Name	1	2	3	4	5	total
	10	10	35	35	35	125

UMBC CMSC 471 Midterm Exam 21 March 2016

Write your answers on this exam, which is closed book and consists of five problems, summing to 125 points. You have the entire class period, seventy-five minutes, to work on this exam. Good luck.

1. True/False [10 points]

Circle either **T** or **F** in the space before each statement to indicate whether the statement is true or false. If you think the answer is simultaneously true and false, quit while you are ahead.

T F Alan Turing proposed his famous Turing test as a technique for deciding whether or not a problem was "Turing computable".

T F A greedy, best-first search algorithm is always complete.

T F A simple breadth-first search always finds a shortest solution if one exists that is of finite length.

T F For a search problem, the path returned by uniform cost search may change if we add a positive constant C to every step cost.

T F The alpha-beta algorithm is preferred to minimax because it provides a better estimation of which move is best for a given lookahead distance.

T F In a two-player, zero-sum game there is always a winner and a loser.

T F In a prisoners' dilemma game, each player chooses a dominant strategy, but each could do better if both chose different strategies.

T F Game theory predicts that players will always have a dominant strategy.

T F Tit-for-tat is a strategy that cannot be applied in repeated games.

T F If h1(s) and h2(s) are admissible A* heuristics, then their average, (h1(s) + h2(s))/2 must also be admissible.

2. Multiple Choice [10 points]

Circle *all* of the correct answers.

2.1 [2] Which of the following search algorithms finds an optimal solution?

a. Breadth first
b. Hill climbing
c. Depth first
d. Greedy search
e. None of the above

2.2 [2] A search algorithm is *complete* if it…

a. Always finds the optimal solution
b. Always finds a solution if there is one
c. Finds all possible solutions
d. Never finds a solution

2.3 [2] Which of these techniques uses randomness to avoiding getting trapped in local maxima?

a. Best first search
b. Local beam search
c. Simulated annealing
d. Gradient descent
e. None of the above

2.4 [2] Which of the following is a problem that occurs in hill climbing search?

a. Cliffs
b. Ridges
c. Valleys
d. "Slippery slopes"
e. Local maxima and minima
f. Cycles in the graph
g. None of the above

2.5 [2] A greedy search uses a heuristic function to expand the node that

a. Appears to be closest to the goal
b. Is closest to a goal
c. Is closest to the start state
d. Is the leftmost node in the search tree
e. None of the above

3. Search I [35 points]

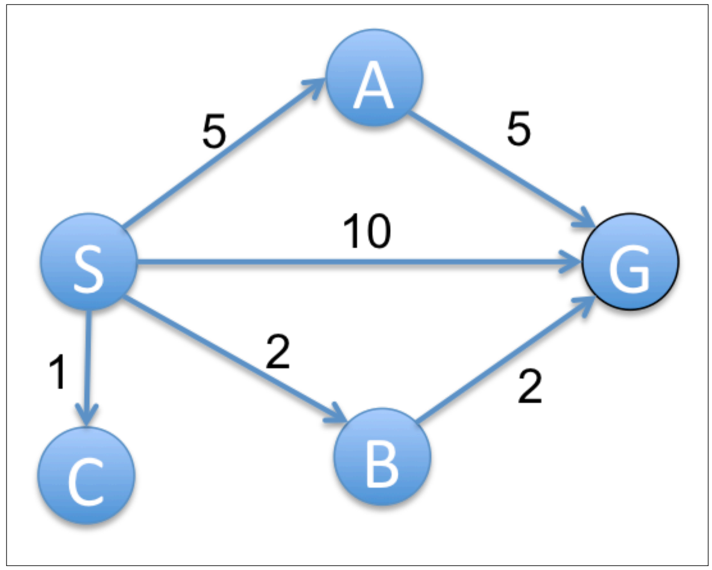

Consider the search space with this graph, where S is the start state and G is the goal. Assume a heuristic is available for the A* algorithm that has the following values:

{S:4, A:3, B:2, C:100, G:0}

For questions a-e specify the path as a sequence of nodes starting at S and ending at G that each algorithm returns. Assume the successor functions work so that nodes are explored in alphabetical order whenever possible.

(a) [5] Breadth-first search

(b) [5] Depth-first search

(c) [5] Uniform-cost search

(d) [5] A* search

(e) [5] Greedy search

(f) [5] Name a node that uniform-cost search will expand, but A* will not.

(g) [5] Judging just from the heuristic function for these nodes, is it an admissible heuristic?

4. Minimax and Expectimax [35]

4.1 [6] Consider a zero-sum game with two players. Each leaf is labeled with the payoff player 1 receives. It is player 1's turn to move. Assume both play optimally at every ply (i.e. player 1 seeks to maximize payoff while Player 2 seeks to minimize it). Circle player 1's best next move on the graph, and show the minimax values for each node in the tree.

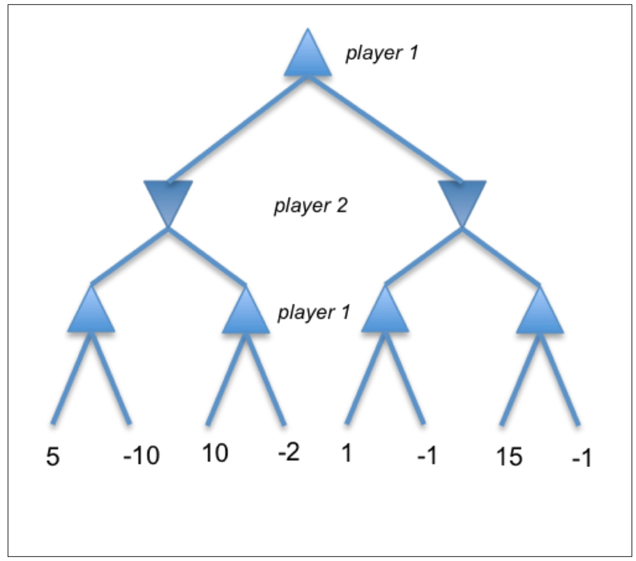

4.2 [6] Consider this game tree. Player 1 moves first, and attempts to maximize the expected payoff. Player 2 moves second, and attempts to minimize the expected payoff. Expand nodes left to right. Cross out both internal and leaf nodes that would be pruned by the alpha-beta algorithm

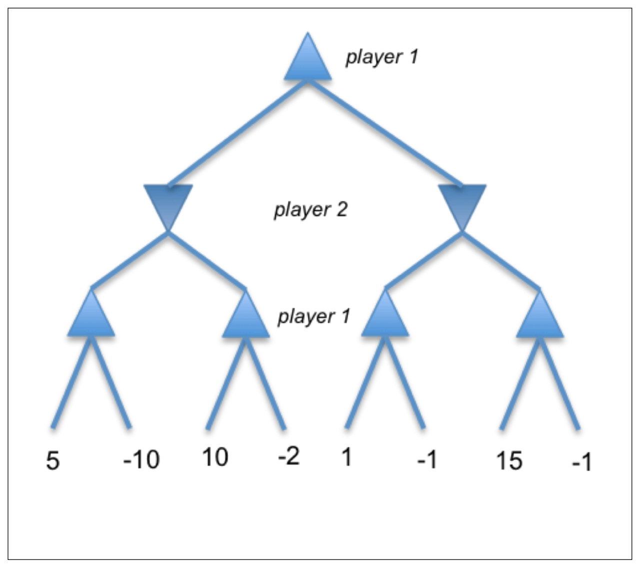

4.3 [6] Assume that player 2 chooses an action uniformly *at random* every turn and that player 1 knows this. Player 1 still seeks to maximize her payoff. Circle player 1's optimal next move, and give her expected payoff. Show the expected value at each node in the tree.

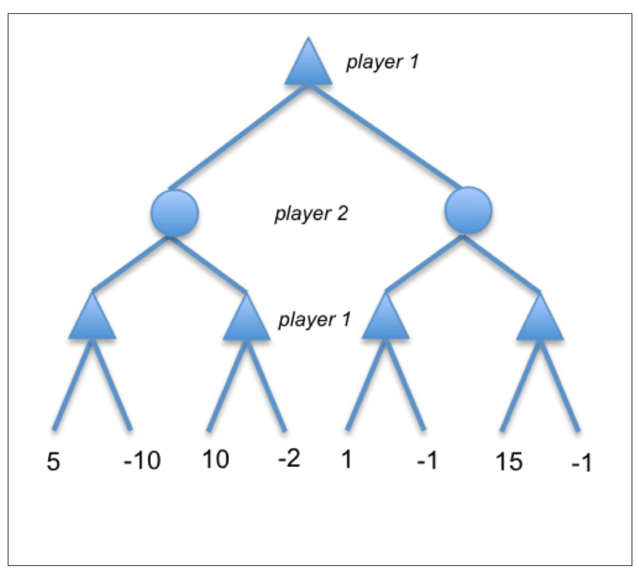

4.4 [6] Now consider a modified version of the game tree, where the leftmost leaf node has an unknown payoff **X**. Player 1 moves first, and attempts to maximize the value of the game.

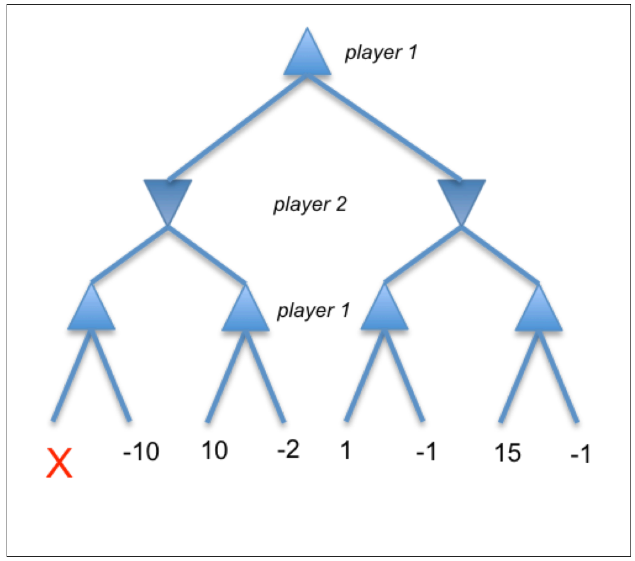

4.5 [3] Assume player 2 is a minimizing agent (and Player 1 knows this). For what values of **X** does player 1 choose the left action?

4.6 [3] Assume Player 2 chooses actions at random (and Player 1 knows this). For what values of **X** does Player 1 choose the left action?

4.7 [5] For what values of **X** is the minimax value backed up (i.e., when player2 chooses rationally) to the tree's root worth more than the expectimax value at the root (i.e., when player 2 chooses randomly)?

X > 10

5. Constraint Satisfaction [35]

Consider the crossword puzzle shown on the right. Suppose we have the following words in our dictionary: *ant, ape, big, bus, car, has, bard, book, buys, hold, lane, year, rank, browns, ginger, symbol, syntax*. The goal is to fill the puzzle with words from the dictionary.

This problem asks you to set the problem up for solution as a constraint satisfaction by specifying the variables, their domains and constraints.

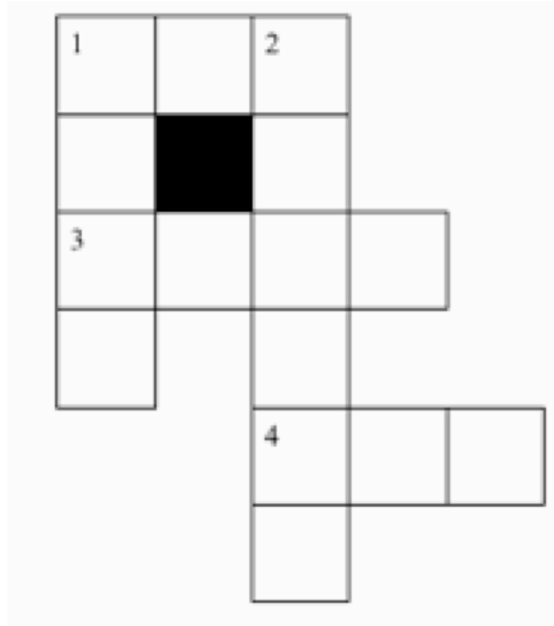

5.1 [10] Describe a set of variable for this problem and give the initial domain for each. If you need more table rows, use the back of this sheet.

var	description	Initial domain
V1		
V2		
V3		
V4		
V5		
V6		
V7		
V8		
V9		
V10		

5.2 [10] Draw a constraint graph for the problem, e.g., a graph with variables as nodes and an edge between variables when they have a constraint between their values.

5.3 [5] Describe how you could represent a variable value in Python, e.g., as an integer, float, character, string, tuple, list, dictionary, set, class instance, etc.

5.3 [10] For this representation describe each of the constraints. Put one constraint in one row, specifying the two variables involved and a short description. Feel free to use Python to describe it or a simple sentence. If you need more table rows, use the back of this sheet.

Var1	Var2	Constraint