

# Final Exam

600.335/435 Artificial Intelligence  
Fall 2015

**Name:**

**Section (335/435):**

## Instructions

Please be sure to write both your name and section in the space above! Some questions will be exclusive to grad students and will be marked as such. Please be sure to read through the entire exam before you start, and be mindful of the time. If one question is taking too long, it may be worth moving on and coming back to the problem question(s) after the rest of the exam is complete. Remember that you are only allowed one sheet (both sides) of notes, everything else besides that and the test itself must be off of your workspace. Please show ALL relevant work for your answers and provide explanations when prompted. Failure to do either will result in loss of credit. You will have until 12 noon to complete this test.

## Intelligent Agents

1. (4 pts) Consider an email spam filter as an intelligent agent. Every 10 seconds the agent runs a script, 'check.py', that checks the user's inbox and returns unread messages (or nothing). For all unread messages, another script is then run called, 'filter.py', which classifies each unread email as either spam or non-spam and then sends that email to either the inbox or the spam folder, depending on the classification. Give an example of a **sensor**, a **percept**, an **effector**, and an **action** for this agent, and then give an example of a function that maps from the agent's percepts to an action. The function should include two possible mappings that map two different percepts to two different actions (i.e. it should be bijective).

2. (6 pts) A student is developing a robot that can roll a die. The robot grips and releases the die with an arm that can also twist, and uses two cameras that can see the entire rolling area at once for feedback. The goal is to develop methods that can roll the die in a way that increases the chance of landing on a desired side of the die. Please consider the environment that is the robot, die, and rolling area.


Is this environment (justify your answers!):

- Accessible?
  
  
  
  
  
  
  
  
  
  
- Deterministic?
  
  
  
  
  
  
  
  
  
  
- Episodic?
  
  
  
  
  
  
  
  
  
  
- Static?
  
  
  
  
  
  
  
  
  
  
- Discrete?


3. (2 pts) Now consider the case that the robot from the previous question is being programmed to roll dice in a competitive game. A group of robots are all collected into one area, and each given a die. If a robot rolls a die that brings the sum of rolled values to an even number, that robot is eliminated from the running. These games are played again and again with the robots that haven't been eliminated until only one winner remains. The rolling area has been expanded to accommodate the other robots, and as a result the robot now has to sweep this larger area in order to see all of the already rolled dice. In addition, the robots don't have to take turns. The robot should be constantly strategizing over what kind of roll to perform depending on the values of the other dice that have already been rolled.

Which of the answers to the previous question have been changed, and why?

## Basic Search




4. (2 pts) Fill in the nodes of the above tree in the order they would be explored in a depth-first search (assume left to right traversal).
5. (6 pts) Compare depth-first searching to breadth-first searching. When are depth-first searches the superior choice and why? Give a real-world example.
6. (6 pts) Give the time and space complexity of depth-first searching in terms of its branching factor,  $b$ , and search depth  $d$ . Is this algorithm optimal? Why/why not?

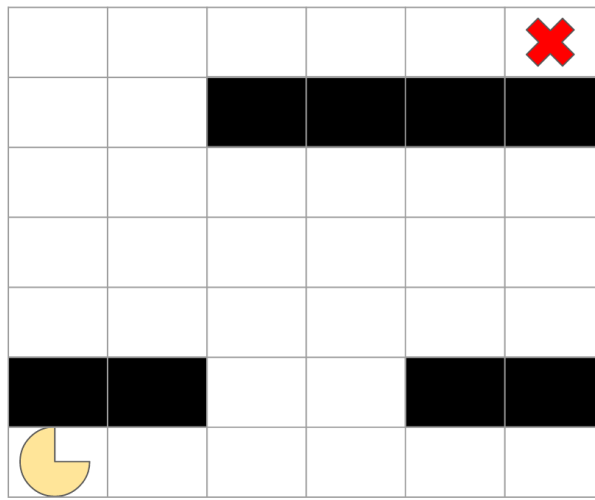


7. (2 pts) Fill in the nodes of the above tree in the order they would be explored in an iterative deepening search (again, assume left to right traversal).
8. (6 pts) Compare iterative deepening searching to breadth-first searching. When are iterative deepening searches the superior choice and why? Give a real-world example.
9. (6 pts) Give the time and space complexity of iterative deepening searching in terms of its branching factor,  $b$ , and search depth  $d$ . Is this algorithm optimal? Why/why not?

### Informed Search




10. (4 pts) Given the above pacman search problem, please show the order of node expansion for a uniform-cost search. Assume west, north, east, south order of traversal.



11. (6 pts) Given the above pacman search problem, please show the order of node expansion for an A\* search with manhattan distance as the heuristic. Assume west, north, east, south order of traversal.
12. (6 pts) Is the manhattan distance function considered admissible? Why or why not? Why is admissibility important for A\* searching?
13. (3 pts) Given two admissible heuristics  $h_1$  and  $h_2$ , circle the following heuristics that are also admissible (note: there may be more than one!).
  - (a)  $\max(h_1, h_2)$
  - (b)  $h_1 + h_2$
  - (c)  $\max(h_1, 2 * h_2)$
  - (d)  $\min(h_1, 3 * h_2)$
  - (e)  $\frac{h_1+h_2}{2}$
14. (2 pts) Prove (briefly)  $\max_{next}(|x_{current} - x_{next}|, |y_{current} - y_{next}|)$  is an admissible heuristic if you are navigating the movement of the king piece on a chessboard with obstacles, moving from the bottom left to the top right.

## Game Playing

15. (6 pts) I am trying to develop a method that will determine the best chess move for any given turn. The way I do this is by constructing a tree, where each node is a configuration of the board and each branch is a different valid move from the parent node's configuration. I fully expand this tree with all possible valid moves to each possible end state (and mark it as a victory or loss). I then simply find the next move that has the greatest possibility of leading to a favorable win state by counting the number of "victory" leaf nodes that each move could lead me towards and going with the move that has the greatest number. There is a problem with this method. Please specify what the writer of this method is misunderstanding, and what modifications to the algorithm he/she could make to avoid this problem.



16. (4 pts) The above tree represents a two-player game where each player takes turns making a move. The first move is made by us, the second is made by the opposing player, and so on. Circle the nodes in this tree that are explored according to a minmax algorithm.
17. (6 pts) Put a strike through the nodes that would be pruned according to  $\alpha - \beta$  pruning.
18. (4 pts) Tic-Tac-Toe is a game that is considered "trivially solved." Please explain what it means when a game is "solved." What is it about the structure of tic-tac-toe that makes this game trivial to solve? Please be specific.

## First Order Logic

19. (6 pts) Simplify the following propositional logic statement as much as possible and translate into english (hint- remove implications):

$$(A \implies B) \wedge (\neg A \implies C)$$

20. (3 pts) Translate the following english sentence into first-order logic:

Not all those who wander are lost

21. (4 pts) Is the following statement valid? What about satisfiable?

$$(A \implies B) \iff (B \implies A)$$

22. (8 pts) Please examine the following statements for this question. Using resolution, prove that statement 1 follows from statement 2. You should either complete the proof (if possible), or continue until the proof breaks down and you cannot continue (if impossible). Show the unifying substitution in each step. If the proof fails, explain where, how, and why it fails.

**1**  $\exists y \forall x (x \geq y)$

**2**  $\forall x \exists y (x \geq y)$

## Probabilistic Reasoning

23. (2 pts) Suppose we flip a fair coin 3 times. What is the probability that all three land on heads?
24. (2 pts) Now suppose we again flip a fair coin 3 times, but this time we know that 2 of them have landed on heads. What is the probability that all three land on heads?
25. (4 pts) Given the probability space of 3 coin flips, describe an event that is disjoint from flipping 3 heads in a row and show why.

For the next three questions, suppose we have two continuous random variables,  $X, Y$ , whose joint density function is defined by:

$$f(x, y) = \begin{cases} \frac{1}{\pi} & \text{if } x^2 + y^2 \leq 1 \\ 0, & \text{otherwise} \end{cases}$$

26. (2 pts) What's the probability that  $X \geq 0 \wedge Y \geq 0$ ?
27. (4 pts) Are these two events independent? Explain why or why not.
28. (4 pts) Are the two events  $X \geq a$  and  $Y \geq b$ , where  $a, b \in \mathbb{R}$  and  $-1 \leq a \leq b \leq 1$ , independent for arbitrary  $a$  and  $b$ ? Explain why or why not (e.g. give a counterexample)?



## Bayesian Networks


29. (3 pts) Draw a Bayesian network that contains 3 variables where the factorization of their joint probability distribution is

$$P(A, B, C) = P(B | A)P(C | A)P(A)$$


30. (3 pts) Draw a Bayesian network that contains 3 variables where the factorization of their joint probability distribution is

$$P(A, B, C) = P(C | A, B)P(A)P(B)$$

31. (3 pts) Write the factorization of the joint probability density of the following Bayesian network:

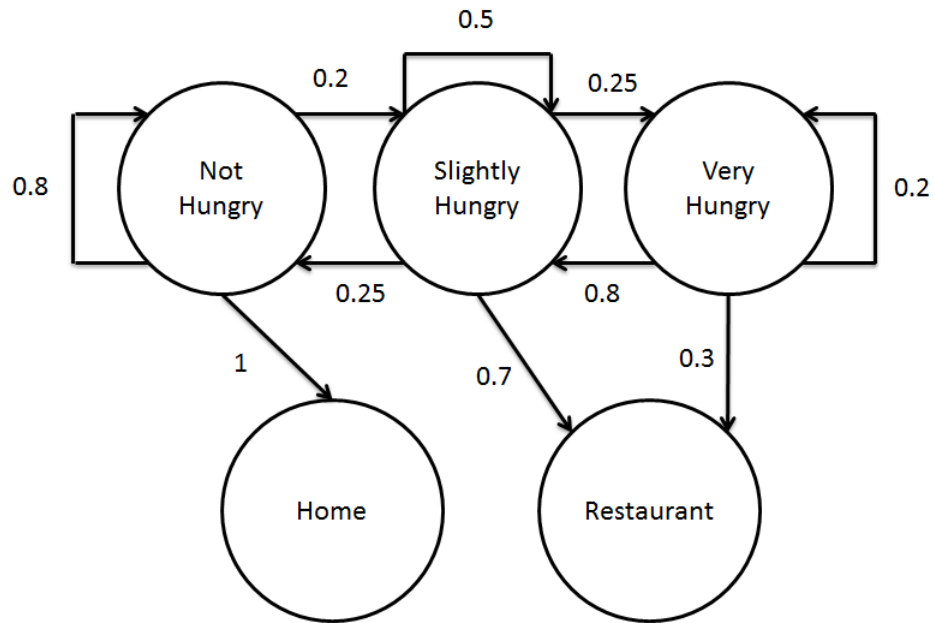


32. (3 pts) Write the factorization of the joint probability density of the following Bayesian network:



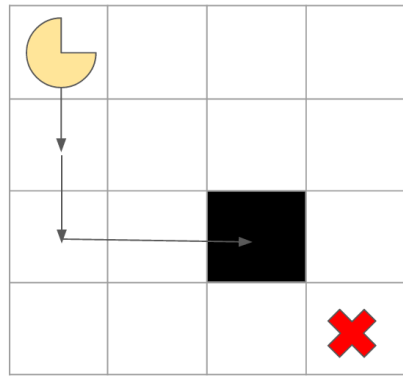
## Markov Decision Processes

Use the hidden markov model below for the following two questions.




33. (10 pts) Examine the hidden markov model. Please describe what it is trying to represent, and what information about this system its structure and probabilities tell you.
34. (4 pts) Given that the hidden state is "not hungry" at time  $t = 1$ , what is the probability that the hidden state is "slightly hungry" at time  $t = 3$ ?


# Reinforcement Learning



Path 1



Path 2



Path 3

The above picture illustrates an icy grid that our agent has stochastic movement within. Note that 60% of the time the agent is able to make its desired motion successfully, but 40% of the time the agent slips and accidentally moves an extra space in the same direction. Suppose we want to use reinforcement learning to learn how to best control our agent in such a situation. We will be examining both passive and active reinforcement learning methods. If the goal state is reached, our agent has completed its task successfully. If the agent enters the hole, however, it is considered a failure.

- 35. (5 pts) In both the passive and active reinforcement learning methods, we will need to define a reward function  $R(s)$ . Please state the purpose of such a function, and define one for this grid. Remember to factor in both reaching end states and normal movements within the grid.
  
- 36. (5 pts) Examine path 1 as shown in the figure. Note that this path results in a failure state. If we update a utility-based passive agent and then a q-learning agent both using this same run, what exactly is updated in each case and how? Please define the purpose of the functions that are modified. No explicit formulas are needed.

37. (6 pts) Say that these agents have just gone through path 2 as shown in the figure. Note that this path results in the goal state. Based on this run, what exactly is updated in each case and how? How are these updates different than the ones based on a failed run?
38. (7 pts) Say that these agents have just gone through path 3 as shown in the figure. Note that this path results in the goal state and has some overlap with path 2. Based on this run, what exactly is updated in each case and how? Please examine the state that path 2 and path 3 have in common and note that they reach the goal after leaving this state with a different number of steps. How does this fact affect the way that the agents' functions are updated based solely on this one state they have in common?