4003-455 Artificial Intelligence, Spring 2013
Prof. Zanibbi

Midterm Examination

Name: _____

April 4, 2013, Duration: 60 minutes, Out of 50 points

Instructions

- **If you have a question, please remain seated and raise your hand.**
- Place all books and coats at the front of the exam room.
- This exam is closed book and notes - no 'cheat sheets' are permitted.
- No electronic devices (laptops, phones, etc.) may be used during the examination.
- You may write your answers using pen or pencil, and you may write on the backs of pages.
- Additional pages are provided at the back of the exam.

1. **(5 points) True-or-False**

T / F Randomness is often used to avoid local minima/maxima in local searches.

T / F If two variables A and B are conditionally independent given variable C, then $\mathbb{P}(A, B \mid C) = \mathbb{P}(A \mid C)\mathbb{P}(B \mid C)$.

T / F Alpha-beta search may return a different move than the minimax algorithm.

T / F The minimax algorithm was formalized by John von Neumann in 1944.

T / F Incremental search may be understood as enumerating possible action sequences until either 1) a goal state for the problem environment is obtained, and the corresponding sequence of actions returned as a solution, or 2) all possible action sequences have been considered.

T / F Admissible heuristics may be created using a more complex version of the original search problem.

T / F A *rational* agent is one which always selects and then executes the optimal (e.g. minimal cost) solution to a problem.

T / F The normalization for $\mathbb{P}(A) = \alpha < 0.01, 0.30, 0.05 >$ is $\mathbb{P}(A) =< 1, 30, 5 >$.

T / F All prior and conditional probabilities of interest for a probabilistic model may be computed from the *joint probability distribution table*.

T / F The term 'Artificial Intelligence' was first coined by John McCarthy at the Dartmouth workshop held in 1956.

2. **(6 points) History and Foundations**

(a) (6 points) Briefly describe the protocol for a *Turing test*. How did Turing suggest the test be used to measure whether a machine exhibits *'intelligent' behavior*?

3. **(20 points) Search**

(a) (2 points) Name the four key components of a formal *incremental* search problem definition.

(b) (6 points) Provide the runtime and space complexity for each of the following *tree search* algorithms (i.e. that don't remember visited states), in terms of b (branching factor), m (maximum search tree depth), and d (depth of the optimal solution). Also, identify which of these searches are *optimal*.

i. Iterative Deepening
Time: **Space:**

ii. Depth-First
Time: **Space:**

iii. Breadth-First
Time: **Space:**

iv. Uniform-Cost
Time: **Space:**

v. *Any* search algorithm, if search tree is finite and no solution exists.
Time: **Space:**

Optimal Search Algorithms (selected from i-iv above):

(c) (8 points) We need to search the *state space* below, starting from A, and trying to reach goal G. Draw the search tree(s) produced by each of the search algorithms listed below. **Assume that child states are visited in alphabetical order**.

i. Iterative deepening using graph search (i.e. remembering visited states) - show the tree at each depth!

ii. Greedy search, with estimated cost *decreasing* alphabetically by state name (e.g. $h(C) < h(B)$).

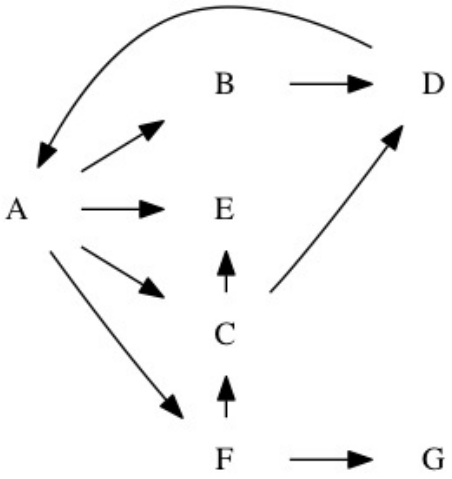

(d) (4 points) Define *admissible heuristic*, and then explain why or why not the search heuristic used for the greedy search is admissible.

4. **(7 points) Minimax**

(a) (5 points) For the game tree shown below, provide the minimax values for the internal nodes and the root of the tree, and then **indicate which action is the minimax action**.

(b) (2 points) Draw a line through the edges of the game tree that would be skipped when using the alpha-beta pruning algorithm.

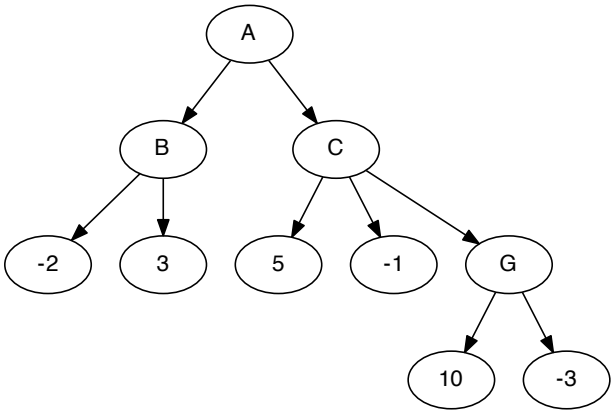

5. **(5 points) Probability**

(a) (5) Consider the *joint probability distribution below*, representing probabilities that a consumer purchases a particular drink from brand X or Y. There are three variables, Brand (brand X or brand Y), Type of drink (cola or tea), and Temperature (hot or cold).

	brand X		brand Y	
	tea	cola	tea	cola
hot	2/16	1/16	3/16	1/16
cold	2/16	3/16	1/16	3/16

i. (2) Is type of drink independent of brand? Why or why not?

ii. (3) Compute the *distribution* $\mathbb{P}(Type \mid Temperature = cold)$ from the table.

6. **(7 points) Decision Trees**

Suppose that we are creating a decision tree that will predict whether a consumer will purchase *tea* or *cola* based on drink brand and temperature.

(a) (2) Based on the joint probability distribution table in the previous question, which is the better *attribute* variable to use at the root of the tree, and why?

(b) (5) Identify the three base cases for the decision tree induction algorithm - give both the situation under which the case applies, and what decision is made at the node for that case.

i. Case 1 (Condition and Decision):

ii. Case 2 (Condition and Decision):

iii. Case 3 (Condition and Decision):

(c) **(+1 Bonus):** Describe the conventional structure of a *training data set* used to construct a decision tree.

[**Additional Space**]