Homework #2: Chapters 3 and 5

The following exercises are due at the beginning of class on **Tuesday, February 11.** In addition to the two problems on the back, there is an extra-credit assignment worth up to an additional 20 points. *This was revised on Feb. 7 (see the highlight on #4 on page 2)*

1. *[15 points]* Use greedy best-first search to find a path from Mehadia to Sibiu, assuming the roads between cities and their lengths are given by the map in Figure 3.2 of the book (p. 68). Use the straight-line distance from each node to Sibiu as your heuristic function, assuming these distances are given by the following table:

Arad	138		Mehadia	146
Bucharest	253		Neamt	237
Craiova	196		Oradea	148
Dobreta	183		Pitesti	162
Eforie	433		Rimnicu Vilcea	78
Fagaras	96		Sibiu	0
Giurgiu	272		Timisoara	141
Hirsova	385		Urziceni	302
Iasi	298		Vaslui	345
Lugoj	97		Zerind	146

   Show your search tree, including the h(n) value for each node, and label each node with the order in which it is expanded (note, this may be different from the order it is generated). In order to reduce unnecessary search, you should use graph-based search.

2. *[15 points]* Now repeat the exercise above, but use A* instead of greedy best-first. Show your search tree, complete with f(n), g(n) and h(n) values for each node, and label each node with the order in which it is expanded. Once again, use a graph-search approach, but remember that if a generated node has a lower f(n) than a frontier node with an identical state, then it must still be explored.

3. *[20 points]* Use A* to solve the 8-puzzle with the initial and goal states shown below. Assume that your path cost is 1 per move and that your heuristic function is the Manhattan distance of all tiles from their correct placement (note, the blank does not count as a tile). Show your search tree, complete with f(n), g(n) and h(n) values for each node and label each node with the order in which it is expanded. Also show the current game board at each node in the tree. Once again, use the graph search version of the algorithm.

Initial State

2	4	3
1		5
7	8	6

Goal State

1	2	3
4	5	6
7	8	

4. *[40 points]* This problem looks at playing the game tic-tac-toe. Assume that X is the MAX player. Let the utility of a win for X be 10, a loss for X be -10, and a draw be 0. There are two parts to this question, each using one of the two game boards given below:

board1

X	O	
	O	X
O	X	

board2

		X
O		
X		O

a) Given the game board **board1** above where it is X's turn to play next, show the entire game tree. Mark the utilities of each terminal state and use the minimax algorithm to calculate the optimal move.

b) Given the game board **board2** on the opposite page where it is X's turn to play next, show the game tree with a cut-off depth of two ply (i.e., stop after each player makes one move). Use the following evaluation function on all leaf nodes:
$Eval(s) = 10X_3(s) + 3X_2(s) + X_1(s) - (10O_3(s) + 3O_2(s) + O_1(s))$
where we define $X_n(s)$ as the number of rows, columns, and diagonals in state *s* with exactly *n* X's and no O's, and similarly define $O_n(s)$ as the number of rows, columns, and diagonals in state *s* with exactly *n* O's and no X's. Use the minimax algorithm to determine X's best move.

5. *[10 points]* Consider the following game, which I call fully-observable, two-player, five-card hearts. Each player is dealt 5 cards, and both players can see each other's cards. The first player must play a card to lead the first "trick." If the other player has a card in the same suit they must play one of those cards. Otherwise, they can play any card. The player with the highest card in the suit that was led takes the trick. This player then leads the next trick. After all cards have been played, utilities are assigned as follows: -1 for each heart that you have taken in a trick, and an additional -13 if you have taken the queen of spades.
   a) In what ways is this game different from the two-player zero sums games we have discussed in class?
   b) For a given deal, what is the maximum number of games (i.e., distinct sequence of moves) that can be played? Be sure to explain how you derived your answer.

Extra Credit *(+20 points)*:
This optional exercise requires you to do some Java programming in order to conduct an experiment that compares uniform-cost, greedy best-first, and A* search. In order to do this, download and read the code I have made available from the course web site (under Additional Class Materials). This code implements all three search algorithms and provides two *abstract* classes: SearchProblem and State. In order to solve a particular problem, you only need to extend these two classes with details specific to your problem.

Your task is to extend the code to solve the following three configurations of the 8-puzzle, where the goal state for all three configurations is as specified to the far right.

Initial State #1	Initial State #2	Initial State #3	Goal State

7	1	4
6	3	2
	8	5

4	8	2
6	3	5
1		7

7	5	3
6		4
8	1	2

1	2	3
8		4
7	6	5

You will then run each of the three search algorithms on each puzzle, recording the path cost of the solution (if found), number of nodes expanded, number of nodes generated, and time to perform the search. To do this, you'll need to write a class that extends State and can record the state of the game (i.e., the current position of each tile). It is important that this class implements an equals() method that can be used to compare the current state to another state. You'll need to write a second class that extends SearchProblem and implements the four methods: getInitialState(), goalTest(), getSuccessors(), and getHeuristicValue(). You may want to include a constructor that allows you to initialize the class with different initial states. Note, getSuccessors() returns a list of Successor elements, where each Successor records the State, a string describing the action to reach it, and the step cost of executing that action. By including the step cost in the successor information, we can avoid providing a separate path cost function. For the heuristic, use the sum of the Manhattan distances of all tiles from their goal positions. Finally, you'll need to write a main() method that runs the tests.

After you collect your data, write an analysis of it. What appears to be the strengths and weaknesses of each algorithm based on your experiment? Did the experimental results agree with the theoretical properties of the algorithms discussed in class and in the book? What, if anything, surprised you?

Attach a hardcopy of your code, the output of your experiment, and your analysis to your homework submission. Submit your source (.java) and compiled (.class) files to via e-mail to heflin@cse.lehigh.edu with subject line: "CSE 327: HW #2 Extra Credit".