

# CSCI-630 Foundations of Intelligent Systems

## Fall 2016, Prof. Zanibbi

### Midterm Examination

Name: \_\_\_\_\_

October 21, 2016. Duration: 50 minutes, Out of 50 points

### Instructions

- If you have a question, please remain seated and raise your hand.
- After leaving the exam room, you may not return for the duration of the exam. If you finish your exam with five minutes or less remaining, remain in your seat and wait until the end of the exam.
- 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 Despite producing potentially suboptimal results, variations of hill climbing are frequently used for optimization, due to the infinite size of metric spaces (e.g.,  $\mathbb{R}^2$ ).
- T / F In their book ‘Perceptrons,’ Minsky and Papert demonstrated that a regular perceptron is unable to learn the XOR function. This nearly killed neural net research for a decade.
- 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 *conditional probability distribution table*.
- 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 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 Overfitting in decision trees is a side-effect of obtaining smaller and smaller samples after repeated splitting during training.
- T / F The minimax algorithm was formalized by John McCarthy in 1944.
- T / F If no solution to a search problem exists, then the time complexity for breadth-first, iterative deepening, and A\* search will be the same.
- T / F A *rational* agent is one which always selects and then executes the optimal (e.g. minimal cost) solution to a problem.

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

- (a) (4 points) Name the four components of an agent problem, or ‘task environment’ (Hint: PEAS).

- (b) (2 points) Name the Nobel prize-winning economist and AI pioneer who devised the notion of ‘satisficing.’

3. (18 points) Search

(a) (4 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. Uniform Cost  
**Time:** \_\_\_\_\_ **Space:** \_\_\_\_\_ **Queue:** \_\_\_\_\_
- ii. Depth-First  
**Time:** \_\_\_\_\_ **Space:** \_\_\_\_\_ **Queue:** \_\_\_\_\_
- iii. Iterative Deepening  
**Time:** \_\_\_\_\_ **Space:** \_\_\_\_\_ **Queue:** \_\_\_\_\_
- iv. Breadth-First  
**Time:** \_\_\_\_\_ **Space:** \_\_\_\_\_ **Queue:** \_\_\_\_\_

(b) (2 points) When do we need to convert a *tree search* to a *graph search* to avoid infinite loops and redundant searches?

(c) (6 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) Search the *state space* below, starting from *C* 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.**



- (e) (2 points) For A\* search to obtain optimal results, the *heuristic* function estimating the distance to a goal must be *admissible*. What is the strategy discussed in the text and lecture for creating admissible heuristics for a search problem?

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 sandwich from shop X or Y. There are three variables, Shop (shop X or shop Y), Type of sandwich (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
others cold	2/16	2/16	1/16	1/16

- i. (1) How many independent entries are there in this table?
- ii. (3) Is Shop independent of Temperature? Why or why not?
- iii. (2) Compute the *distribution*  $\mathbb{P}(\text{Shop} \mid \text{Type} = \text{cucumber})$  from the table.

6. (10 points) Machine Learning

(a) These questions relate to decision trees used for classification.

i. (3) Provide the **three base cases** for the decision tree induction algorithm, at which point we stop splitting training samples, and create a leaf node.

ii. (2) Why is *information gain* used to select among attributes for splitting a node during decision tree learning?

iii. (2) Decision trees have a tendency to overfit training data, in the worst case ‘memorizing’ training examples while extracting few predictive patterns. We looked at two techniques for preventing this problem in class - name them.

(b) Recall that AdaBoost is a binary classification *ensemble* algorithm.

i. (1) Which types of binary classifiers may be combined using AdaBoost?

ii. (2) During training, how do we select the next classifier to add to the ensemble?

**Bonus (+1)**

Provide the formula for computing the *entropy* of a probability distribution.

[ Additional Space ]



