Midterm Preparation.

Notice that the following questions are a sample of the kind of questions that could be on the midterm. However, they certainty do not cover everything that could be on the midterm. 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, and the book.

TRUE of FALSE?

No search method that makes use of heuristic functions can GUARANTEE to find a shortest path from the start to a goal: true or false.

Breadth-first search always terminates if there is a path from the start to the goal: true or false.

Depth-first search always terminates if there is a path from the start to the goal: true or false.

Alpha-beta search never expands more nodes than minimax search: true or false.

Games like chess are usually solved with forward search techniques: true or false.

The order in which successor states are generated affects the number of node expansions of minimax search: true or false.

A* is an informed version of depth-first search, that is, a version of depth-first search that uses a heuristic function to bias the search towards the goal states: true or false.

Given two admissible heuristic functions, at least one dominates the other one: true or false.

Games like chess can be solved with A* by limiting the lookahead of A*. This is necessary because the game tree is extremely large: true or false.

It is possible to use A* with a heuristic function that assigns zero to every state: true or false.

Suppose that you have a search problem where you are more concerned with minimizing the search effort than with minimizing the solution cost. Which of the following f-values would you use in conjunction with the A* method and the pathmax equation? f(s) = 0, f(s) = g(s), or f(s) = h(s).

Depth-first search and uniform cost search behave identically if all actions have a cost of one, all else being equal: true or false.

In the following maze the successors of a cell include any cell directly to the east, south, west or north of the current cell except that no transition may pass through the central barrier. For exam-

ple successors(m) = {d, n, g}. The search problem is to find a path from s to g. We are going to examine the order in which cells are expanded by various search algorithms. For example, one possible expansion order that breadth-first search might use is: s h f k p c q a r b t d g. There are other possible orders depending on which of two equal-distance-from-start states happen to be expanded first. For example s f h p k c q r a t b g is another possible answer.

a) Assume you run depth-first search until it expands the goal node. Assume that you always try to expand East first, then South, then West, then North. Assume your version of depth-first search avoids loops: it never expands a state on the current path. What is the order of state expansion?

b) You decide to use a Manhattan-Distance heuristic function, where h(state) = shortest number of steps from state to g if there were no barriers. So, for example, h(k) = 2, h(s) = 4, and h(g) = 0. Is this heuristic function admissible? Is it consistent?

c) Assume that you use best-first greedy search using heuristic h (a version that never re-explores the same state twice). Again, give all the states expanded, in the order they are expanded, until the algorithm expands the goal state.

d) Finally, assume you use A* search with heuristic h, and run it until it terminates using the conventional A* termination rule. Again, give all the states expanded, in the order they are expanded. (Note that depending on the method that A* uses to break ties, more than one correct answer is possible).

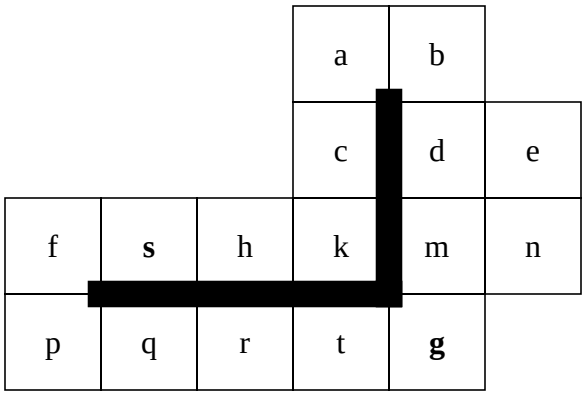

What are the minimax values of the root nodes of the following two game trees? Cross out the node(s), if any, whose value(s) the alpha-beta method never determines, assuming that it always evaluates the leftmost successor node first.

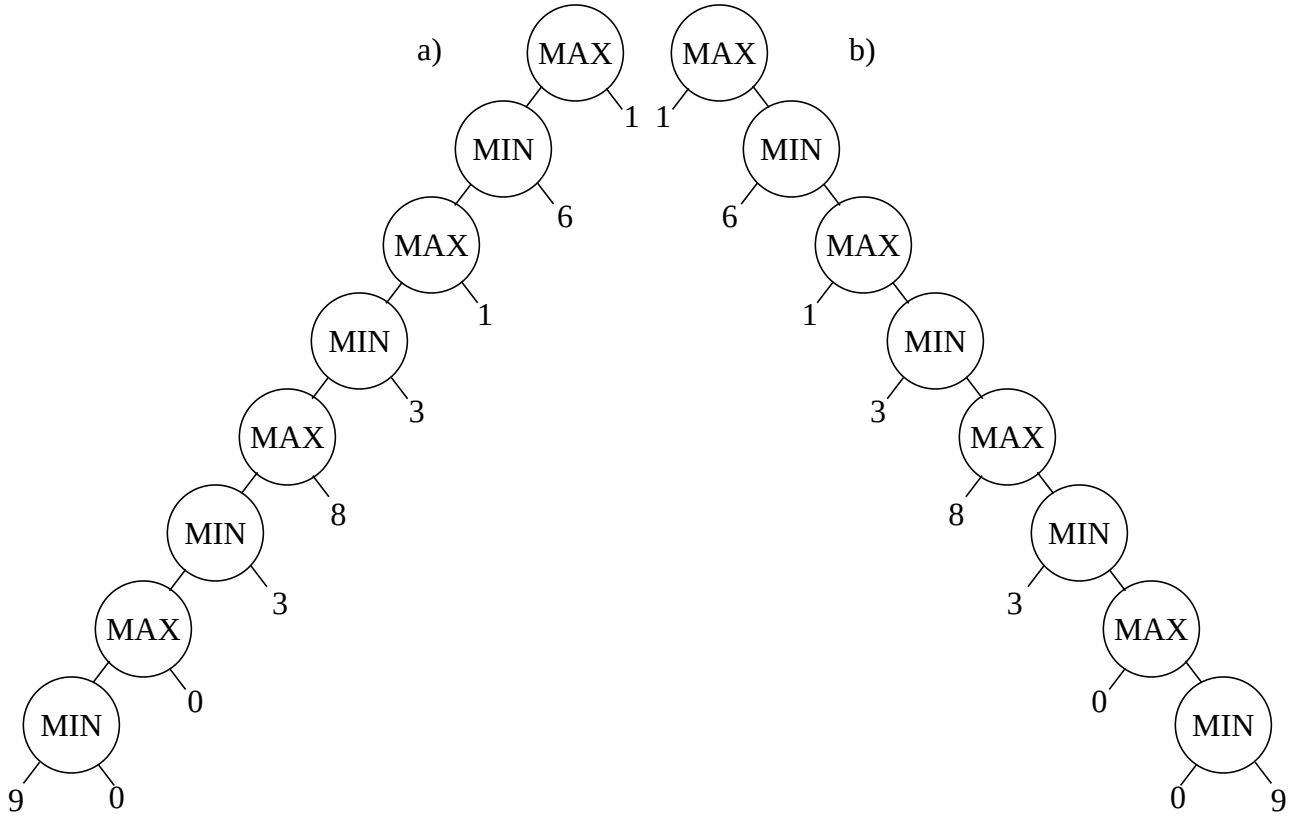

Consider the following probabilities.

P(PassExam AND WildParty) = 0.2
P(PassExam AND NOT WildParty) = 0.5
P(NOT PassExam AND WildParty) = 0.2
P(NOT PassExam AND NOT WildParty) = 0.1

a) What is P(PassExam)?
b) What is P(WildParty)?
c) What is P(WildParty OR PassExam)?
d) What is P(PassExam | WildParty)?
e) What is P(PassExam | NOT WildParty)?

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.

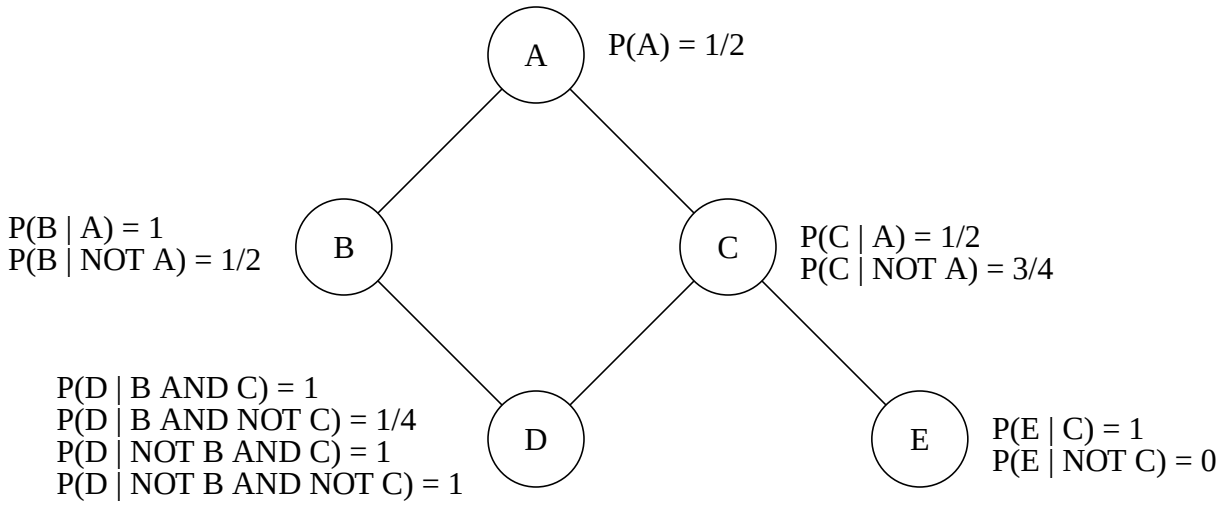

a) What is P(A AND B AND C AND D AND E)?
b) What is P(NOT A AND NOT B AND NOT C AND NOT D AND NOT E)?
c) What is P(C | A AND B)?
d) There are 32 different truth-value assignments to A, B, C, D, and E. List which of these 32 truth-value assignments have non-zero probability.

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.

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.

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.

d) Calculate the value of information of the test, and derive an optimal conditional plan for the buyer.

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

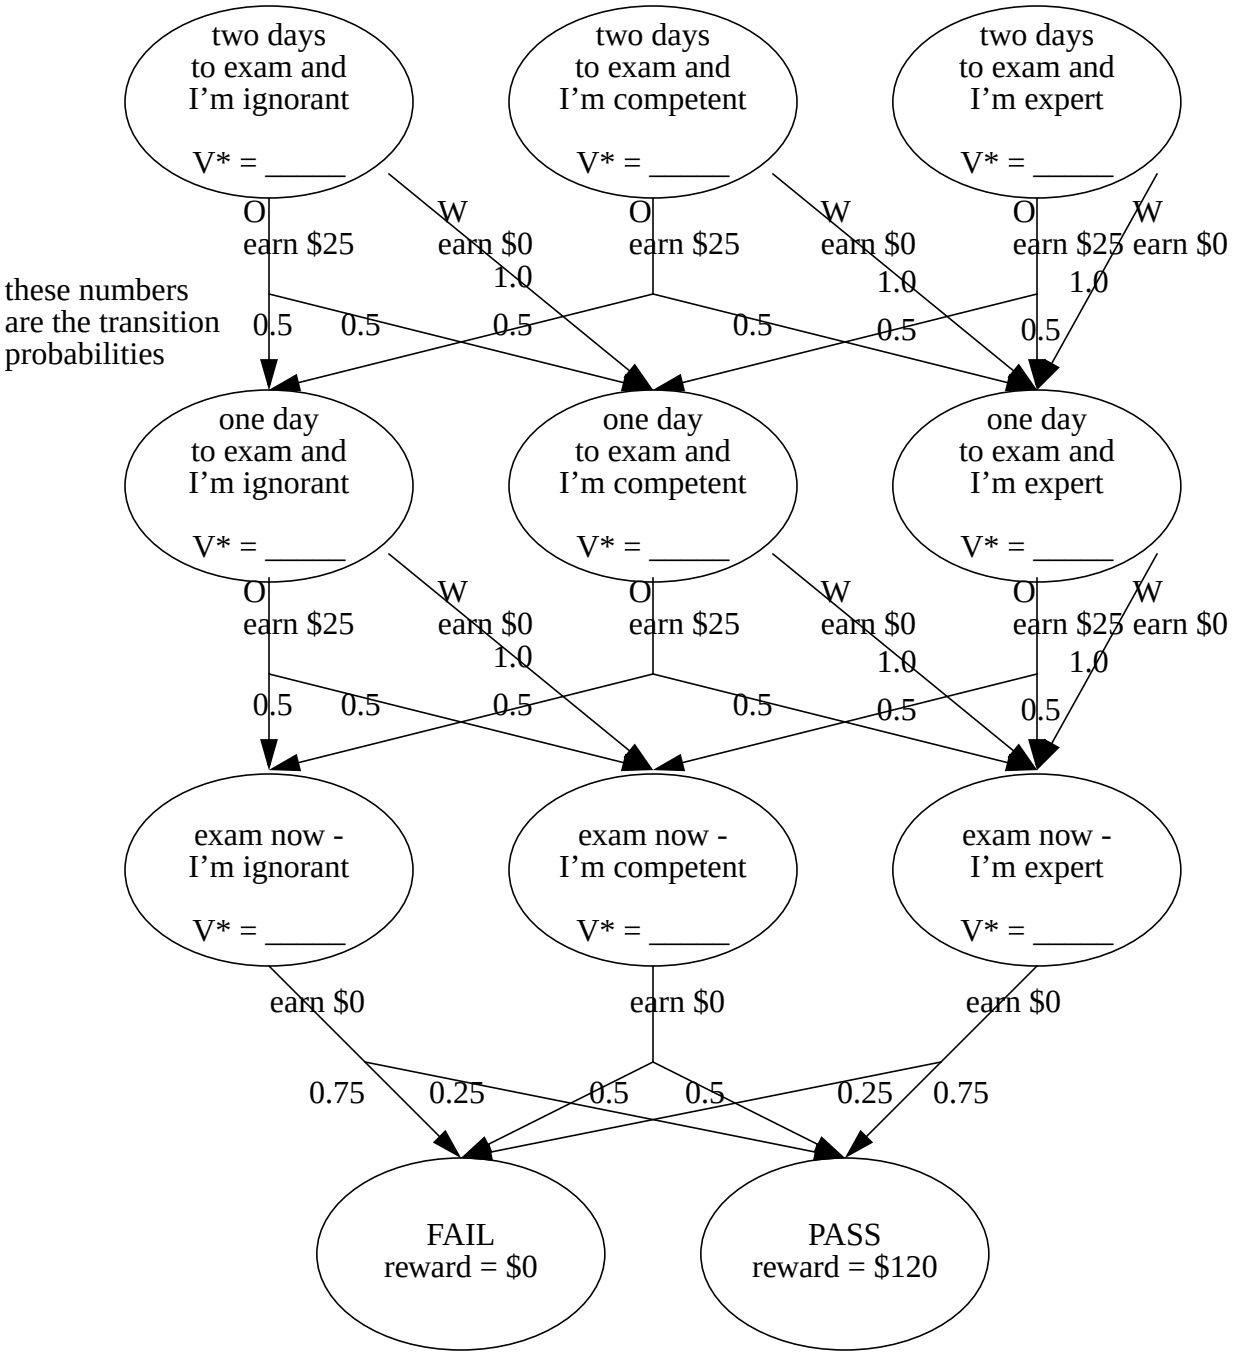

V*(state) is the largest expected total amount of money that you can possibly receive when you start in that state. For each state in the figure, fill in its V*(state) value. (Double-check your answer for arithmetic errors). What is the definition of a "policy"? Describe the optimal policy for the example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

Consider the following Markov decision process model and, using value-iteration, determine a behavior that (approximately) minimizes the expected total (undiscounted) cost.

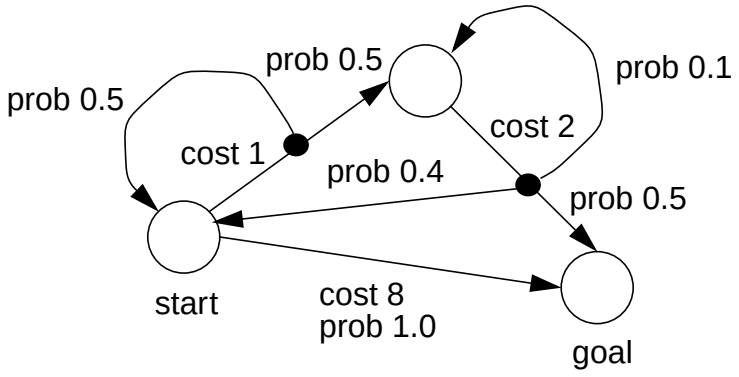

Explain the differences between:
a) situation space planners and plan space planners
b) progressive and regressive planners
c) total and partial order planners.

Model a simple world in STRIPS in which there are rooms with doors that can be either open (and unlocked), closed and locked, or closed and unlocked. There is an agent that can go to different rooms, pick up some or all of the keys that are in its current room, take them with it, use them to lock and unlock doors, and drop them anywhere. The agent cannot go through closed doors and thus might have to open and close them, as well as unlock and lock them.

Consider the following partial-order plan. How many operators does this plan have? Is this plan consistent? Is this plan complete? What are the preconditions of operators that are currently not satisfied (not achieved)? Does the dashed link denote a causal link or an ordering constraint that is not a causal link?

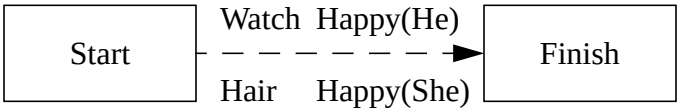

(Thanks to Andy Moore for providing some of the questions.)