Name: ________________________

CS 331 Final
Spring 2017

You have 110 minutes to complete this exam. You are only allowed to use your textbook, your notes, your assignments and solutions to those assignments during this exam. If you find that you are spending a large amount of time on a difficult question, skip it and return to it when you've finished some of the easier questions. Total marks for this exam is 72.

Section	Marks
Pre-Midterm	/ 32
Probability	/ 10
Bayesian Networks	/ 30
Total	/ 72

Section I: Pre-Midterm questions [32 points]

1. Circle true or false below each question.
a) There exist environments in which a simple-reflex agent must have randomized actions in order to be rational. [2 points]

True False

b) The Othello agent from programming assignment 2 operated in a static environment. [2 points]

True False

c) The model in a model-based agent can use propositional logic. [2 points]

True False

d) The only difference between breadth-first search and depth-first search is that BFS uses a FIFO queue and DFS uses a LIFO queue. [2 points]

True False

e) A* search with g(n)=0 for all nodes n is optimally efficient. [2 points]

True False

f) All consistent heuristics are also admissible. [2 points]

True False

g) The biggest advantage of IDDFS over BFS is a better time complexity. [2 points]

True False

h) Local search is a good strategy for learning the structure of Bayes nets because it is guaranteed to find the optimal DAG. [2 points]

True False

i) Genetic algorithms can be interpreted as a form of gradient descent. [2 points]

True False

j) The expectiminimax algorithm avoids worst-case scenarios. [2 points]

True False

k) Alpha-beta pruning can be applied with both evaluation functions and utility functions. [2 points]

True False

l) The minimax algorithm is suitable for strategic environments. [2 points]

True False

m) The resolution algorithm operates on sentences in disjunctive normal form. [2 points]

True False

n) A knowledge base consisting of only the sentence False entails all sentences in propositional logic. [2 points]

True False

o) If a sentence P in propositional logic is valid, then $\neg P$ is unsatisfiable. [2 points]

True False

p) $(P \implies Q) \wedge (Q \implies P) \equiv P \vee Q.$ [2 points]

True False

Section II: Probability [10 points]

2. In this problem, we will consider the famous probability puzzle called the "Monty Hall problem." Imagine that you are a contestant on a game show, in which you must select one of three doors and receive the item behind the door. Behind one of the doors is a new car (the prize to be won), and behind the other two doors are goats. After you make an initial selection of which door to open, the host of the game show opens one of the two remaining doors. The host knows where the car is, and he must open a door to a goat (i.e. he will never reveal the car). After showing you a goat, the host asks whether you would like to switch from your initial selection to the other remaining door. The puzzle is: is it to your advantage to switch?

Let $C \in \{1,2,3\}$ be a random variable for the location of the car, $D \in \{1,2,3\}$ be a random variable for the door you initially choose, and $H \in \{1,2,3\}$ be a random variable for the door that the host opens for you. Suppose $P(C) = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. You may assume anything you wish about $P(D)$.

a) What is the probability that the car is behind door 1, given that you initially select door 1? (Hint: C and D are independent.) [2 points]

b) What is the probability that the host opens door 3, given the position of the car and that you selected door 1? (Hint: you need to specify a value for each position of the car.) [2 points]

c) To decide whether to switch, you need to reason about $P(C = c|D = 1, H = 3)$. Use the chain rule (or equivalently, the definition of conditional probability), to write down a symbolic expression for this probability. (Hint: The joint distribution $P(C, D, H) = P(C)P(D)P(H|C, D)$.) [2 points]

d) Use your solutions above to compute numeric values for $P(C = 1|D = 1, H = 3)$ and $P(C = 2|D = 1, H = 3)$. [2 points]

e) Having initially chosen door 1 and then being shown a goat behind door 3, should you switch to door 2? [2 points]

Section III: Bayesian Networks [30 points]

3. You are investigating the effects of two treatments on a disease. You collect patients with both low and high severity cases that received the treatments. For each patient, you record the severity of their case (low or high), the treatment they received (a or b), and whether or not they recover. The data you collect looks like this:

Severity (S)	**Treatment (T)**	**Recover (R)**	**Number of cases in dataset**
low	a	true	9
low	a	false	1
low	b	true	23
low	b	false	5
high	a	true	19
high	a	false	7
high	b	true	5
high	b	false	2

a) First, you decide to look at the overall effects of the treatments. Use this dataset to estimate the conditional probabilities in the table for R conditional on T. Use Uniform Dirichlet priors as in Programming Assignment #3. Which treatment is better? [5 points]

T	**R**	**P(R|T)**
a	true	
a	false	
b	true	
b	false	

b) Next, you focus on the severity of the cases. Use the dataset to estimate the conditional probabilities in the table for R conditional on S and T. This time, you only compute the set of independent parameters for $R = true$, knowing that the probabilities for $R = false$ in each case can be computed from these as needed. Use Uniform Dirichlet priors as in Programming Assignment #3. Which treatment is better if you have a low-severity case? Which treatment is better if you have a high-severity case? [5 points]

S	**T**	**R**	**P(R|S,T)**
low	a	true	
low	b	true	
high	a	true	
high	b	true	

c) "Simpson's paradox" tells us that it is possible for one treatment to look better in the subgroups and another treatment to look better in the whole dataset. One factor that can contribute to this phenomenon occurs when the severity of the case determines which treatment is given. Draw a Bayes Net in which severity of the case determines the treatment, which in turn determines recovery. In your network, severity should be independent of recovery given treatment. [2 points]

4. Consider the following Bayes net, with variables for the season (summer or winter) and binary variables indicating whether it is raining, whether the sprinkler is on, whether the pavement is wet, and whether the pavement is slippery. For the binary variables, only the independent parameters are supplied in the CPTs; the rest of the values can be inferred from these.

S	**P(S)**
summer	0.5
winter	0.5

S	**R**	**P(R|S)**
winter	true	0.8
summer	true	0.1

S	**P**	**P(P|S)**
winter	true	0.05
summer	true	0.6

R	**P**	**W**	**P(W|R,P)**
true	true	true	0.99
true	false	true	0.95
false	true	true	0.75
false	false	true	0.1

W	**L**	**P(L|W)**
true	true	0.7
false	true	0.2

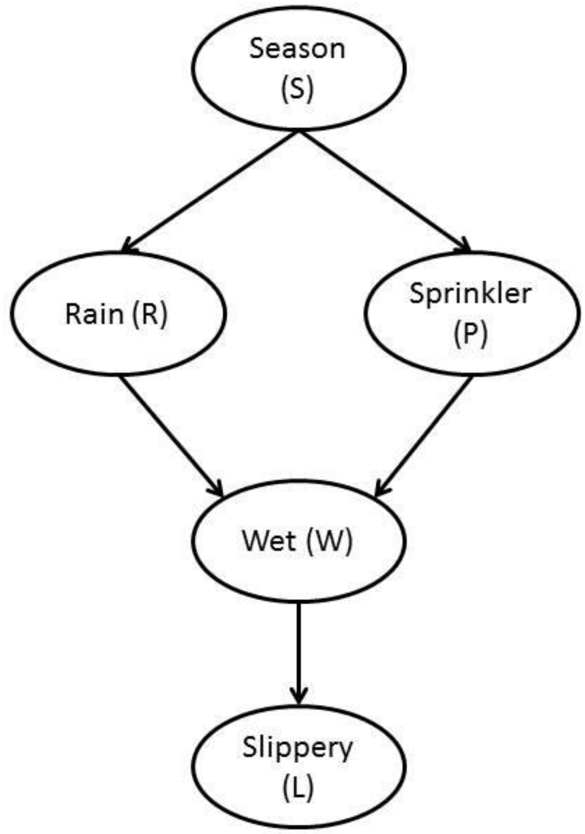

Using these CPTs and the Bayes net structure, compute the joint probability that it is summer ($S = summer$), the sprinkler is off ($P = false$), and the pavement is wet ($W = true$). Show your work. [10 points]

5. Use the Bayesian network below to determine whether or not the following conditional independence relationships hold or not. Show the blocked/unblocked paths for partial credit.

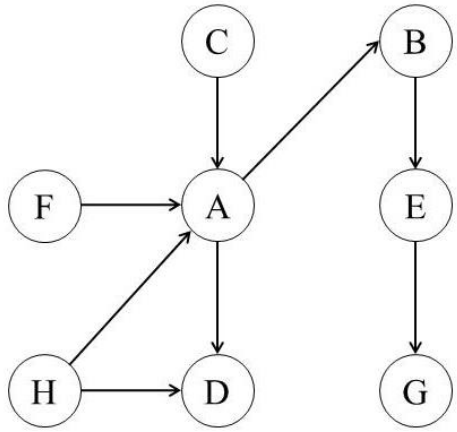

a) I(H,B|A) [2 points]

b) I(H,B|D) [2 points]

c) I(H,B|{A,D}) [2 points]

d) I(D,G|A) [2 points]