CSCI-630 Foundations of Intelligent Systems
Fall 2015, Prof. Zanibbi

Midterm Examination

Name: ____

October 16, 2015. 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 term 'Artificial Intelligence' was first coined by John McCarthy for the Dartmouth workshop held in 1956.

T / F The average for a list of numbers is also the *expected value* when the numbers in the list have a uniform probability.

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

T / F If $P(a,b|c) = P(a|c)P(b|c)$, then a is conditionally independent of b given c.

T / F For two-player zero sum games, the minimax algorithm is optimal against any opponent.

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

T / F Research in neural networks decreased significantly after the book *Perceptrons* was published in 1969, until the (re-)discovery of the backpropagation learning rule in the mid-1980's.

T / F Admissible heuristics for A* may be created by computing the solution to a more complex version of the original search problem.

T / F Over the course of AI's history, models used for analysis and selecting actions have been progressively more complex up to the present day.

T / F Random restart hill climbing, simulated annealing and genetic algorithms use randomness to increase *exploration* of the state space, to avoid local maxima/minima reached by greedily *exploiting* the best available action at every state.

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

(a) (4 points) In class we discussed four approaches to studying artificial intelligence. Name them, and identify which approach is being emphasized in this course.

(b) (2 points) Name the IBM researcher who developed a checkers-playing program in the 1950's that used a minimax-like algorithm, along with a learning algorithm to tune weights in the program's board evaluation function.

3. **(18 points) Search**

(a) (5 points) Provide the worst-case 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).

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

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

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

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

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

Time:

(b) (1 point) Identify which algorithms in part (a) are *optimal*.

(c) (3 points) Name the four components of an *incremental* search problem definition. Then identify how components change for 1) *game* search for turn-based games (e.g. tic-tac-toe), and 2) *local* search. Use '+' for added problem components, and '-' for removed components (e.g. '+name' and '–name').

	Incremental	**Game**	**Local**
1.			
2.			
3.			
4.			

(d) (4 points) Explain how 'parents' are selected for combination in a *genetic algorithm*.

(e) (5 points) We need to search the *state space* below, starting from D, and trying to reach goal G. Draw the search trees produced by iterative deepening using tree search (i.e. **not** remembering visited states). **Child states are visited in alphabetical order.**

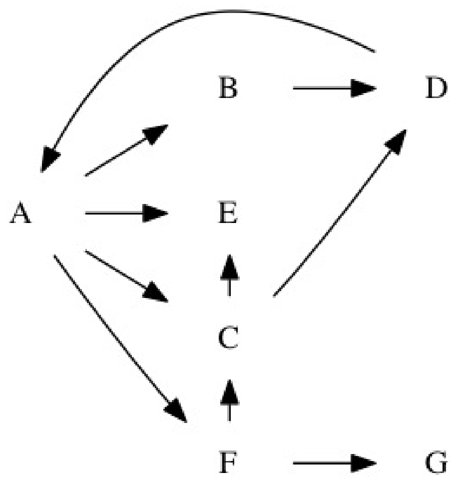

4. **(5 points) Minimax**

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

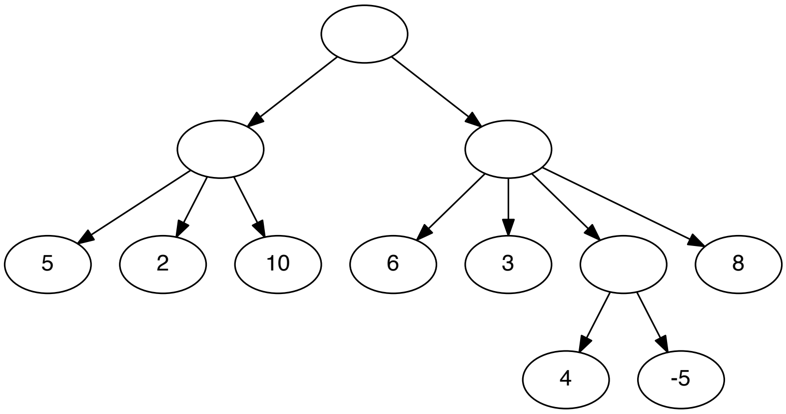

5. **(6 points) Probability**

(a) Consider the *joint probability distribution below*, representing probabilities that a consumer purchases a particular sandwhich from shop X or Y. There are three variables, Shop (shop X or shop Y), Type of sandwhich (cucumber or cheese), and Temperature (hot or cold).

	shop X		shop Y	
	cucumber	cheese	cucumber	cheese
hot	1/16	3/16	2/16	4/16
cold	2/16	2/16	1/16	1/16

i. (1) How many independent entries are there in this table?

ii. (3) Is Type of sandwhich independent of Shop? Why or why not?

iii. (2) Compute the *distribution* $\mathbb{P}(Type \mid Temperate = cold)$ from the table.

6. **(10 points) Miscellaneous Topics**

(a) (3) Explain the difference between an *informed* and an *uninformed* search, and name one example for each type of search algorithm.

(b) (2) What advantages of depth-first and breadth-first search are combined in iterative deepening?

(c) (5) In class, we talked about being able to represent joint probability distribution tables using combinations of *conditional* and *prior* probability tables, and how independence assumptions may be used to reduce the number of table entries that need to be defined.

We are given three binary variables, A, B and C, and that $\mathbb{P}(A|B, C) = \mathbb{P}(A)$. Using the *product* rule, we get $\mathbb{P}(A, B, C) = \mathbb{P}(A|B, C)\mathbb{P}(B|C)\mathbb{P}(C)$. Use this to convert the joint probability table to the smallest possible combination of conditional and prior probability tables, and **identify the number of independent entries**. Sketch your smaller tables, indicating clearly which distribution each belongs to.

Finally, explain how to use your tables to compute the probability $P(a, b, \neg c)$.

Bonus (+1)

T / F Greedy search algorithms are unpopular in practice because they are often sub-optimal.

[**Additional Space**]