


Please answer clearly and succinctly. Show your work clearly for full credit. If an explanation is requested, think carefully before writing. Points will be removed for rambling answers with irrelevant information (and may be removed in cases of messy and hard to read answers). If a question is unclear or ambiguous, feel free to make the additional assumptions necessary to produce the answer. State these assumptions clearly; you will be graded on the basis of the assumption as well as subsequent reasoning.

There are 10+0 problems worth 78 points on 7 pages

Problem 0 (1 point) Write your name on the top of each page.

Problem 1 (11 points) Suppose that an agent is in a 3x3 maze environment similar to the ones shown in the illustration below. The agent knows the size of the maze, that the initial state is (1,1) in the lower left and that the goal is (3,3). There are four actions, **Left**, **Right**, **Up** and **Down**, which have their usual effects, except when blocked by a wall or by the maze boundaries; in this case, the agent does not move. The agent does *not* (initially) know if there are any walls present in the maze nor where such walls may be. After each action, the agent receives a single percept, which tells it if it has failed to move; by using this percept, the agent may deduce the presence of walls.




- a) (5 points) Describe the problem of finding a path from the initial state to the goal as a search problem in belief state space.

- b) (3 points) How large is the *initial belief state*?

- c) (3 points) Approximately how many distinct belief states are there? (Some of the possible belief states may not be reachable from the initial state, but include them anyway).

Problem 2 (10 points) The figure below shows a problem-space graph, where A is the initial state and G denotes the goal. Edges are labeled with their true cost. We have a heuristic function, $f()$, written in the standard form: $f(n)=g(n)+h(n)$ where $g(n)$ is the cost to get from A to n and $h(n)$ is an estimate of the remaining distance to G.




- a) (2 points) In the graph above, is $f()$ admissible? Why or why not?

- b) (2 points) Is $f()$ monotonic? Why or why not?

c) (6 points) Suppose we use IDA* to search the graph and that states having the same f values are visited in alphabetical order. In what order does the algorithm consider f-limits and visit states? Fill out the table below:

f-limit	States visited (in order)
2	A

Problem 3 (8 points) Label the intermediate nodes in the search tree, below, with values and draw a line across edges to denote any pruning done by alpha-beta search. Which action should the agent choose?



Problem 4 (4 points) Encode “All Germans speak the same language” in first-order logic using the following notation: **Speaks(p,l)** denotes that person p speaks language l; **Nationality(p, c)** denotes that person p is from country c; the constant **G** denotes Germany.

Problem 5 (4 points)

Which of the following are semantically and syntactically correct translations of “Everyone’s zipcode within a state has the same first digit”?

$$\text{i. } \forall x, s, z_1 [State(s) \wedge LivesIn(x, s) \wedge Zip(x) = z_1] \Rightarrow [\forall y, z_2 LivesIn(y, s) \wedge Zip(y) = z_2 \Rightarrow Digit(1, z_1) = Digit(1, z_2)].$$

$$\text{ii. } \forall x, s [State(s) \wedge LivesIn(x, s) \wedge \exists z_1 Zip(x) = z_1] \Rightarrow [\forall y, z_2 LivesIn(y, s) \wedge Zip(y) = z_2 \wedge Digit(1, z_1) = Digit(1, z_2)].$$

$$\text{iii. } \forall x, y, s State(s) \wedge LivesIn(x, s) \wedge LivesIn(y, s) \Rightarrow Digit(1, Zip(x)) = Digit(1, Zip(y)).$$

$$\text{iv. } \forall x, y, s State(s) \wedge LivesIn(x, s) \wedge LivesIn(y, s) \Rightarrow Digit(1, Zip(x)) = Digit(1, Zip(y)).$$

Answer	i	yes	no
	ii	yes	no
	iii	yes	no
	iv	yes	no

Problem 6 (4 points) After your yearly checkup, your doctor has bad news and good news. The bad news is that you tested positive for a serious disease is 99% accurate (ie the probability of testing positive when you do have the disease is 0.99 as is the probability of testing negative when you don’t have the disease). The good news is that this is a rare disease, striking only 1 in 10,000 people of your age. Use Bayesian reasoning to calculate the chance that you actually have the disease?

Problem 7 (4 points) Suppose you are given a sack containing n unbiased coins. You are told that $n-1$ of the coins are normal with heads on one side and tails on the other. One coin in the sack is fake, having heads on both sides. Suppose that you reach into the sack, pick a coin uniformly at random, flip it twice, and get heads both times. What is the (conditional) probability that you picked the fake coin?

Problem 8 (13 points) Suppose an agent inhabits a world with two states, S and $\neg S$, can do exactly one of two actions, a and b . Action a does nothing and action b flips from one state to the other. Suppose that S is initially true.

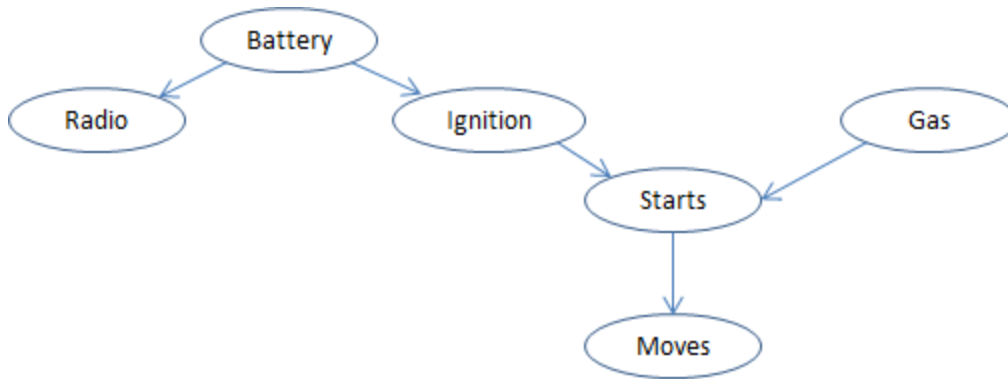
- a) (6 points) Draw the complete planning graph (with mutex relations) for the initial propositional level, the first action level and the next propositional level.

- b) (6 points) Consider this world as an MDP (with deterministic actions) and let $R(S)=3$, $R(\neg S)=2$ and $\gamma=0.5$. Complete the columns of the following table. (Note we are assuming that reward is a function of the destination state and is independent of the starting state or action executed.)

Time	$Q(S,a)$	$Q(S,b)$	$V(S)$	$Q(\neg S,a)$	$Q(\neg S,b)$	$V(\neg S)$
0			0			0
1	3					
2						
3						


- c) (1 point) Based on the V_3 values you have computed, what is the policy an agent should choose?

Problem 9 (7 points) Consider the Bayesian network shown below.




- (1 points) Is Radio conditionally independent of Gas given Battery?
- (1 points) Is Radio conditionally independent of Gas given Starts?
- (1 points) Is radio conditionally independent of Gas given Moves?
- (4 points) Give reasonable probability tables for all the nodes (draw them by the nodes in the network, above)


Problem 10 (12 points) Two astronomers in different parts of the world make measurements, M_1 and M_2 , of the number of stars, N , in some small region of the sky, using their telescopes. Normally, there is a small possibility ϵ of error in either direction so $\epsilon/2$ chance of being one too high and $\epsilon/2$ chance of being one too low. Each telescope can also (with a much smaller probability f) be badly out of focus (events F_1 and F_2), in which case the scientist will undercount by three or more stars (or if N is less than 3, fail to detect any stars at all). Consider the three networks shown below:



(a)



(b)



(c)

a) (6 points) Which of these Bayesian Networks are correct (but not necessarily efficient) representations of the preceding information?

b) (1 points) Which is the best network (explain)?

c) (5 points) Write out a conditional distribution for $P(M_1 | N)$ for the case where $N \in \{1, 2, 3\}$ and $M_1 \in \{0, 1, 2, 3, 4\}$. Each entry in the conditional distribution should be expressed as a function of the parameters ϵ and/or f .