
CS 188
Fall 2008

Introduction to
Artificial Intelligence

Midterm Exam

INSTRUCTIONS

- You have 80 minutes. 70 points total. Don't panic!
- The exam is closed book, closed notes except a one-page crib sheet, 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.
- **Question 0: Fill out the following grid and write your name, SID, login, and GSI at the top of each subsequent page.** (-1 points if done incorrectly!)

Last Name	
First Name	
SID	
Login	
GSI	
<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	Total
/12	/11	/15	/17	/15	/70

THIS PAGE INTENTIONALLY LEFT BLANK

1. (12 points.) Search: Mr. and Ms. Pacman

Pacman and Ms. Pacman are lost in an $N \times N$ maze and would like to meet; *they don't care where*. In each time step, *both* simultaneously move in one of the following directions: {NORTH, SOUTH, EAST, WEST, STOP}. They do *not* alternate turns. You must devise a plan which positions them together, somewhere, in as few time steps as possible. Passing each other does not count as meeting; they must occupy the same square at the same time.

(a) (4 points) Formally state this problem as a *single-agent* state-space search problem.

States:

Maximum size of state space:

Maximum branching factor:

Goal test:

(b) (3 points) Give a non-trivial admissible heuristic for this problem.

(c) (3 points) Circle all of the following graph search methods which are guaranteed to output optimal solutions to *this problem*:

- (i) DFS
- (ii) BFS
- (iii) UCS
- (iv) A* (with a consistent and admissible heuristic)
- (v) A* (with heuristic that returns zero for each state)
- (vi) Greedy search (with a consistent and admissible heuristic)

(d) (2 points) If h_1 and h_2 are admissible, which of the following are also guaranteed to be admissible? Circle all that apply:

- (i) $h_1 + h_2$
- (ii) $h_1 * h_2$
- (iii) $\max(h_1, h_2)$
- (iv) $\min(h_1, h_2)$
- (v) $(\alpha)h_1 + (1 - \alpha)h_2$, for $\alpha \in [0, 1]$

2. (11 points.) CSPs: Finicky Feast

You are designing a menu for a special event. There are several choices, each represented as a variable: (A)ppetizer, (B)everage, main (C)ourse, and (D)essert. The domains of the variables are as follows:

- A: (v)eggies, (e)scargot
 B: (w)ater, (s)oda, (m)ilk
 C: (f)ish, (b)eef, (p)asta
 D: (a)pple pie, (i)ce cream, (ch)eese

Because all of your guests get the same menu, it must obey the following dietary constraints:

- (i) Vegetarian options: The appetizer must be veggies or the main course must be pasta or fish (or both).
- (ii) Total budget: If you serve the escargot, you cannot afford any beverage other than water.
- (iii) Calcium requirement: You must serve at least one of milk, ice cream, or cheese.

(a) (3 points) Draw the constraint graph over the variables A, B, C, and D.

(b) (2 points) Imagine we first assign $A=e$. Cross out eliminated values to show the domains of the variables after forward checking.

$$\begin{array}{l} A \ [\quad e \quad] \\ B \ [\ w \ s \ m \] \\ C \ [\ f \ b \ p \] \\ D \ [\ a \ i \ ch \] \end{array}$$

(c) (3 points) Again imagine we first assign $A=e$. Cross out eliminated values to show the domains of the variables after arc consistency has been enforced.


$$\begin{array}{l} A \ [\quad e \quad] \\ B \ [\ w \ s \ m \] \\ C \ [\ f \ b \ p \] \\ D \ [\ a \ i \ ch \] \end{array}$$

(d) (1 point) Give a solution for this CSP or state that none exists.

(e) (2 points) For general CSPs, will enforcing arc consistency after an assignment *always* prune at least as many domain values as forward checking? Briefly explain why or why not.

3. (15 points.) Game Trees: The Balancer


Consider the following zero-sum game, in which the utilities $U_A(s)$ are shown for the first player (A). Assume the second player (B) is a minimizer: B holds the opposite utilities to A, $U_B(s) = -U_A(s)$. In this case, B's maximization of U_B is equivalent to minimization of U_A (i.e. the computation is standard minimax).



(a) (2 points) In each node, write $U_A(s)$, the (minimax) utility of that state for player A, assuming that B is a minimizer.

(b) (3 points) Cross off any nodes which will be skipped by alpha-beta pruning, assuming left-to-right ordering.

Assume now that B is not a minimizer, but a *balancer*. A balancer does not try to minimize A's score, but rather wishes the outcome of the game to be as balanced as possible. Formally, assume B's utility for a state s is defined as $U_B(s) = -|U_A(s)|$. The game tree is shown here, with hexagons indicating player B's control.



(c) (3 points) In each node, write $U_A(s)$, the utility of that state for player A, assuming that B is a balancer.

(d) (3 points) Write pseudocode for the functions which compute the $U_A(s)$ values of game states in the general case of multi-turn games where B is a balancer. Assume you have access to the following functions: `successors(s)` gives the possible next states, `isTerminal(s)` checks whether a state is a terminal state, and `terminalValue(s)` returns A's utility for a terminal state. *Careful:* As in minimax, be sure that both functions compute and return player A's utilities for states – B's utility can always be computed from A's utility.

```
def maxValue(s): // compute  $U_A(s)$  assuming that A is next to move.
```

```
def balanceValue(s): // compute  $U_A(s)$  assuming that B is next to move.
```

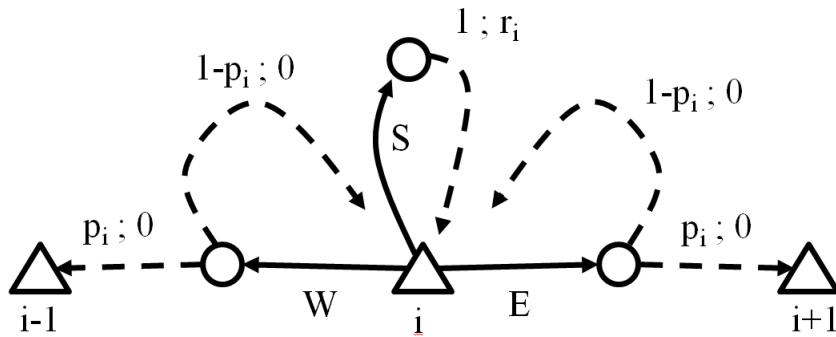
(h) (2 points) Consider pruning children of a B node in this scenario. On the tree on the bottom of the previous page, cross off any nodes which can be pruned, again assuming left-to-right ordering.

(i) (2 points) Again consider pruning children of a B node s . Let α be the best option for an A node higher in the tree, just as in alpha-beta pruning, and let v be the U_A value of the best action B has found so far from s . Give a general condition under which `balanceValue(s)` can return without examining any more of its children.

4. (17 points.) MDPs and RL: Wandering Merchant

There are N cities along a major highway numbered 1 through N . You are a merchant from city 1 (that's where you start). Each day, you can either travel to a neighboring city (actions *East* or *West*) or stay and do business in the current city (action *Stay*). If you choose to travel from city i , you successfully reach the next city with probability p_i , but there is probability $1 - p_i$ that you hit a storm, in which case you waste the day and do not go anywhere. If you stay to do business in city i , you get $r_i > 0$ in reward; a travel day has reward 0 regardless of whether or not you succeed in changing cities.

The diagram below shows the actions and transitions from city i . Solid arrows are actions; dashed arrows are resulting transitions labeled with their probability and reward, in that order.



(a) (2 points) If for all i , $r_i = 1$, $p_i = 1$, and there is a discount $\gamma = 0.5$, what is the value $V^{stay}(1)$ of being in city 1 under the policy that always chooses *stay*? Your answer should be a real number.

(b) (2 points) If for all i , $r_i = 1$, $p_i = 1$, and there is a discount $\gamma = 0.5$, what is the optimal value $V^*(1)$ of being in city 1?

(c) (2 points) If the r_i 's and p_i 's are known positive numbers and there is almost no discount, i.e. $\gamma \approx 1$, describe the optimal policy. You may define it formally or in words, e.g. "always go east," but your answer should precisely define how an agent should act in any given state. *Hint:* You should not need to do any computation to answer this question.

Suppose we run value iteration. Recall that $V_k(s)$ is the value of state s after k rounds of value iteration and all the values are initialized to zero.

(d) (2 points) If the optimal value of being in city 1 is positive, i.e. $V^*(1) > 0$, what is the largest k for which $V_k(1)$ could still be zero? Be careful of off-by-one errors.

(e) (2 points) If all of the r_i and p_i are positive, what is the largest k for which $V_k(s)$ could still be zero for some state s ? Be careful of off-by-one errors.

Suppose we don't know the r_i 's or the p_i 's, so we decide to do Q-learning.

(f) (3 points) Suppose we experience the following sequence of states, actions, and rewards: $(s=1, a=stay, r=4)$, $(s=1, a=east, r=0)$, $(s=2, a=stay, r=6)$, $(s=2, a=west, r=0)$, $(s=1, a=stay, r=4, s=1)$. What are the resulting $Q(s, a)$ values if the learning rate is 0.5, the discount is 1, and we start with all $Q(s, a) = 0$? Fill in the table below; each row should hold the q-values after the transition specified in its first column. You may leave unchanged values blank.

(s,a,r,s')	Q(1,S)	Q(1,E)	Q(2,W)	Q(2,S)
initial	0	0	0	0
(1,S,4,1)				
(1,E,0,2)				
(2,S,6,2)				
(2,W,0,1)				
(1,S,4,1)				

Circle *true* or *false*; skipping here is worth 1 point per question.

(g) (2 points) (*True/False*) Q-learning will only learn the optimal q-values if actions are eventually selected according to the optimal policy.

(h) (2 points) (*True/False*) In a deterministic MDP (i.e. one in which each state / action leads to a single deterministic next state), the Q-learning update with a learning rate of $\alpha = 1$ will correctly learn the optimal q-values.

5. (15 points.) Probability and Utilities: Wheel of Fortune

You are playing a simplified game of Wheel of Fortune. The objective is to correctly guess a three letter word. Let X , Y , and Z represent the first, second, and third letters of the word, respectively. There are only 8 possible words: X can take on the values 'c' or 'l', Y can be 'a' or 'o', and Z can be 'b' or 't'.

Before you guess the word, two of the three letters will be revealed to you. In the first round of the game, you choose one of X , Y or Z to be revealed. In the second round, you choose one of the remaining two letters to be revealed. In the third round, you guess the word. If you guess correctly, you win. The utility of winning is 1, while the utility of losing is 0.

You watch the game a lot and determine that the eight possible words occur with the probabilities shown on the right. Your goal is to act in such a way as to maximize your chances of winning (and thereby your expected utility).

$P(X, Y, Z)$

X	Y	Z	P
c	a	b	0.10
c	a	t	0.10
c	o	b	0.20
c	o	t	0.20
l	a	b	0.05
l	a	t	0.15
l	o	b	0.05
l	o	t	0.15

(a) (3 points) What is the distribution $P(Y, Z)$? Your answer should be in the form of a table.

(b) (2 points) Are the second and third letters (Y and Z) independent? Show a specific computation that supports your claim.

(c) (2 points) Are the second and third letters (Y and Z) independent if you know the value of the first letter (X)? Show a specific computation that supports your claim.

Suppose that in the first round, you ask about X and are told that $X = c$. It is the second round and you can now either ask the host to reveal Y or to reveal Z .

(d) (2 points) If you ask the host to reveal Y , what is the probability that you will win in the third round?

(e) (1 point) What letter *should* you ask the host about in the second round to maximize your chance of winning, Y or Z ?

(f) (3 points) What is your expected utility if you act optimally from the state where $X=c$?

(g) (2 points) Suppose that the host is allowed to pick any distribution over the three variables but has to tell you what the distribution is before the game starts. What distribution should the host pick to minimize your chances of winning? Justify your answer briefly.