
CS 188
Fall 2011

Introduction to
Artificial Intelligence

Midterm Exam

INSTRUCTIONS

- You have 3 hours.
- The exam is closed book, closed notes except a one-page crib sheet.
- Please use non-programmable calculators only.
- Mark your answers **ON THE EXAM ITSELF**. If you are not sure of your answer you may wish to provide a *brief* explanation. All short answer sections can be successfully answered in a few sentences at most.

Last Name	
First Name	
SID	
Login	
<i>All the work on this exam is my own.</i> (please sign)	


For staff use only

Q. 1	Q. 2	Q. 3	Q. 4	Q. 5	Q. 6	Q. 7	Q. 8	Total
/12	/12	/13	/7	/7	/11	/12	/6	/80

THIS PAGE INTENTIONALLY LEFT BLANK

1. (12 points) Search

Consider the search graph shown below. S is the start state and G is the goal state. All edges are bidirectional.



For each of the following search strategies, give the path that would be returned, or write *none* if no path will be returned. If there are any ties, assume alphabetical tiebreaking (i.e., nodes for states earlier in the alphabet are expanded first in the case of ties).

- (a) (1 pt) Depth-first graph search

- (b) (1 pt) Breadth-first graph search


- (c) (1 pt) Uniform cost graph search

- (d) (1 pt) Greedy graph search

- (e) (2 pt) A* graph search


For the following question parts, all edges in the graphs discussed have cost 1.

- (f) (3 pt) Suppose that you are designing a heuristic h for the graph on the right. You are told that $h(F) = 0.5$, but given no other information. What ranges of values are possible for $h(D)$ if the following conditions must hold? Your answer should be a range, e.g. $2 \leq h(D) < 10$. You may assume that h is nonnegative.



- i. h must be admissible
- ii. h must be admissible and consistent

- (g) (3 pt) Now suppose that $h(F) = 0.5$, $h(E) = 1.1$, and all other heuristic values except $h(B)$ are fixed to zero (as shown on the right). For each of the following parts, indicate the range of values for $h(B)$ that yield an admissible heuristic AND result in the given expansion ordering when using A* graph search. If the given ordering is impossible with an admissible heuristic, write *none*. Break ties alphabetically. Again, you may assume that h is nonnegative.



- i. B expanded before E expanded before F
- ii. E expanded before B expanded before F

2. (12 points) Formulation: Holiday Shopping

You are programming a holiday shopping robot that will drive from store to store in order to buy all the gifts on your shopping list. You have a set of N gifts $G = \{g_1, g_2, \dots, g_N\}$ that must be purchased. There are M stores, $S = \{s_1, s_2, \dots, s_M\}$ each of which stocks a known inventory of items: we write $g_k \in s_i$ if store s_i stocks gift g_k . Shops may cover more than one gift on your list and will never be out of the items they stock. Your home is the store s_1 , which stocks no items.

The actions you will consider are travel-and-buy actions in which the robot travels from its current location s_i to another store s_j in the fastest possible way and buys whatever items remaining on the shopping list that are sold at s_j . The time to travel-and-buy from s_i to s_j is $t(s_i, s_j)$. You may assume all travel-and-buy actions represent shortest paths, so there is no faster way to get between s_i and s_j via some other store. The robot begins at your home with no gifts purchased. You want it to buy all the items in as short a time as possible and return home.

For this planning problem, you use a state space where each state is a pair (s, u) where s is the current location and u is the set of unpurchased gifts on your list (so $g \in u$ indicates that gift g has not yet been purchased).

(a) (1 pt) How large is the state space in terms of the quantities defined above?

(b) (4 pt) For each of the following heuristics, which apply to states (s, u) , circle whether it is admissible, consistent, neither, or both. Assume that the minimum of an empty set is zero.

- | | |
|--|---|
| (neither / admissible / consistent / both) | The shortest time from the current location to any other store:
$\min_{s' \neq s} t(s, s')$ |
| (neither / admissible / consistent / both) | The time to get home from the current location:
$t(s, s_1)$ |
| (neither / admissible / consistent / both) | The shortest time to get to any store selling any unpurchased gift:
$\min_{g \in u} (\min_{s': g \in s'} t(s, s'))$ |
| (neither / admissible / consistent / both) | The shortest time to get home from any store selling any unpurchased gift:
$\min_{g \in u} (\min_{s': g \in s'} t(s', s_1))$ |
| (neither / admissible / consistent / both) | The total time to get each unpurchased gift individually:
$\sum_{g \in u} (\min_{s': g \in s'} t(s, s'))$ |
| (neither / admissible / consistent / both) | The number of unpurchased gifts times the shortest store-to-store time:
$ u (\min_{s_i, s_j \neq s_i} t(s_i, s_j))$ |

You have waited until very late to do your shopping, so you decide to send an swarm of R robot minions to shop in parallel. Each robot moves at the same speed, so the same store-to-store times apply. The problem is now to have all robots start at home, end at home, and for each item to have been bought by at least one robot (you don't have to worry about whether duplicates get bought). Hint: consider that robots may not all arrive at stores in sync.

(c) (4 pt) Give a minimal state space for this search problem (be formal and precise!)

One final task remains: you still must find your younger brother a stuffed Woozle, the hot new children's toy. Unfortunately, no store is guaranteed to stock one. Instead, each store s_i has an initial probability p_i of still having a Woozle available. Moreover, that probability drops exponentially as other buyers scoop them up, so after t time has passed, s_i 's probability has dropped to $\beta^t p_i$. You cannot simply try a store repeatedly; once it is out of stock, that store will stay out of stock. Worse, you only have a single robot that can handle this kind of uncertainty! Phrase the problem as a single-agent MDP for planning a search policy for just this one gift (no shopping lists). You receive a single reward of +1 upon successfully buying a Woozle, at which point the MDP ends (don't worry about getting home); all other rewards are zeros. You may assume a discount of 1.


(d) (3 pt) Give a minimal state space for this MDP (be formal and precise!)

3. (13 points) CSPs: Trapped Pacman

Pacman is trapped! He is surrounded by mysterious corridors, each of which leads to either a pit (P), a ghost (G), or an exit (E). In order to escape, he needs to figure out which corridors, if any, lead to an exit and freedom, rather than the certain doom of a pit or a ghost.

The one sign of what lies behind the corridors is the wind: a pit produces a strong breeze (S) and an exit produces a weak breeze (W), while a ghost doesn't produce any breeze at all. Unfortunately, Pacman cannot measure the strength of the breeze at a specific corridor. Instead, he can stand *between* two adjacent corridors and feel the max of the two breezes. For example, if he stands between a pit and an exit he will sense a strong (S) breeze, while if he stands between an exit and a ghost, he will sense a weak (W) breeze. The measurements for all intersections are shown in the figure below.

Also, while the total number of exits might be zero, one, or more, Pacman knows that two neighboring squares will *not* both be exits.



Pacman models this problem using variables X_i for each corridor i and domains P, G, and E.


(a) (3 pt) State the binary and/or unary constraints for this CSP (either implicitly or explicitly).

(b) (4 pt) Cross out the values from the domains of the variables that will be deleted in enforcing arc consistency.

X_1	P	G	E
X_2	P	G	E
X_3	P	G	E
X_4	P	G	E
X_5	P	G	E
X_6	P	G	E

- (c) (1 pt) According to MRV, which variable or variables could the solver assign first?
- (d) (1 pt) Assume that Pacman knows that $X_6 = G$. List all the solutions of this CSP or write *none* if no solutions exist.

The CSP described above has a circular structure with 6 variables. Now consider a CSP forming a circular structure that has n variables ($n > 2$), as shown below. Also assume that the domain of each variable has cardinality d .




- (e) (2 pt) Explain precisely how to solve this general class of circle-structured CSPs efficiently (i.e. in time linear in the number of variables), using methods covered in class. Your answer should be at most two sentences.
- (f) (2 pt) If standard backtracking search were run on a circle-structured graph, enforcing arc consistency at every step, what, if anything, can be said about the worst-case backtracking behavior (e.g. number of times the search could backtrack)?

4. (7 points) Games: Multiple Choice

In the following problems please choose **all** the answers that apply. You may circle more than one answer. You may also circle no answers, if none apply.

- (a) (2 pt) In the context of adversarial search, α - β pruning
- (i) can reduce computation time by pruning portions of the game tree
 - (ii) is generally faster than minimax, but loses the guarantee of optimality
 - (iii) always returns the same value as minimax for the root of the tree
 - (iv) always returns the same value as minimax for all nodes on the leftmost edge of the tree, assuming successor game states are expanded from left to right
 - (v) always returns the same value as minimax for all nodes of the tree
- (b) (2 pt) Consider an adversarial game in which each state s has minimax value $v(s)$. Assume that the maximizer plays according to the optimal minimax policy π , but the opponent (the minimizer) plays according to an unknown, possibly suboptimal policy π' . Which of the following statements are true?
- (i) The score for the maximizer from a state s under the maximizer's control could be greater than $v(s)$.
 - (ii) The score for the maximizer from a state s under the maximizer's control could be less than $v(s)$.
 - (iii) Even if the opponent's strategy π' were known, the maximizer should play according to π .
 - (iv) If π' is optimal and known, the outcome from any s under the maximizer's control will be $v(s)$.
- (c) (3 pt) Consider a very deep game tree where the root node is a maximizer, and the complete-depth minimax value of the game is known to be v_∞ . Similarly, let π_∞ be the minimax-optimal policy. Also consider a depth-limited version of the game tree where an evaluation function replaces any tree regions deeper than depth 10. Let the minimax value of the depth-limited game tree be v_{10} for the current root node, and let π_{10} be the policy which results from acting according to a depth 10 minimax search at every move. Which of the following statements are true?
- (i) v_∞ may be greater than or equal to v_{10} .
 - (ii) v_∞ may be less than or equal to v_{10} .
 - (iii) Against a perfect opponent, the actual outcome from following π_{10} may be greater than v_∞ .
 - (iv) Against a perfect opponent, the actual outcome from following π_{10} may be less than v_∞ .

5. (7 points) MDPs: Bonus level!



Pacman is in a bonus level! With no ghosts around, he can eat as many dots as he wants. He is in the 5×1 grid shown. The cells are numbered from left to right as $1, \dots, 5$. In cells 1 through 4, the actions available to him are to move *Right* (R) or to *Fly* (F) out of the bonus level. The action *Right* deterministically lands Pacman in the cell to the right (and he eats the dot there), while the *Fly* action deterministically lands him in a terminal state and ends the game. From cell 5, *Fly* is the only action. Eating a dot gives a reward of 10, while flying out gives a reward of 20. Pacman starts in the leftmost cell (cell 1).

We write this as an MDP where the state is the cell that Pacman is in. The discount is γ .

Consider the following 3 policies:

$$\pi_0(s) = F \text{ for all } s$$

$$\pi_1(s) = R \text{ if } s \leq 3, F \text{ otherwise}$$

$$\pi_2(s) = R \text{ if } s \leq 4, F \text{ otherwise}$$

(a) (4 pt) Assume $\gamma = 1.0$. What is:

i. $V^{\pi_0}(1)$?

ii. $V^{\pi_1}(1)$?

iii. $V^{\pi_2}(1)$?

iv. $V^*(1)$?

(b) (3 pt) Now consider an arbitrary value for γ .

i. Does there exist a value for γ such that π_0 is strictly better than both π_1 and π_2 ? If yes, give a value for γ . If no, write *none*.

ii. Does there exist a value for γ such that π_1 is strictly better than both π_0 and π_2 ? If yes, give a value for γ . If no, write *none*.

iii. Does there exist a value for γ such that π_2 is strictly better than both π_0 and π_1 ? If yes, give a value for γ . If no, write *none*.

6. (11 points) MDPs and RL: Return to Blackjack

Armed with the power of Q -learning, you return to the CS188 casino! In this question, you will play a simplified version of blackjack where the deck is infinite and the dealer always has a fixed count of 15. The deck contains cards 2 through 10, J , Q , K , and A , each of which is equally likely to appear when a card is drawn. Each number card is worth the number of points shown on it, the cards J , Q , and K are worth 10 points, and A is worth 11. At each turn, you may either *hit* or *stay*. If you choose to hit, you receive no immediate reward and are dealt an additional card. If you stay, you receive a reward of 0 if your current point total is exactly 15, +10 if it is higher than 15 but not higher than 21, and -10 otherwise (i.e. lower than 15 or larger than 21). After taking the *stay* action, the game enters a terminal state *end* and ends. A total of 22 or higher is referred to as a *bust*; from a bust, you can only choose the action *stay*.

As your state space you take the set $\{0, 2, \dots, 21, \text{bust}, \text{end}\}$ indicating point totals, “bust” if your point total exceeds 21, and “end” for the end of the game.

- (a) (3 pt) Suppose you have performed k iterations of value iteration. Compute $V_{k+1}(12)$ given the partial table below for $V_k(s)$. Give your answer in terms of the discount γ as a variable. Note: do not worry about whether the listed V_k values could actually result from this MDP!

s	$V_k(s)$
13	2
14	10
15	10
16	10
17	10
18	10
19	10
20	10
21	10
bust	-10
end	0

$$V_{k+1}(12) =$$

You suspect that the cards do not actually appear with equal probability and decide to use Q -learning instead of value iteration.

- (b) (4 pt) Given the partial table of initial Q -values below, fill in the partial table of Q -values on the right after the following episode occurred. Assume a learning rate of 0.5 and a discount factor of 1. The initial portion of the episode has been omitted. Leave blank any values which Q -learning does not update.

Initial values

s	a	$Q(s, a)$
19	hit	-2
19	stay	5
20	hit	-4
20	stay	7
21	hit	-6
21	stay	8
bust	stay	-8

Episode

s	a	r	s	a	r	s	a	r
19	hit	0	21	hit	0	bust	stay	-10

Updated values

s	a	$Q(s, a)$
19	hit	
19	stay	
20	hit	
20	stay	
21	hit	
21	stay	
bust	stay	

Unhappy with your experience with basic Q -learning, you decide to featurize your Q -values, representing them in the form $\sum_i w_i f_i(s, a)$ for some feature functions $f_i(s, a)$.

First, consider the two feature functions

$$f_1(s, a) = \begin{cases} 0, & \text{if } a = \text{stay}; \\ +1, & \text{if } a = \text{hit and } s \geq 15; \\ -1, & \text{if } a = \text{hit and } s < 15; \end{cases} \quad \text{and} \quad f_2(s, a) = \begin{cases} 0, & \text{if } a = \text{stay}; \\ +1, & \text{if } a = \text{hit and } s \geq 18; \\ -1, & \text{if } a = \text{hit and } s < 18. \end{cases}$$

(c) (3 pt) Circle all of the following partial policy tables for which it is possible to represent Q -values in the form $w_1 f_1(s, a) + w_2 f_2(s, a)$ that imply that policy unambiguously (i.e., without having to break ties).

(i)		(ii)		(iii)		(iv)		(v)	
s	$\pi(s)$	s	$\pi(s)$	s	$\pi(s)$	s	$\pi(s)$	s	$\pi(s)$
14	hit	14	stay	14	hit	14	hit	14	hit
15	hit	15	hit	15	hit	15	hit	15	hit
16	hit	16	hit	16	hit	16	hit	16	hit
17	hit	17	hit	17	hit	17	hit	17	stay
18	hit	18	stay	18	stay	18	hit	18	hit
19	hit	19	stay	19	stay	19	stay	19	stay

You find these features limiting, so you want a *single* set of features that can represent *any* arbitrary policy for this game via Q -values in the form $\sum_{i=1}^n w_i f_i(s, a)$. Remember that policies are defined over all states with choices: $\{0, 2, \dots, 21\}$, not just the states listed in the previous part.


(d) (1 pt) What is the minimum number N of features needed to be able to encode all policies in this manner? Briefly justify your answer (no more than one sentence!).

7. (12 points) Mumps Outbreak There has been an outbreak of mumps at UC Berkeley. You feel fine, but you're worried that you might already be infected and therefore won't be healthy enough to take your CS 188 midterm. You decide to use Bayes nets to analyze the probability that you've contracted the mumps.

You first think about the following two factors:

- You think you have immunity from the mumps (+*i*) due to being vaccinated recently, but the vaccine is not completely effective, so you might not be immune (-*i*).
- Your roommate didn't feel well yesterday, and though you aren't sure yet, you suspect they might have the mumps (+*r*).

Denote these random variables by *I* and *R*. Let the random variable *M* take the value +*m* if you have the mumps, and -*m* if you do not. You write down the following Bayes net to describe your chances of being sick:




(a) (2 pt) Fill in the following table with the joint distribution over *I*, *M*, and *R*, $P(I, M, R)$.

<i>I</i>	<i>R</i>	<i>M</i>	$P(I, R, M)$
+ <i>i</i>	+ <i>r</i>	+ <i>m</i>	0
+ <i>i</i>	+ <i>r</i>	- <i>m</i>	
+ <i>i</i>	- <i>r</i>	+ <i>m</i>	0
+ <i>i</i>	- <i>r</i>	- <i>m</i>	
- <i>i</i>	+ <i>r</i>	+ <i>m</i>	0.056
- <i>i</i>	+ <i>r</i>	- <i>m</i>	0.024
- <i>i</i>	- <i>r</i>	+ <i>m</i>	0.024
- <i>i</i>	- <i>r</i>	- <i>m</i>	0.096

(b) (2 pt) What is the marginal probability $P(+m)$ that you have the mumps?

(c) (3 pt) Assuming you do have the mumps, you're concerned that your roommate may have the disease as well. What is the probability $P(+r \mid +m)$ that your roommate has the mumps given that you have the mumps? Note that you still don't know whether or not you have immunity.

You're still not sure if you have enough information about your chances of having the mumps, so you decide to include two new variables in the Bayes net. Your roommate went to a frat party over the weekend, and there's some chance another person at the party had the mumps (+ f). Furthermore, both you and your roommate were vaccinated at a clinic that reported a vaccine mix-up. Whether or not you got the right vaccine (+ v or - v) has ramifications for both your immunity (I) and the probability that your roommate has since contracted the disease (R). Accounting for these, you draw the modified Bayes net shown on the right.



(d) (5 pt) Circle all of the following statements which are *guaranteed* to be true for this Bayes net:

(i) $V \perp\!\!\!\perp M \mid I, R$

(ii) $V \perp\!\!\!\perp M \mid R$

(iii) $M \perp\!\!\!\perp F \mid R$

(iv) $V \perp\!\!\!\perp F$

(v) $V \perp\!\!\!\perp F \mid M$

(vi) $V \perp\!\!\!\perp F \mid I$

8. (6 points) The Sound of (Silent) Music

You are an exobiologist, studying the wide range of life in the universe. You are also an avid dancer and have an excellent model of the way species invent dancing. The key variables are:


Sound sensing (S): Whether or not a species has the ability to sense sound

Cold climate (C): Whether or not the native planet of the species has cold weather

Music (M): Whether or not the species invented music

Non-verbal communication (N): Whether or not the species has any form of non-verbal communication

You model the relationships between these variables and **dancing (D)** using the Bayes net specified to the right.



You want to know how likely it is for a dancing, sound-sensing species to invent music, according to this Bayes net. However, you're a doctor, not a dynamic programmer, and you can't bear the thought of variable elimination this late in the exam. So, you decide to do inference via sampling.

You use prior sampling to draw the samples below:

S	C	M	N	D
-s	+c	+m	-n	+d
+s	+c	-m	-n	-d
+s	-c	-m	+n	-d
+s	+c	+m	-n	+d
+s	+c	-m	+n	+d
+s	-c	-m	+n	-d
+s	-c	-m	-n	-d
+s	+c	+m	+n	+d
+s	+c	-m	+n	-d
-s	-c	-m	-n	-d

- (a) (2 pt) Based on rejection sampling using the samples above, what is the answer to your query, $P(-m \mid +d, +s)$?

While your sampling method has worked fairly well in many cases, for rare cases (like species that can't sense sound) your results are less accurate as rejection sampling rejects almost all of the data. You decide to use likelihood weighting instead. The conditional probabilities of the Bayes net are listed below.

C	M	N	D	P(D C , M , N)
+c	+m	+n	+d	0.9
+c	+m	+n	-d	0.1
+c	+m	-n	+d	0.8
+c	+m	-n	-d	0.2
+c	-m	+n	+d	0.8
+c	-m	+n	-d	0.2
+c	-m	-n	+d	0.2
+c	-m	-n	-d	0.8
-c	+m	+n	+d	0.8
-c	+m	+n	-d	0.2
-c	+m	-n	+d	0.5
-c	+m	-n	-d	0.5
-c	-m	+n	+d	0.6
-c	-m	+n	-d	0.4
-c	-m	-n	+d	0.1
-c	-m	-n	-d	0.9

S	M	P(M S)
+s	+m	0.8
+s	-m	0.2
-s	+m	0.1
-s	-m	0.9

S	P(S)
+s	0.9
-s	0.1

S	N	P(N S)
+s	+n	0.7
+s	-n	0.3
-s	+n	0.9
-s	-n	0.1

C	P(C)
+c	0.5
-c	0.5

You now wish to compute the probability that a species that has no sound-sensing ($-s$) or dancing ($-d$) nonetheless has music ($+m$), using likelihood weighting. I.e., you want $P(+m \mid -s, -d)$.

(b) (2 pt) You draw the samples below, using likelihood weighting. For each of these samples, indicate its weight.

S	C	M	N	D	weight
-s	+c	+m	+n	-d	
-s	+c	-m	-n	-d	
-s	-c	-m	+n	-d	
-s	+c	+m	+n	-d	

(c) (2 pt) Compute the answer to your query, $P(+m \mid -s, -d)$, using likelihood weighting with these samples.