CSCI-630 Foundations of Intelligent Systems
Spring 2014, Prof. Zanibbi

Midterm Examination

Name: ____

March 21, 2014, Duration: 50 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 The Turing test is the accepted standard for identifying machines that may be considered 'intelligent.'

T / F The availability of 'big data' has led to an unexpected trend toward using simpler decision models, and combinations of simple decision models (*ensembles*). This is partly due to an increased accuracy in probabilistic estimates obtained from large, representative samples.

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

T / F For discrete variables, all probabilities of interest may be computed from the *joint probability distribution table*.

T / F The basic operations used to produce the next generation in a genetic algorithm are selection, crossover, and mutation.

T / F A heuristic for a search problem A is said to *dominate* another heuristic B if in every reachable state s, $A(s) \leq B(s)$.

T / F Intelligent agents, which may be comprised of a combination of classifiers, regressors, search and other algorithms may be represented by a *function* from percepts (i.e. inputs) to chosen actions (i.e. outputs).

T / F The probability distribution $\mathbb{P}(A) = \alpha < 0.4, 0.8, 0.1 >$ normalizes to $\mathbb{P}(A) =< 0.2, 0.4, 0.05 >$.

T / F $P(a, b, c) = P(a|b, c)P(b|c)P(c)$.

T / F Supervised learning algorithms are evaluated using a training data set.

2. **(6 points) History**

(a) (4 points) Briefly describe the setup and procedure for executing a *Turing test*.

(b) (2 points) What does it mean for an agent to be *rational*?

3. **(18 points) Search**

(a) (6 points) Provide the runtime complexity, space complexity, and fringe data structure (queue type) for each of the following *tree search* algorithms, in terms of b (branching factor), m (maximum search tree depth), and d (depth of the optimal solution). Finally, identify which of these searches are *optimal*.

i. Bidirectional (using Breadth-First search)
Time: **Space:** **Queue:**

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

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

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

v. *Any* tree search algorithm, if the state space is finite and no solution exists.

Time:

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

(b) (2 points) Name the four components of an *incremental* search problem definition.

(c) (6 points) Name three local search algorithm that use randomness to avoid local minima during search, and indicate how randomness is used in each search algorithm.

Search Algorithm	**Use of Randomness**
1.	
2.	
3.	

(d) (4 points) We need to search the *state space* below, starting from C, and trying to reach goal G. Draw the search trees produced by iterative deepening using graph search (i.e. remembering visited states). Assume that child states are visited in alphabetical order.

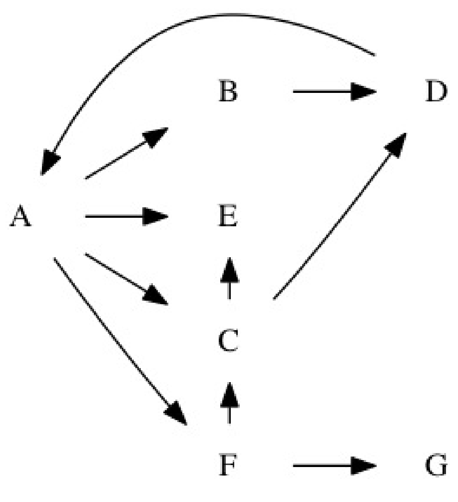

4. **(6 points) Minimax**

(a) (4 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.

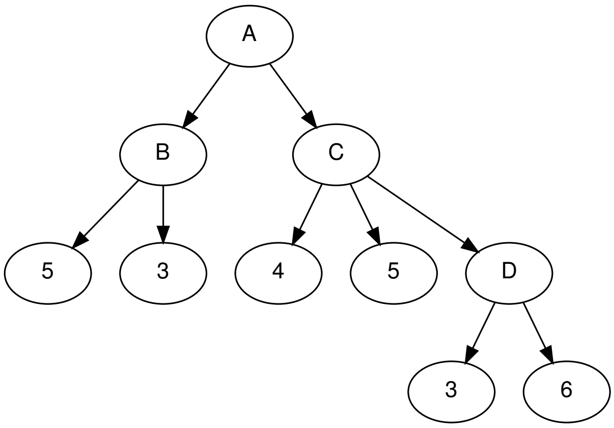

5. **(5 points) Probability**

(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. (3) Is Brand independent of Temperature? Why or why not?

ii. (2) Compute the *distribution* $\mathbb{P}(Brand \mid Type = tea)$ from the table.

6. **(10 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) Show how to compute the *entropy* of a probability distribution $\mathbb{P}(V) = \langle 0.2, 0.7, 0.1 \rangle$ for a discrete variable V with three possible values. **You do not need to compute the final value.**

(b) (8) Identify the three base and recursive cases for the decision tree induction algorithm. Describe the tree node created in each case.

i. Base Case 1 (Condition and Node Created):

ii. Base Case 2 (Condition and Node Created):

iii. Base Case 3 (Condition and Node Created):

iv. Recursive Case (Node Created):

(+1 Bonus)

T / F All approaches to creating intelligent systems studied in the course so far have involved *human* intelligence.

[**Additional Space**]