Introduction to Artificial Intelligence — Dieter Fox
CSE 473, Spring 2017

Midterm

May 3, 2017

Directions

This exam has 6 problems worth 93 points and you have 50 minutes to complete it.

- The exam is closed book. No calculators are needed.
- If you have trouble with a question, by all means move on to the next problem—or the next part of the same problem.
- If a question is unclear, feel free to ask for clarification!
- On True/False questions, incorrect answers will incur negative points.
- **Please do not turn the page until I indicate that it is time to begin.**

Name: ______________________________

Number: ______________________________

1	/20
2	/25
3	/10
4	/10
5	/16
6	/12
Total	/93

1. (20 points, 2pt each, -2pt for wrong answer) **True or False**

   Please circle the correct answer.

   (a) Search (b denotes the branching factor of a search, d the depth of the solution, and m the maximum depth of a search space):

      i. The space complexity of breadth-first search is $O(b^m)$. . . . . T F

      ii. The space complexity of iterative deepening depth-first search is $O(bd)$. . . . . T F

   (b) Depth-first search is complete for infinite state spaces when cycle elimination is applied and each node has only a finite number of direct successors . . . . T F

   (c) If a heuristic is admissible, then it is also consistent. . . . . T F

   (d) If run long enough, hill climbing always finds the optimal solution, even without re-starts. . . . . T F

   (e) For a Constraint Satisfaction Problem, given a choice of variable, we want to consider the most constraining value first (the one that rules out the most values in the remaining variables) . . . . T F

   (f) Given preferences that satisfy the axioms of rationality, there always exists a real-valued function U such that values assigned by U preserve preferences of prizes and lotteries. . . . . T F

   (g) An agent that uses Expectimax search may achieve a better score when playing against a suboptimal adversary than a Minimax agent would. . . . . T F

   (h) For smaller values of γ, the policy computed by value iteration becomes more myopic, or short sighted. . . . . T F

   (i) Policy iteration is suboptimal because it may terminate too early. . . . . T F

2. (25 points, 5pt each) **Graph Search**

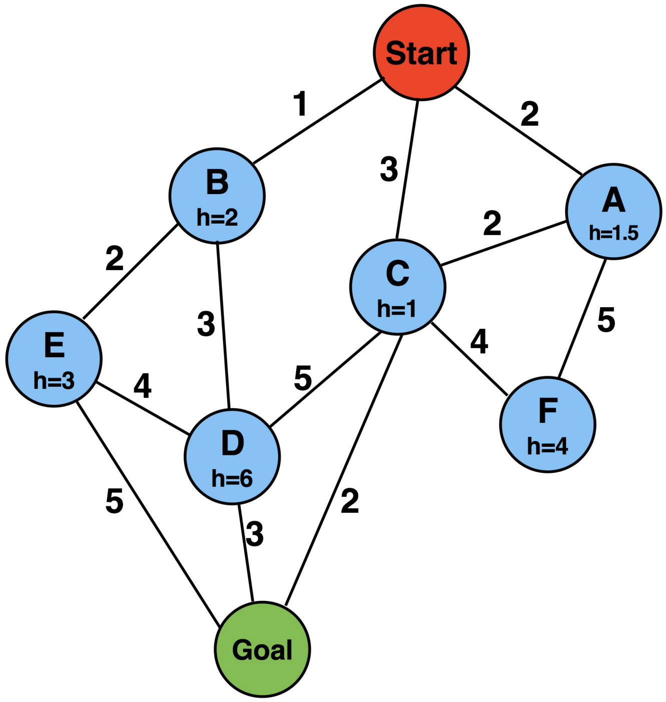

Given the graph above, show the order in which the states are visited by the search algorithms listed below. **Please do graph search with duplication check, not tree search.** Path cost of an edge is given by the number next to the edge. The heuristic estimate of path cost from a state to the goal state is indicated in the circles.

If a state is visited more than once, make sure to write it down each time. Ties (e.g. which child to first explore in the depth first search) should be resolved according to alphabetic order (i.e. A is expanded before Z). **Remember to include the start(S) and goal(G) states in your answer.** Treat the goal state as G when you break ties.

(a) Depth First Search

(b) Breath First Search

(c) Iterative Deepening Depth First Search

(d) Uniform Cost Search

(e) A* Search

3. (10 points, 5/5 pt) **Minimax**

(a) Fill in the values for each of the interior nodes in the following minimax search tree (no pruning).

(b) Indicate which edges would be pruned if you had filled out the tree using alpha-beta pruning. Assume nodes are expanded left to right. You should draw an X through any such edge(s).

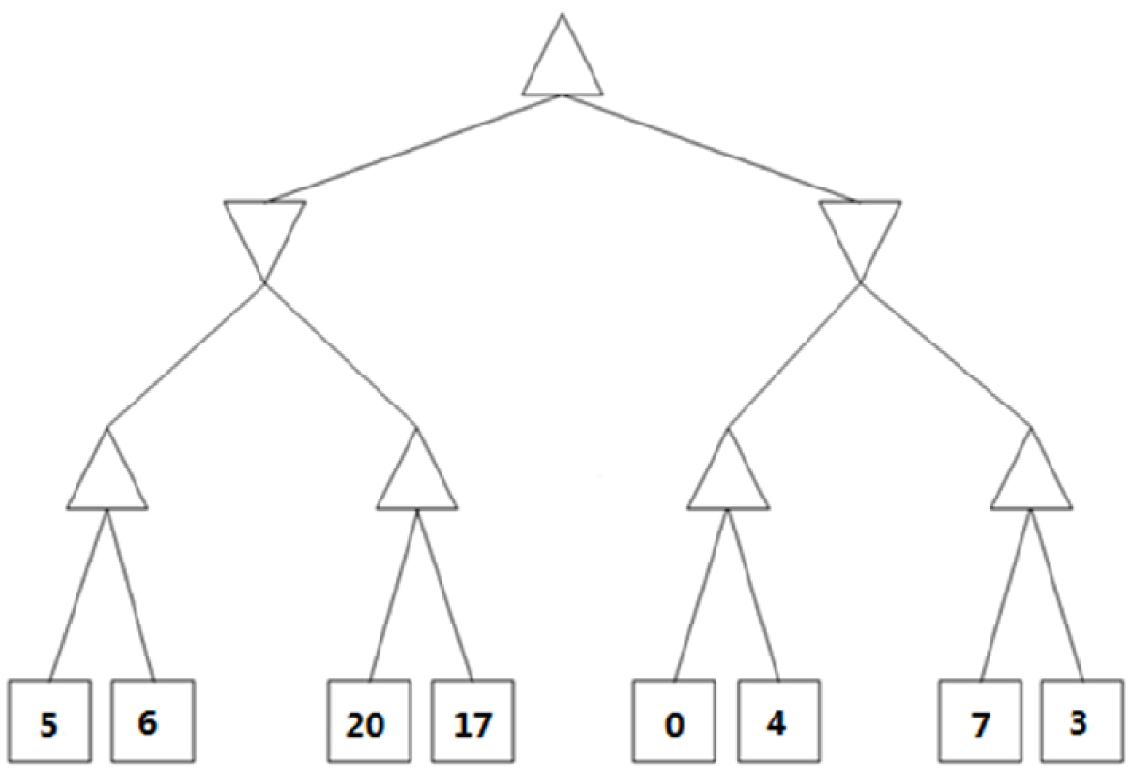

4. (10 points, 5/5 pt) **Expectimax**

(a) Fill in the values for each of the interior nodes in the following expectimax search tree. The maximizing nodes are triangular, and the expectation nodes are circular. Assume a uniform distribution over outcomes for the expectation nodes.

(b) Now, let p be the probability of following the left branch from an expectation node, and $(1 - p)$ be the probability of following the right branch. Give the range of values of p that change the optimal action.

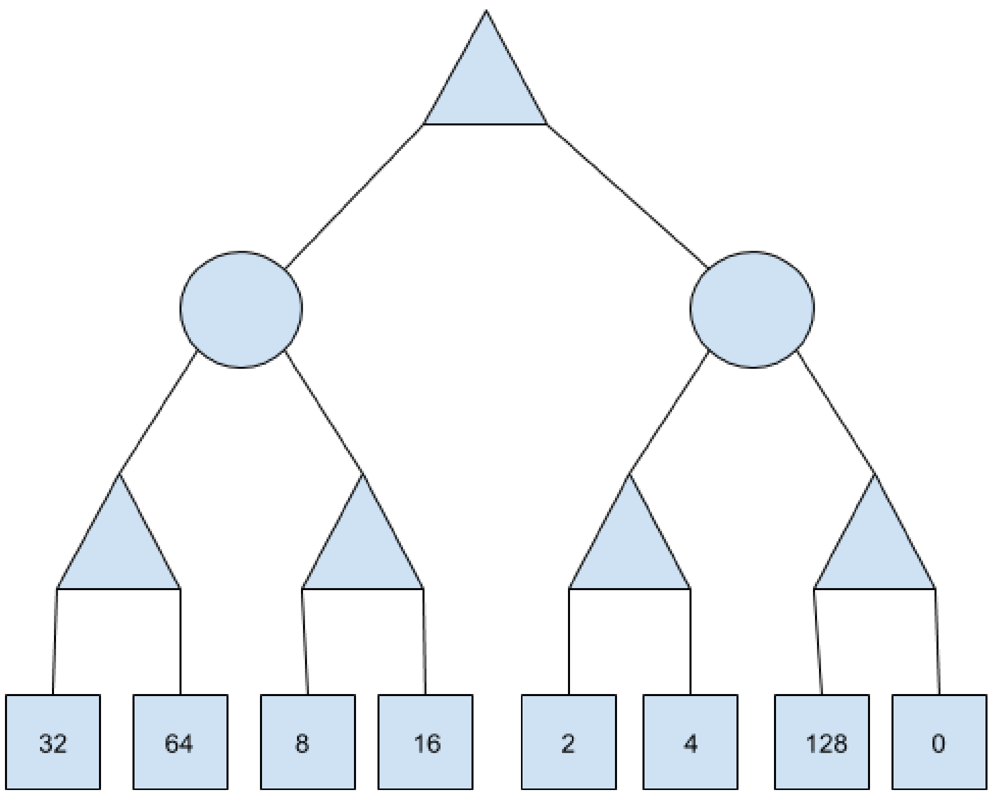

5. (16 points, 5/5/4/2 pt) **Markov Decision Processes**

You are a mountain climber looking to find a treasure at the summit of Mt. Fox. At the top of the mountain (S_5) is a prize with value 100, and at the base there is an exit (S_0) with value 1. The mountain is an MDP with the following layout:

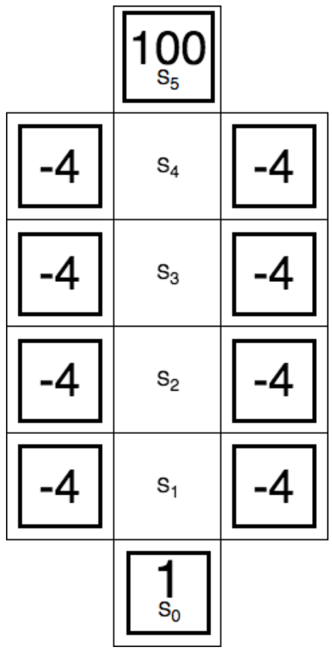

Double-rectangle states are exit states, from which the only action is to *Exit*. Moving into one of these states immediately gives the reward shown.

At each nonterminal state, you can choose to climb either *Up* (U) or *Down* (D). Climbing down moves down with probability 1. Climbing up will move upwards with probability 0.5, and with probability 0.5 will result in slipping off the mountain to certain death and a penalty of -4.

Use the value $\gamma = 1$ as your discount factor.

(a) **Q-Value iteration**: Complete the table by filling in the values computed over the first time steps of Q-value iteration:

t	S_0	S_1, U	S_1, D	S_2, U	S_2, D	S_3, U	S_3, D	S_4, U	S_4, D	S_5
0	1	0	0	0	0	0	0	0	0	100
1	1									100
2	1									100

Policy iteration: You approach the base of the mountain and discover a map left by a previous traveler. Scribbled on the map you see a policy and some values:

$$\pi(S_1) = U \qquad \pi(S_2) = D \qquad \pi(S_3) = D \qquad \pi(S_4) = U$$

$$V(S_1) = -4 \qquad V(S_2) = -4 \qquad V(S_3) = 8 \qquad V(S_4) = 10$$

(b) Using this fixed policy, perform policy *evaluation* by computing a value for each nonterminal state. Use the same transition and reward functions as before. The table is initialized with the provided values.

t	S_0	S_1	S_2	S_3	S_4	S_5
0	1	-4	-4	8	10	100
1	1					100
2	1					100

(c) What policy is produced based on the new values of V^π? (If there is a tie between actions, keep the original policy.)

$\pi(S_1) =$ __________

$\pi(S_2) =$ __________

$\pi(S_3) =$ __________

$\pi(S_4) =$ __________

(d) Will policy iteration continue another loop or will it terminate here? (Circle one)

Continue / Terminate

6. (12 points, 2/5/5 pt) **Constraint Satisfaction Problems - CSP**

For the graph below, we would like to color each node with either Red, Green, or Blue, such that no two adjacent nodes have the same color.

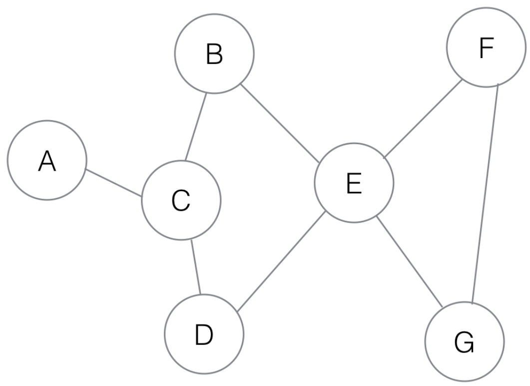

(a) According to the degree heuristic, which node would you consider first? Why?

(b) You may have noticed that this graph is nearly tree-structured. Recall that a nearly tree-structured CSP can be transformed into tree-structured CSPs by removing node(s) that form a cutset. (i) What is the smallest cutset for this problem? (ii) What is the computational complexity of the resulting problem? Determine the number of operations given that each variable has 3 values (no need to simplify your answer).

(c) Suppose we fix node C to be colored Red, D to be colored Blue, and F to be colored Green. Cross out illegal colors for all remaining nodes by checking arc consistency.

	A	B	C	D	E	F	G
Red				-		-	
Green			-	-			
Blue			-			-	

Additional space for notes.

Additional space for notes.