

Name

1	2	3	4	5	total
10	20	30	35	35	130

UMBC CMSC 471 Midterm Exam

15 March 2017

Write your answers on this exam, which is closed book and consists of five problems, summing to 120 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 A simple reflex agent can be rational
FALSE. A reflex agent always responds to a given situation the same way without considering its expected outcome

T F It is impossible for an agent to act rationally in a partially observable environment
FALSE. Rationality implies choosing an action that optimizes expected outcome given what you know, which may only be partial

T F Iterative deepening search involves re-running breadth-first search repeatedly with increasing depth limits
FALSE. It does this with depth-first search, not breadth first search

T F Simulated annealing is a variation on hill climbing search that can prevent getting stuck in local minima
TRUE. SA is like hill-climbing with a varying probability of choosing a sub-optimal move

T F An optimal solution path for a search problem with only positive costs will never have repeated states
TRUE. In any solution path with repeated states, removing a cycle yields a solution path with lower cost

T F The alpha-beta algorithm is preferred to minimax because it computes the same answer as minimax while usually doing so without examining as much of the game tree
TRUE. Alpha-beta and minimax always return the same values but alpha-beta can avoid having to process parts of the tree

T F There exist constraint satisfaction problems that can be expressed using trinary (three-variable) constraints, but not binary constraints.
FALSE. Such a problem can be transformed into an equivalent one with additional variables so that only binary constraints are needed

T F Game theory only works for non zero-sum games.
FALSE. An important feature of the game theory framework is that it works for both zero-sum and non-zero sum games

T F Arc consistency is a more powerful constraint propagation algorithm than forward checking
TRUE. Forward checking only applies to variables with only a single possible value and eliminates values from the domain of variables that share a constraint with it. Arc consistency works on the entire set of variables and eliminates values from them that cannot satisfy associated constraints.

T F If h_1 and h_2 are admissible search heuristics, then $h_3 = 2 \cdot h_1 - h_2$ must also be admissible.
FALSE. A counter example is when h_1 is h^* (i.e., the true distance) and h_2 be 0 (i.e., the null heuristic)

2. Short answers [20: 5 points each]

(a) Iterative deepening is sometimes used as an alternative to breadth first search. Give one important advantage of iterative deepening over BFS, and give one significant disadvantage of iterative deepening as compared with BFS. Be concise and specific!

Advantage: iterative deepening requires less memory

Disadvantage: iterative deepening repeats computations and therefore can require additional run time

(b) Game theory suggests that a rational agent should always defect in a prisoner's dilemma game, but if both players follow this reasoning, it results in a situation where both are worse off than if they cooperated. Describe a scenario where a rational agent may choose to behave differently than always defecting.

If the players engage in multiple games with each other, they can adapt their strategy to take into account the past behavior of other players and do better. For example using a tit-for-tat strategy that takes into account the last decision of their current player.

(c) What is a main advantage of A* search over hill-climbing search?

A* search is not susceptible to local minima and is, unlike hill-climbing, is complete and admissible.

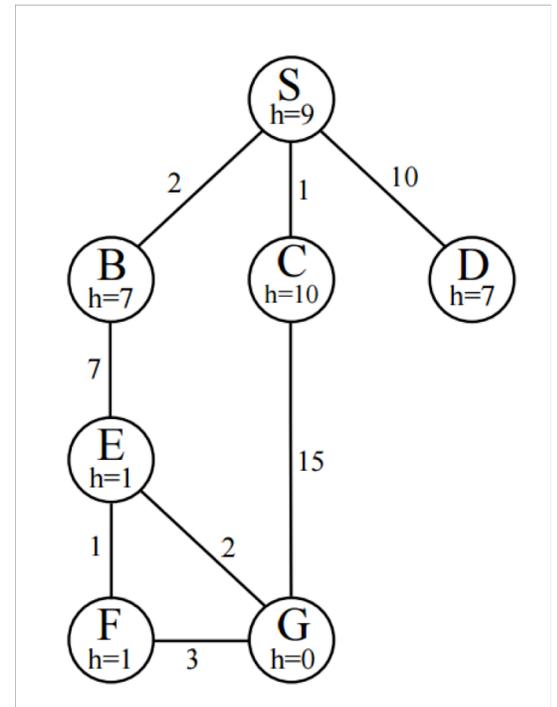
(d) What is a main advantage of hill-climbing search over A* search?

Space efficiency. hill-climbing stores only the current node on the OPEN list, whereas A* stores all nodes on the current frontier.

3. Problem solving as search [30 points]

Consider the search graph shown below. S is the start state and G is the goal state. All edges are bidirectional and are annotated with their cost. A node's heuristic value is shown inside the node (e.g., $h=9$).

For each of the following search strategies, give the path that would be returned, or write none if no path will be returned. If there are any ties, assume alphabetical tiebreaking (i.e., nodes for states earlier in the alphabet are expanded first in the case of ties).



(a) [5] Depth-first graph search

S-B-E-F-G

(b) [5] Breadth-first graph search

S-C-G

(c) [5] Uniform cost graph search

S-B-E-G

(d) [5] Greedy graph search (aka best first graph search)

S-B-E-G

(e) [5] Algorithm A graph search using the heuristic function values shown

S-B-E-G

(f) [5] is the heuristic as shown for this graph admissible? Explain why or why not.

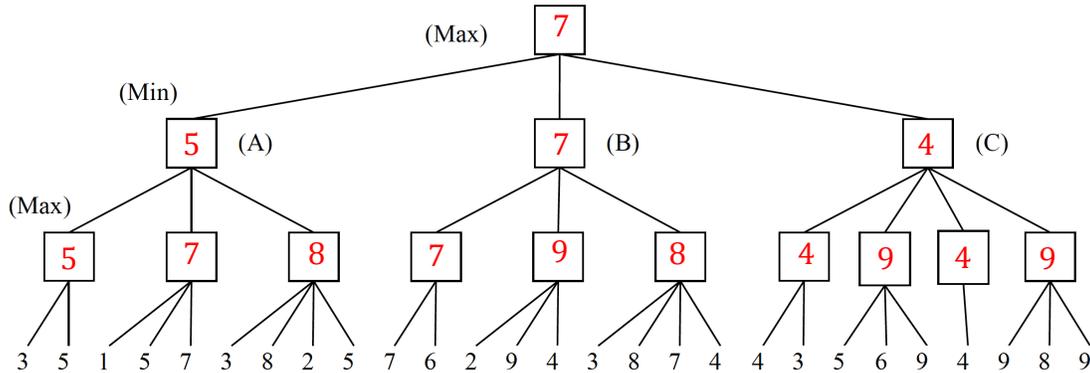
Yes, for each of the nodes, the heuristic value is less than or equal to the actual cost of the shortest path from the start state to a goal state.

4. Game tree search [35 points]

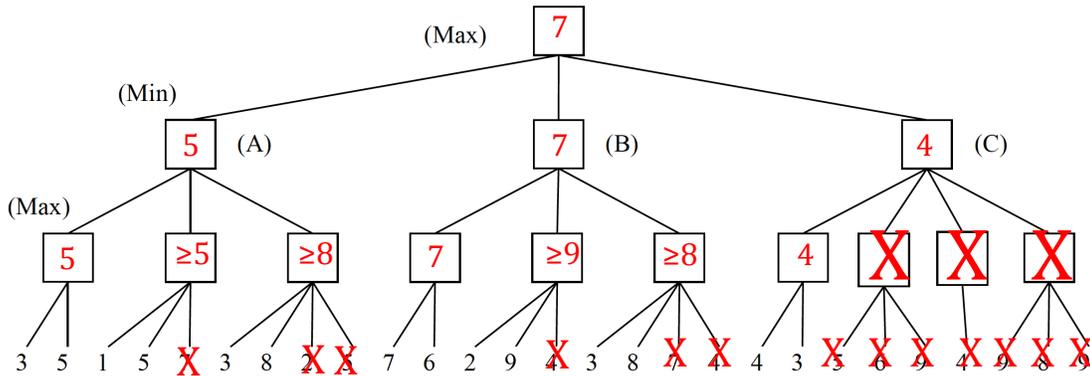
Consider the game tree below in which the first player is trying to maximize her score and the number at the leaves are the values returned by a static evaluator for the board positions reached.

(a) [10] Fill in each box with the value returned by the standard minimax algorithm

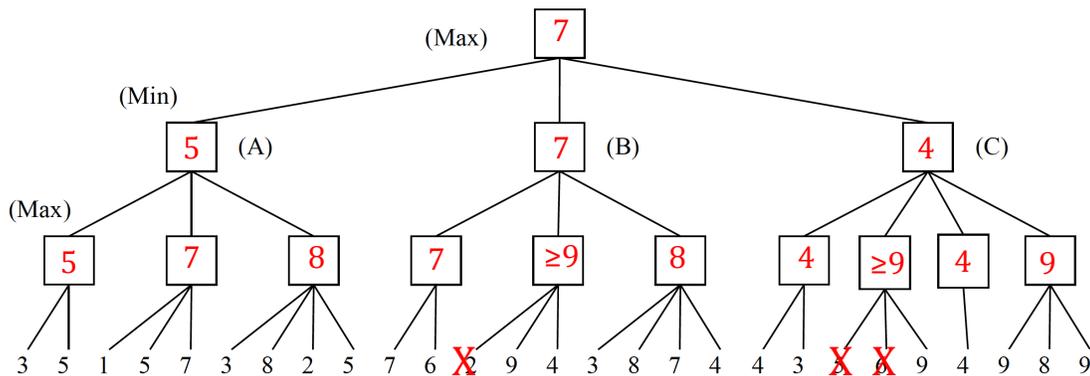
(b) [5] Circle the best initial move for the first player: A B C



c) [10] In the copy of this game tree below, fill in each box with the value returned by the standard **alpha-beta algorithm** if the tree is processed from **left to right**. And cross out both leaves and nodes that need not be examined or considered.

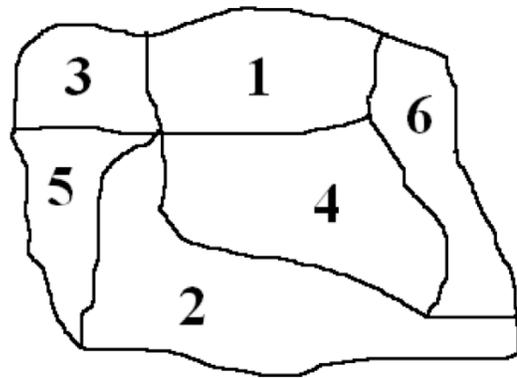


d) [10] In the copy of this game tree below, fill in each box with the value returned by the standard **alpha-beta algorithm** if the tree is processed from **right to left**, crossing out both leaves and nodes that need not be examined or considered.



5. Constraint satisfaction [35 points]

Consider coloring the six-region map shown to the right with three colors: R, G, B so that no two adjacent regions that share a border have the same color. Regions that only meet at a corner are not considered adjacent (e.g., 5 and 1).

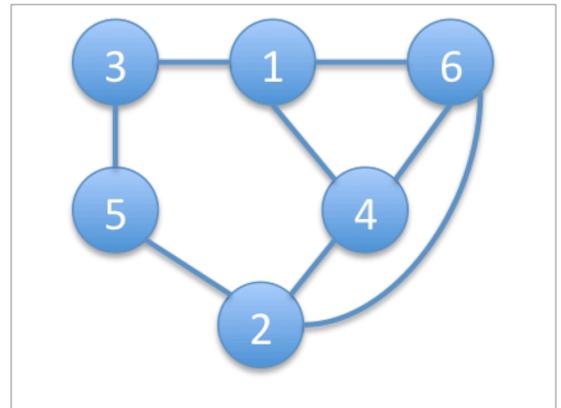


(a) [5] Identify the variables that should be used to set this up as a CSP problem and the domain of possible values for each variable.

There are six variables, which we could name as 1,2,3,4,5,6

The domain of each is {R, G, B}

(b) [10] Draw an abstract constraint graph with a node for each variable and an edge between two variables if there is a constraint between them



(c) [10] If we assign region 1 to be R and region 2 to be G. Perform forward checking and show the variables and their possible values for all six regions. Recall that when a value is assigned to a variable V, forward checking eliminates inconsistent values in variables connected to V in the constraint graph. Is a solution possible from this state?

- $1=\{R\}, 2=\{G\}, 3=\{G, B\}, 4=\{B\}, 5=\{R, B\}, 6=\{B\}$
- **No solution is possible**

(d) [10] Assume the initial domains of the regions in the map above are given as $1=\{R,G,B\}$, $2=\{R,G\}$, $3=\{R,G,B\}$, $4=\{R\}$, $5=\{R,G,B\}$, and $6=\{R\}$. What is the result of applying the Arc Consistency algorithm, AC-3? Is a solution possible from this state?

There are two possible answers:

- $1=\{G, B\}, 2=\{G\}, 3=\{R, G, B\}, 4=\{R\}, 5=\{R, B\}, 6=\{R\}$
- $1=\{G, B\}, 2=\{G\}, 3=\{R, G, B\}, 4=\{R\}, 5=\{R, B\}, 6=\{R\}$

No solution is possible