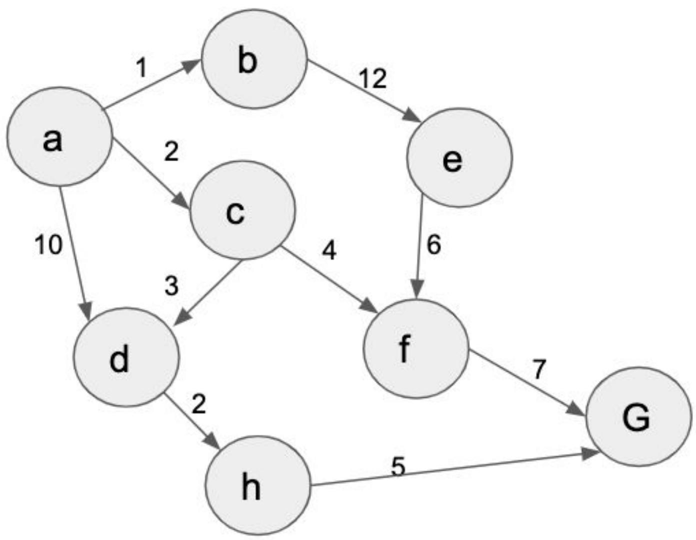

Node	H1	H2
a	10	8
b	11	11
c	9	3
d	5	8
e	10	13
f	5	9
h	4	4

A) What is the path returned by Depth First Search (DFS) if the agent orders the nodes alphabetically (ignoring costs on edges)? Show the expanded search tree. (5 points)

B) What is the path returned by Breadth First Search (BFS) if the agent orders the nodes alphabetically (ignoring costs on edges)? Show the expanded search tree. (5 points)

C) What is the path returned by Uniform Cost Search (UCS) given the costs of the edges in the figure? Show the expanded search tree. (5 points)

D) Are the functions H_1 and H_2 admissible heuristics for the graph. Why or why not? If not provide the node that has inadmissible heuristic value and provide the interval of values that make the heuristic admissible. (10 points)

E) What is the path returned by greedy search using heuristic H_2. Show the expanded tree. (5 points)

F) What is the path returned by A* search using Heuristics H_1. Show the expanded tree (5 points)

Problem 2. Heuristics for Informed Search [15+5 points]

A delivery robot is moving in a n*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 in 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 costs of move actions are 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 the minimum number of moves.

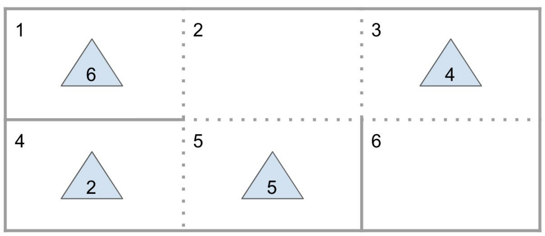

A) If the robot can carry only one parcel at a time, define an admissible heuristic function for searching the space. Explain in plain english why the heuristic is admissible. Is your heuristic consistent? Why? Make sure your heuristic is not h(x) =0. (7 points)
B) If the robot can carry multiple parcels at a time, define an admissible heuristic function. Explain in plain english why the heuristic is admissible. Is your heuristics consistent? Why? Make sure your heuristic is not h(x) =0. (8 points)
C) **Extra credit.** Design the best heuristic you can think of. Is it admissible/consistent? Explain in plain english how would the heuristic guide the agent at each node in the search tree. Make sure the heuristic 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. (5 points)

Problem 3. Alpha-Beta pruning [15 points]

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

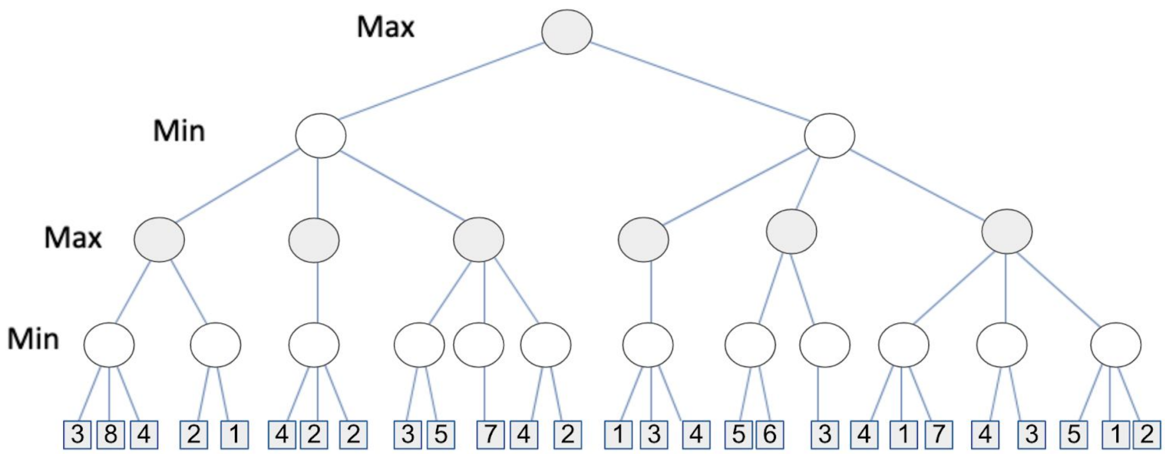

A) Show the values of every intermediate node after performing the minimax algorithm. (5 points)
B) Use the $\alpha - \beta$ Pruning algorithm to determine the branches that need to be cut. (5 points)
C) Change the value of one leaf node so as to maximize the number of leaves Alpha-Beta needs to explore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

Problem 4 - Expectimax [25 points]

In the map below, the pacman and the ghost can move to any adjacent squares (the action space is $\{U, D, R, L\}$). 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.

- 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.

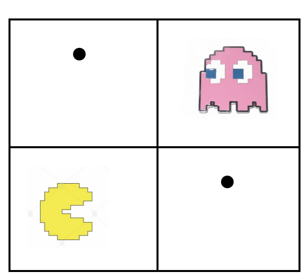

A) Draw the expectimax search tree for the first four turns (meaning Pacman moves twice, and then the ghost moves twice) -- assuming the ghost moves horizontally with probability of *q* and it moves vertically with probability of *1-q*. Use the letter 'P' for pacman and 'G' for ghost when drawing game states. On the tree please distinguish the max-nodes, terminal nodes and the expectation nodes. (10 points)
B) Calculate the expected score of Pacman if it plays optimally for $q=⅓$. (10 points)
C) Explain in plain English what Pacman's strategy should be (5 points).

Problem 5. Value Function [10+10 points]

Game of Hex was invented by John Nash in the 1940s. Hex game 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 draw is impossible, hence there's always a winner.

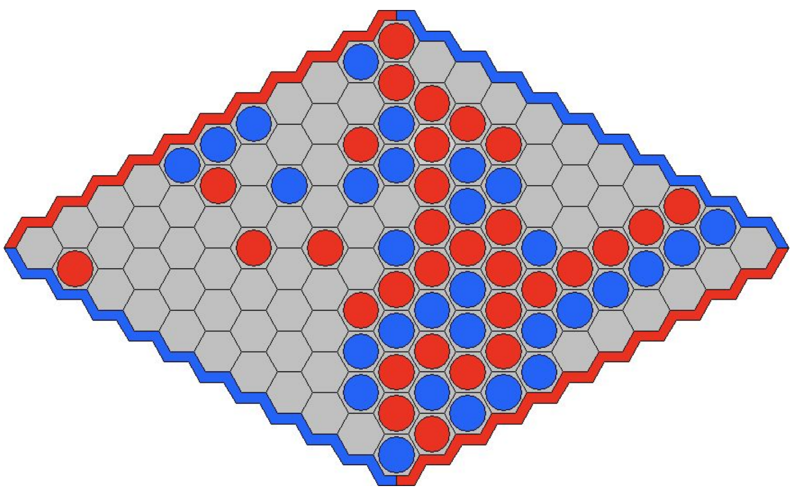

A) Players can play optimally using a minimax algorithm; Why is expanding the whole game tree not practical? How can one approximate the value of states under resource limits? What properties should the evaluation function have for this game? (5 points)
B) Define an evaluation function that approximates the value of each state. (5 points)
C) **Extra credit:** If the size of the game is $n * n$ and each time the agent considers next three moves (2 of itself and 1 for minimizer). What is the time cost of each decision? (Extra 10 points)