Question 1

a) What is one advantage of greedy best-first search over breadth-first search?
b) What is one advantage of breadth-first search over greedy best-first search?
c) Can depth-first search ever find an optimal solution by exploring fewer states than A* search with an admissible and consistent heuristic? Why or why not?

Answers

a) TODO
b) TODO
c) TODO

Question 2

Consider the below maze, where grey cells indicate walls. A search algorithm was run on this maze, and found the yellow highlighted path from point A to B. In doing so, the red highlighted cells were the states explored but that did not lead to the goal.

Of the four search algorithms discussed in lecture — depth-first search, breadth-first search, greedy best-first search with Manhattan distance heuristic, and A* search with Manhattan distance heuristic — which one (or multiple, if multiple are possible) could be the algorithm used below and why?

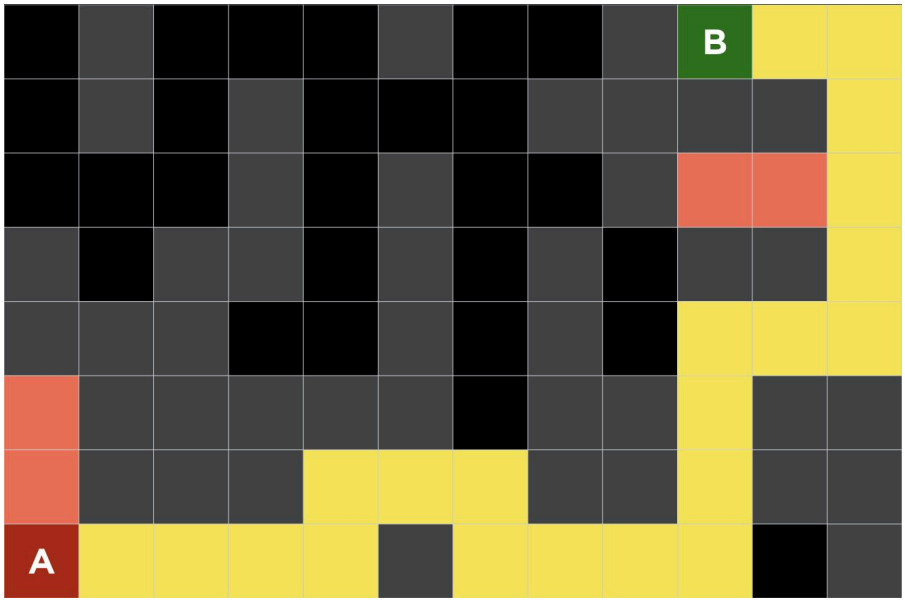

Answer

TODO

Question 3

a) Between Minimax and depth-limited Minimax with an evaluation function, can either of them produce a more optimal move than the other? Why or why not?
b) Between Minimax and Minimax with alpha-beta pruning, can either of them produce a more optimal move than the other? Why or why not?

Answers

a) TODO
b) TODO

Question 4

Consider the Minimax tree below, where the green up arrows indicate the MAX player and red down arrows indicate the MIN player. The leaf nodes are each labelled with their value, and each leaf node is also labelled with a letter for ease of reference.

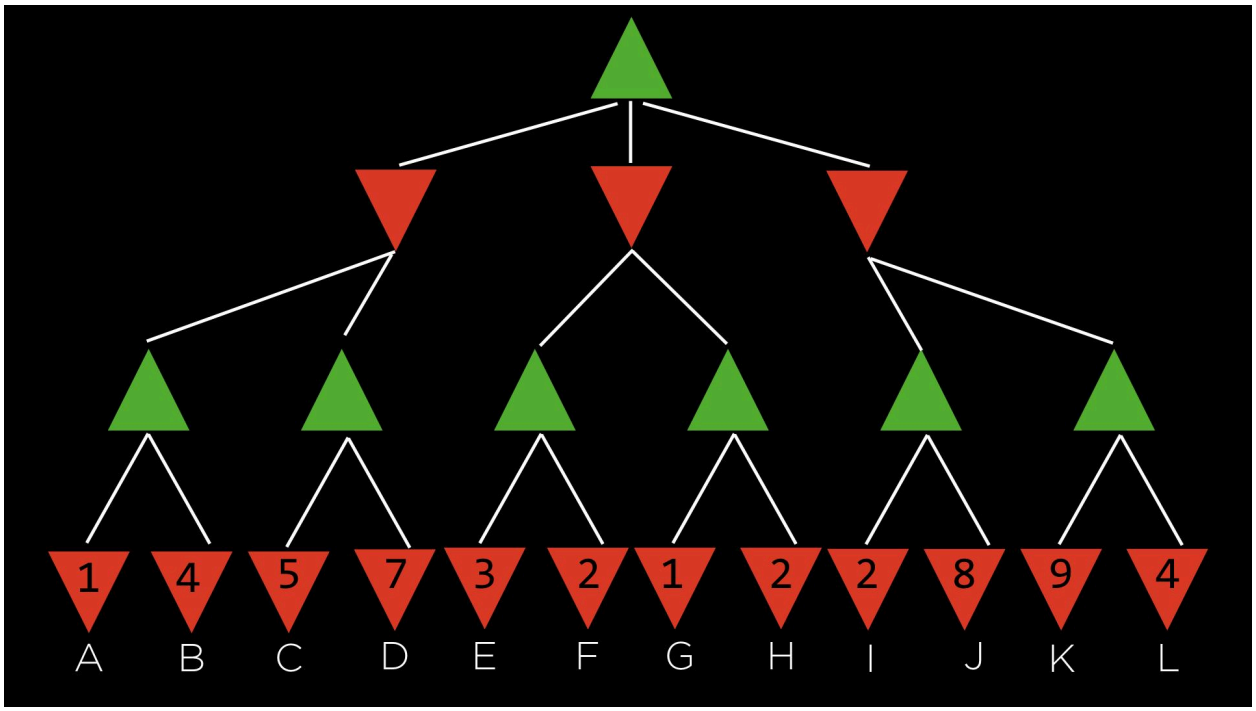

a) What is the value of the root node?
b) If using alpha-beta pruning, which nodes can be pruned?

Answer

a) TODO
b) TODO