Final Preparation.

The final will test **all** material, not only material after the second midterm. Notice that the following questions are a sample of the kind of questions that could be on the final. However, they certainty do not cover everything that could be on the final. They are merely intended to give you some idea of the kind of questions that you can expect. In general, you probably want to look at all of the assignments (including sample solutions), handouts, slides, the book, and all midterm prep questions. Please do not forget to bring a calculator to the final.

After your yearly checkup, the doctor has bad news and good news. The bad news is that you tested positive for a serious disease, and that the test is 99% accurate (that is, the probability of testing positive given that you have the disease is 0.99, as is the probability of testing negative given that you don't have the disease). The good news is that this is a rare disease, striking only one in 10,000 people. Why is it good news that the disease is rare? What are the chances that you actually have the disease?

Consider the following Bayes net.

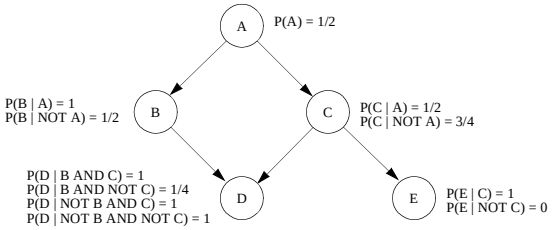

a) What is P(A AND B AND C AND D AND E)?
b) What is P(NOT A AND NOT B AND NOT C AND NOT D AND NOT E)?
c) What is P(C | A AND B)?
d) There are 32 different truth-value assignments to A, B, C, D, and E. List which of these 32 truth-value assignments have non-zero probability.

True or false: The amount of memory needed to represent a Bayes net is proportional to the number of nodes times the maximum number of parents of any node? Justify your answer.

You are an AI consultant for a credit card company, and your task is to construct a belief net that will alow the company to determine whether or not to grant a person a card.

a) What are the evidence variables? These are the variables for which you can obtain information, on the basis of which it is legal to make decisions, and that are relevant to the decision.

b) What is the output variable, that is, what proposition is the company going to examine the probabilities of in order to determine whether or not to grant a person a card?

c) Construct your network by incrementally adding variables in causal order. You may wish to add intermediate nodes such as Reliability and FutureIncome.

A used-car buyer can decide to carry out various tests with various costs, and then, depending on the outcome of the tests, decide which car to buy. We will assume that the buyer is deciding whether to buy car c1, that there is time to carry out at most one test, and that t1 is the test of c1 and costs $50. A car can be in good shape (quality q+) or bad shape (quality q-), and the test may help to indicate what shape the car is in. Car c1 costs $1,500, and its market value is $2,000 if it is in good shape; if not, $700 in repairs will be needed to make it in good shape. The buyer's estimate is that c1 has a 70% chance of being in good shape.

a) Calculate the expected net gain from buying c1, given no test.

b) Tests can be described by the probability that the car will pass or fail given that the car is in good or bad shape. We have the following information: P(pass(c1,t1)|q+(c1)) = 0.80 and P(pass(c1,t1)|q-(c1)) = 0.35. Use Bayes' theorem to calculate the probability that the car will pass (or fail) its test, and hence the probability that it is in good (or bad) shape given each possible test outcome.

c) Calculate the optimal decisions given either a pass or a fail, and their expected net gains.

You are an opera singer who is about to take a Chemical Engineering exam. Within the last two days running up to the exam you can be in one of three states of knowledgability about Chemical Engineering: ignorant, competent, or expert. Every day you must decide what you will do today: opera singing (O), or work on your chemical studies (W). If you pass the exam you will receive a one-time payment of $120. Every day that you do opera singing you will earn $25. You earn nothing for working on your chemical studies or failing the exam. As the following figure shows, you can be in 11 states. In six of the states you get to make the O or W decision. For these states, the next state will be a probabilistic choice with the probabilities dependent on your decision. These probabilities are depicted in the figure. You are motivated by money. You want to maximize the expected total (undiscounted) amount of money you will receive

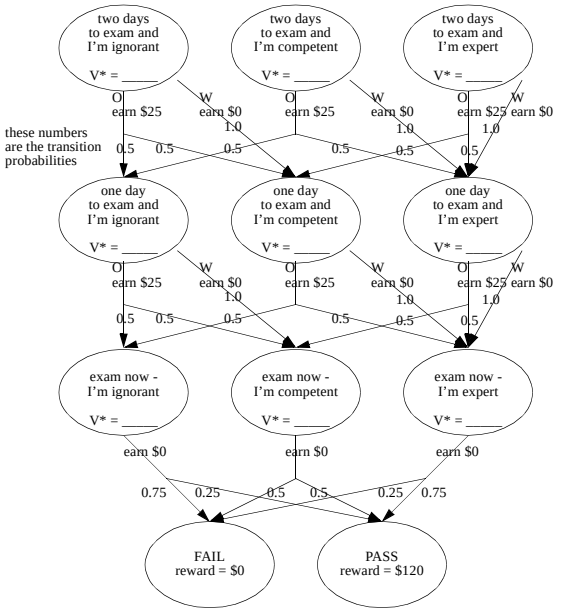

V*(state) is the largest expected total amount of money that you can possibly receive when you start in that state. For each state in the figure, fill in its V*(state) value. (Double-check your answer for arithmetic errors). What is the definition of a "policy"? Describe the optimal policy for the example (in English).

Consider the following Markov decision process model and, using (a) value-iteration and (b) policy iteration, determine a behavior that (approximately) minimizes the expected total (undiscounted) cost. The following question is only relevant if we get to discount factors: Does the answer chance if the task is to minimize the expected total discounted cost for discount factor 0.8?

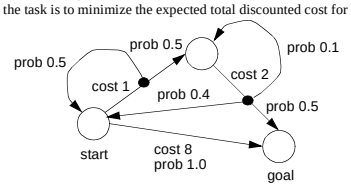

True or false: You can compute any conditional probability from a Bayes net in time logarithmic in the number of nodes? Justify your answer using your knowledge about the complexity of probabilistic inference in Bayes nets.