## Midterm

May 6, 2015

## Directions

This exam has 6 problems worth 84 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: $\qquad$
Number: $\qquad$

| 1 | $/ 20$ |
| ---: | ---: |
| 2 | $/ 20$ |
| 3 | $/ 10$ |
| 4 | $/ 10$ |
| 5 | $/ 8$ |
| 6 | $/ 16$ |
| Total | $/ 84$ |

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

Please circle the correct answer.
(a) Let $b$ be the branching factor of a search, $d$ the depth of the solution, and $m$ the maximum depth of the search space. Then the space complexity of breadth-first search is $b^{m}$.
(b) The time complexity of depth-first search is $b^{m} \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots$...............................
(c) To guarantee the optimality of $A^{*}$ search, the heuristic function must never overestimate the actual cost to reach the goal state.

T F
(d) Uniform-cost search is complete and optimal......................................... T F
(e) If $h$ is a consistent heuristic, then $h$ is also an admissible heuristic.

T F
(f) When using alpha-beta pruning, it is possible to get an incorrect value at the root node by choosing a bad ordering when expanding children.
(g) Least constraining value is a heuristic in constraint satisfaction
(h) Arc consistency detects failure earlier than forward checking...................... T F
(i) The running-time of an efficient solver for tree-structured constraint satisfaction problems is linear in the number of variables.
(j) The Bellman equation has a unique fixed point, and that fixed point can be computed via value iteration.


Given the graph above, show the order in which the states are visited by the search algorithms listed below. 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 and goal states in your answer (yYou can denote them as S and G, respectively). Treat the goal state as $G$ when you break ties. Finally, assume that all search algorithms execute the goal check when nodes are visited, not when their parent is expanded to create them as children.
(a) Depth First Search
(b) Breath First Search
(c) Iterative Deepening Depth First Search
(d) Uniform Cost Search
(e) A*Search
3. (10 points, $4 / 4 / 2 \mathrm{pt})$ CSP

We have a classic map-coloring problem, where we want to color each region of the map shown below. There are two constraints. First, we need to color the map such that two adjacent regions always have different colors. Second, we are also required to color region B with blue and region F with red.

(a) Given the map above, complete the constraint graph below by adding edges between states.

(b) On the table below, cross out values that violate arc-consistency by running the algorithm AC-3. The unary constraints on B and F are already applied for you.

|  | $A$ | $B$ | $C$ | $D$ | $E$ | $F$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Red |  | - |  |  |  |  |
| Green |  | - |  |  |  | - |
| Blue |  |  |  |  |  | - |

(c) Assume that we are using minimum remaining values (MRV) variable ordering. What is the first variable assigned? What is its value? Break ties using alphanumerical ordering (A comes before Z).
4. (10 points, 5pt each) 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 alphabeta pruning. Assume nodes are expanded left to right. You should draw an X through any such edge(s).

5. (8 points, 4pt each) Expectimax
(a) Fill in the values for each of the interior nodes in the following expectimax search tree. Assume a uniform distribution over actions for circular (expectimax) nodes. Ignore the scissors while filling out these values.
(b) What prior constraint would need to exist on possible values of the overall utility function such that you could prune the edge indicated by the scissors? Please give the tightest possible bound(s) as your answer.

6. (16 points, 5/2/5/4pt) Markov Decision Processes

The following problems take place in various scenarios of a gridworld MDP.
In all cases double-rectangle states are exit states. From an exit state, the only action available is Exit, which results in the listed immediate reward and ends the game (by moving into a terminal state; not shown).
From non-exit states, the agent can choose either Left $(L)$ or Right $(R)$ actions, which move the agent in the corresponding direction. There are no living rewards; the only non-zero rewards come from exiting the grid. Throughout this problem, assume that value iteration begins with initial values $V_{0}(s)=0$ for all states.

Consider the following scenario:

(a) Let the discount factor be $\gamma=0.5$, and transitions deterministic. Fill in the missing values for each state following the value iteration algorithm in the following table:

| $t$ | $A$ | $B$ | $C$ | $D$ | $E$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 0 | 0 | 0 |
| 1 |  |  |  |  |  |
| 2 |  |  |  |  |  |
| 3 |  |  |  |  |  |

(b) What is the optimal policy for states $B, C$, and $D$ according to these value iteration results?
(c) Instead of finding the optimal policy, assume we want to do policy evaluation for the policy

$$
\pi(B)=L \quad \pi(C)=L \quad \pi(D)=R
$$

Note that there is no need to specify a policy for $A$ or $E$ as they are exit states.
Fill out the following tables of values for this policy:

| $t$ | $A$ | $B$ | $C$ | $D$ | $E$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 0 | 0 | 0 |
| 1 |  |  |  |  |  |
| 2 |  |  |  |  |  |
| 3 |  |  |  |  |  |

(d) Recalling that the discount factor must be in the range $0 \leq \gamma \leq 1$, for what range of values for $\gamma$ is the optimal policy $\pi^{*}(B)=R$ (move right)?

Additional space for notes.

Additional space for notes.

