15-381 Spring 2007
Final Exam SOLUTIONS

Spring 2007
May 11

Name:____________________________________ Andrew ID:____________________

- This is an open-book, open-notes examination. You have 180 minutes to complete this examination.
- Unless explicitly requested, we do not need to see the details of how you arrive at the answer. Just write the answer and do not waste time writing the details of the derivation.
- Write your answers legibly in the space provided on the examination sheet. If you use the back of a sheet, indicate clearly that you have done so on the front.
- Write your name and Andrew ID on this page and your andrew id on the top of each successive page in the space provided.
- Calculators are allowed but laptops and PDAs are not allowed.
- Good luck!

Question	Points
1	/10
2	/20
3	/20
4	/10
5	/20
6	/18
7	/12
8	/10
9	/10
10	/20
TOTAL	/150

Question 1 T/F Questions (10 Points) (1 point each)

1. Breadth-first search and iterative-deepening search always find the same solution.
   False

2. Breadth-first search is a special case of uniform cost search.
   True

3. Best-first search can be thought of as a special case of A*.
   False

4. A heuristic that always evaluates to $h(s) = 1$ for non-goal search nodes s is always admissible.
   False

5. A simultaneously-played zero-sum game with hidden information never has a pure strategy Nash Equilibrium.

   False. It may not but some do.

6. MDP instances with small discount factors tend to emphasize near-term rewards.

   True.

7. For sufficiently high dimensions, the joint distribution is approximately the product of the marginal distributions.
   False

8. Bayes networks can be completely specified by the conditional probabilities of the linked nodes.
   False

9. The entropy of a binary random variable decreases as data become less ordered.

   `False`

10. The k-Nearest Neighbors approach scales well to high dimensional spaces.

    `False`

Question 2 Search and CSP (20 Points)

1. Consider the search graph below, where S is the start node and G1, G2, and G3 are goal states. Arcs are labeled with the cost of traversing them and the heuristic cost to a goal is shown inside the nodes. For each of the three search strategies below, indicate which of the goal states is reached:

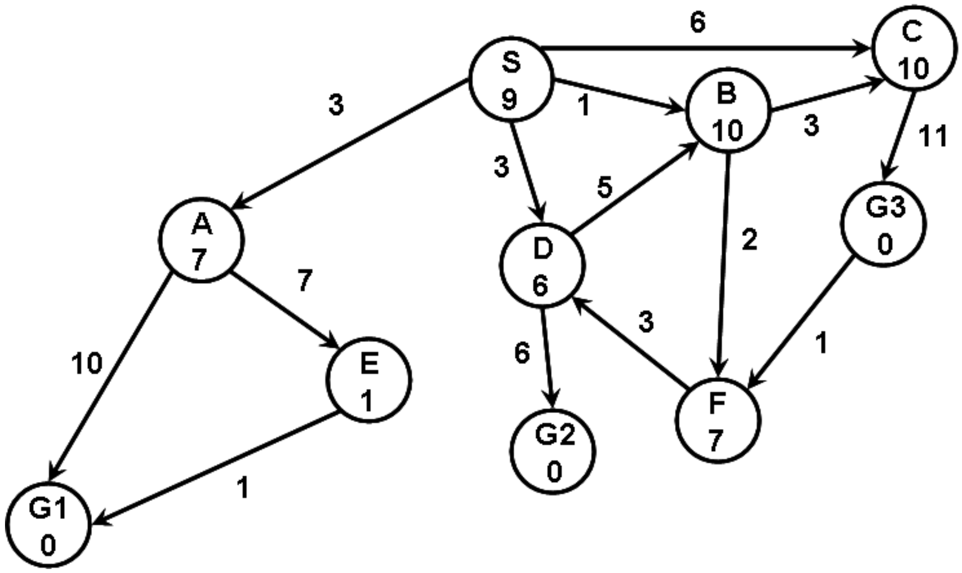

   (a) (2 Points) Breadth-first search. Goal reached: G1

   (b) (2 Points) Uniform cost search. Goal reached: G2

   (c) (2 Points) A* search. Goal reached: G2

2. For a general search problem, state which of breadth-first search (BFS) or depth-first search (DFS) is preferred under which of the following conditions:

   (a) (2 Points) A shallow solution (path from initial state to goal state) is preferred.
   BFS

   (b) (2 Points) The search tree may contain large or possibly infinite branches.
   BFS

(c) (2 Points) Very large memory space to store the search tree (or the queue) is available.
BFS

3. For a general search problem, state which of iterative deepening (ID) or depth-first search (DFS) is preferred under which of the following conditions:

   (a) (2 Points) A shallow solution is preferred.
   ID

   (b) (2 Points) The search tree may contain large or infinite branches.
   ID

4. (4 Points) Consider the Minesweeper game. Each square contains either a zero (touching no bombs in its eight neighboring squares); a number n (touching exactly n bombs); or nothing (unknown). As discussed in class, the problem of determining which squares contain bombs can be formulated as a constraints satisfaction problem, with a variable for each "unknown" square touching a numbered square and a constraint for each numbered square. The possible values for the unknown squares are B (bomb) and E (empty).

   For the following configuration, state whether performing constraint propagation will uniquely discover the values of the squares without requiring depth first search.

0	0	1	
0	0	2	
1	2	3	
B			

   No

Question 3 MDPs (20 Points)

1. A boy is being chased around the school yard by bullies and must choose whether to Fight or Run.
   - There are three states:
     - Ok (O), where he is fine for the moment.
     - Danger (D), where the bullies are right on his heels.
     - Caught (C), where the bullies catch up with him and administer noogies.
   - He begins in state O 75% of the time.
   - He begins in state D 25% of the time.

   The graph of the MDP is given here:

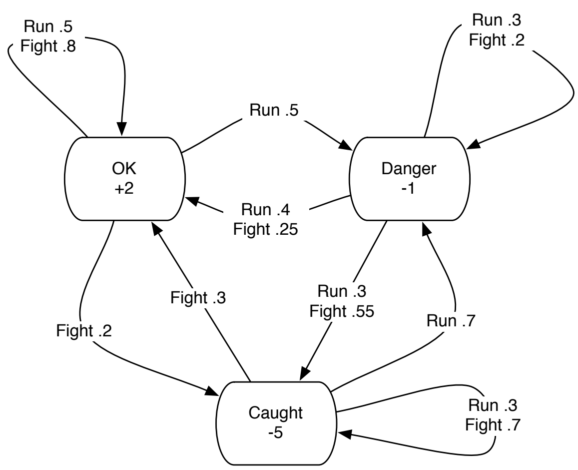

(a) (9 points) Fill out the table with the results of value iteration with a discount factor $\gamma = .9$:

k	$J^k(O)$	$J^k(D)$	$J^k(C)$
1	2	-1	-5
2	2.54	-1.9	-6.98

$$J^2(O) = max(2 + .9((.5 * 2) + (.5 * -1)), 2 + .9((.8 * 2) + (.2 * -5))$$
$$= max(2.45, 2.54)$$
$$= 2.54$$
$$J^2(D) = max(-1 + .9((.4 * 2) + (.3 * -1) + (.3 * -5)), -1 + .9((.25 * 2) + (.2 * -1) + (.55 * -1)))$$
$$= max(-1.9, -3.2050)$$
$$= -1.9$$
$$J^2(C) = max(-5 + .9((.7 * -1) + (.3 * -5)), -5 + .9((.3 * 2) + (.7 * -5))$$
$$= max(-6.98, -7.61)$$
$$= -6.98$$

(b) (5 points) At $k = 2$ with $\gamma = .9$ what policy would you select? Is it necessarily true that this is the optimal policy? At $k = 3$ what policy would you select? Is it necessarily true that this is the optimal policy?

- From O choose Fight.
- From D choose Run.
- From C choose Run.

Value iteration that has not converged is not guaranteed to find the optimal policy, so this policy is not necessarily optimal.

2. (a) (3 points) Suppose you have a robot trying to reach a goal and avoid cliffs in a small grid world. It can only move North, South, East, or West, but occasionally fails to move in the intended direction. If you were to model this using an MDP and were trying to solve it optimally, should you use value iteration or policy iteration? Justify your answer in one sentence.

   Generally, use value iteration. We have many states and a few actions, and value iteration is generally cheaper than policy iteration.

   (b) (3 points) Now suppose that the robot can teleport to any grid cell but the teleportation causes it to land in neighboring grid cells near the target with some probability. Of you were to model this using an MDP and were trying to solve it optimally should you use value iteration or policy iteration? Justify your answer in one sentence.

   Now we should use policy iteration. Instead of 4 actions we now have an action for teleporting to each different square of the grid, which is now a lot of actions; policy iteration is better when we have many many actions.

Question 4 Game Theory (10 Points)

1. Consider the following non-zero sum game in matrix-normal form (with Player A's reward first)

Player A \ Player B	X	Y	Z
X	4, 6	3, 1	1, 5
Y	2, 2	1, 4	1, 4
Z	3, 3	0, 0	2, 6

   (a) (1 point) Are there dominant strategies for either player?

   No

   (b) (2 points) What is the pure Nash Equilibrium for this game?

   (X,X) and/or (Z,Z)

   (c) (1 point) Is the NE for this game significant in any other way?

   (X,X) is also the (strongly) pareto optimal outcome for the game. (All received full credit for this question).

   (d) (2 points) What is the 2x2 matrix that results from the using iterative elimination of dominated strategies?

   Row Y is dominated by Row X, and Column Y is dominated by Column Z. This leaves

Player A \ Player B	X	Z
X	4, 6	1, 5
Z	3, 3	2, 6

(e) (4 points) What is a mixed strategy NE for the 2x2 reduced version of this game?

We find mixed strategies as described in the notes - in the following p is the percentage X is taken by player A and q is the percentage X is taken by player B:

$$\begin{aligned}
u_A &= 4pq + p(1-q) + 3(1-p)q + 2(1-p)(1-q) \\
&= 4pq + p - pq + 3q - 3pq + 2 - 2p + 2pq + 2q \\
&= 2pq - p + 5q + 2 \\
u'_A(\delta p) &= 2q - 1 \\
2q &= 1 \\
q^* &= \frac{1}{2} \\
u_B &= 6pq + 5p(1-q) + 3(1-p)q + 6(1-p)(1-q) \\
&= 6pq + 5p - 5pq + 3q - 3pq + 6 - 6p - 6q + 6pq \\
&= 4pq - 1p - 3q + 6 \\
u'_B(\delta q) &= 4p - 3 \\
4p &= 3 \\
p^* &= \frac{3}{4}
\end{aligned}$$

So $(p^*, q^*) = (\frac{3}{4}, \frac{1}{2})$.

Question 5 Bayes Nets and Naive Bayes (20 points)

1. (3 points) Write down the joint distribution as it factorizes according to the graph below.

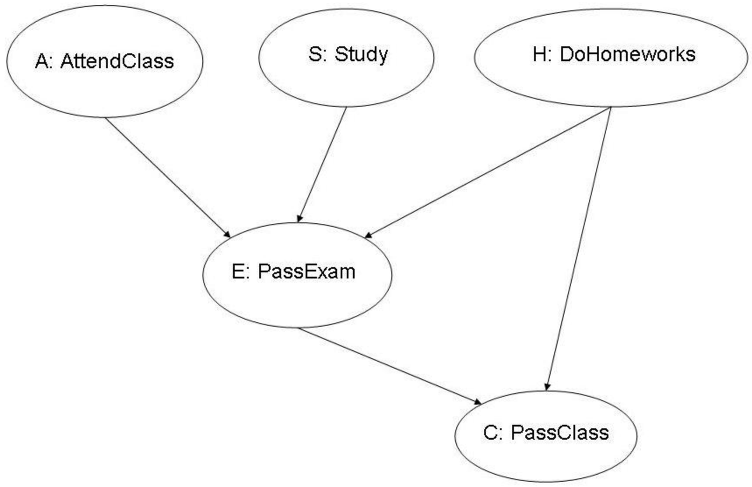

$P(A, S, H, E, C) = P(A) * P(S) * P(H) * P(E|A, S, H) * P(C|E, H)$

2. (5 points) Use variable elimination and your result from the previous question to write down the expression for the probability of passing the class, given that you attend class and study, but don't do the homeworks.

$$\begin{aligned}
P(C|A, S, \neg H) &= \frac{P(C, A, S, \neg H)}{P(A, S, \neg H)} && (1)\\
&= \frac{\Sigma_e P(A, S, \neg H, E = e, C))}{\Sigma_{e,c} P(A, S, \neg H, E = e, C = c)} && (2)\\
&= \frac{\Sigma_e P(A) * P(S) * P(\neg H) * P(E = e|A, S, \neg H) * P(C|E = e, \neg H)}{\Sigma_{e,c} P(A) * P(S) * P(\neg H) * P(E = e|A, S, \neg H) * P(C = c|E = e, \neg H)} && (3)\\
&= \frac{P(A) * P(S) * P(\neg H) * \Sigma_e P(E = e|A, S, \neg H) * P(C|E = e, \neg H)}{P(A) * P(S) * P(\neg H) * \Sigma_e P(E = e|A, S, \neg H) * \Sigma_c P(C = c|E = e, \neg H)} && (4)\\
&= \frac{P(A) * P(S) * P(\neg H) * \Sigma_e P(E = e|A, S, \neg H) * P(C|E = e, \neg H)}{P(A) * P(S) * P(\neg H)} && (5)\\
&= \Sigma_e P(E = e|A, S, \neg H) * P(C|E = e, \neg H) && (6)\\
& && (7)
\end{aligned}$$

3. (5 points) Use the following CPTs for the graph of question 1 to compute $P(A|C, H)$.

$P(A) = 0.5, P(S) = 0.7, P(H) = 0.9$

A	S	H	$P(E\|A, S, H)$
0	0	0	0.2
0	0	1	0.5
0	1	0	0.4
0	1	1	0.8
1	0	0	0.3
1	0	1	0.7
1	1	0	0.6
1	1	1	0.9

E	H	$P(C\|E, H)$
0	0	0.1
0	1	0.4
1	0	0.3
1	1	0.9

$$
\begin{aligned}
P(A|C, H) &= \frac{P(A, C, H)}{P(C, H)} \\
&= \frac{\Sigma_{e,s} P(A, S = s, H, E = e, C)}{\Sigma_{a,e,s} P(A = a, S = s, H, E = e, C)} \\
&= \frac{\Sigma_{e,s} P(A) * P(S = s) * P(H) * P(E = e|A, S = s, H) * P(C|E = e, H)}{\Sigma_{a,e,s} P(A = a) * P(S = s) * P(H) * P(E = e|A = a, S = s, H) * P(C|E = e, H)} \\
&= \frac{P(A) * P(H) * \Sigma_s P(S = s) * \Sigma_e P(E = e|A, S = s, H) * P(C|E = e, H)}{P(H) * \Sigma_a P(A = a) * \Sigma_s P(S = s) * \Sigma_e P(E = e|A = a, S = s, H) * P(C|E = e, H)} \\
&= \frac{P(A) * \Sigma_s P(S = s) * \Sigma_e P(E = e|A, S = s, H) * P(C|E = e, H)}{\Sigma_a P(A = a) * \Sigma_s P(S = s) * \Sigma_e P(E = e|A = a, S = s, H) * P(C|E = e, H)} \\
&= \frac{.5(.7(.9 * .9 + .1 * .4) + .3(.7 * .9 + .3 * .4))}{(.5(.7(.9 * .9 + .1 * .4) + .3(.7 * .9 + .3 * .4)) + .5(.7(.8 * .9 + .2 * .4) + .3(.5 * .9 + .5 * .4)))} \\
&= \frac{0.41}{0.7875} \\
&= 0.5206
\end{aligned}
$$

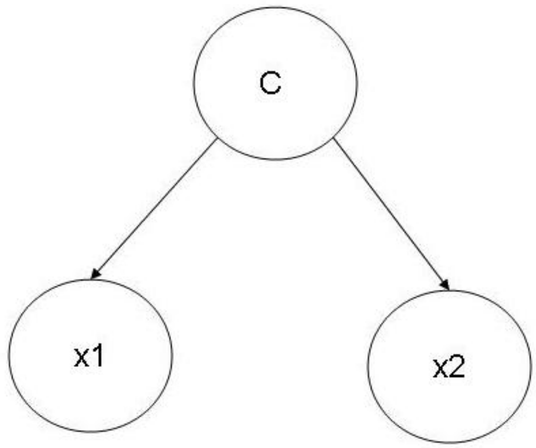

4. (2 points) Above is the Bayes net structure for a Naive Bayes classifier, where C is the class (1 or 2) and x1 and x2 are binary features. Why isn't there an arrow between x1 and x2 in the graph?

   The 'Naive' assumption states that the features are conditionally independent given the class variable. A link between these nodes would signify dependence.

5. (5 points) Use the following CPTs to classify the following observation according to the naive Bayes classifier in question 4: x1=1 and x2=1. You must show your work to get credit.

   $P(C = 1) = 0.6$

C	$P(x1\|C)$
1	0.7
2	0.6

C	$P(x2\|C)$
1	0.5
2	0.8

$$\begin{aligned} P(C=1|x1=1,x2=1) &\ \alpha\ P(x1=1,x2=1|C=1)P(C=1) && (8)\\ &\ \alpha\ P(x1=1|C=1)P(x2=1|C=1)P(C=1) && (9)\\ &\ \alpha\ 0.7*0.5*0.6 && (10)\\ &\ \alpha\ 0.21 && (11)\\ & && (12)\end{aligned}$$

$$\begin{aligned} P(C=2|x1=1,x2=1) &\ \alpha\ P(x1=1,x2=1|C=2)P(C=2) && (13)\\ &\ \alpha\ P(x1=1|C=2)P(x2=1|C=2)P(C=2) && (14)\\ &\ \alpha\ 0.6*0.8*0.4 && (15)\\ &\ \alpha\ 0.192 && (16)\\ & && (17)\end{aligned}$$

$P(C = 1|x1 = 1, x2 = 1) > P(C = 2|x1 = 1, x2 = 1)$ so classify this example as Class 1.

Question 6 Neural Networks (18 Points)

Suppose you are given predictions of n different experts (or, automated learners), whether a given email message is SPAM (1), or EMAIL (0). Your goal is to output a single prediction per message, that would be as accurate as possible. For this purpose, you'd like to implement a *majority voting* mechanism. That is, if more than half of the experts predict SPAM, than your final prediction should be SPAM for that instance. Otherwise, the final prediction should be EMAIL.

1. (5 points) Suggest a neural network, that implements majority voting. For simplicity, we will assume there are 4 experts overall (named A,B,C,D). Specify the network structure and weights.

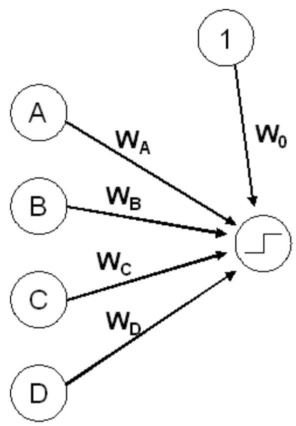

   Consider the weights $W_0 = 0.5$, and $W_A = W_B = W_C = W_D = \frac{1}{4}$.

   Then, whenever $\sum_i W_i - 0.5 > 0$ the output would give 1. This is the case if 3 or more experts give a positive prediction. Otherwise, the output would be 0 (that is, the default here is 0, in case of a tie).

2. (3 points) Explain shortly (1-2 lines) how to adapt the network structure and weights to the general case of n experts.

   The network should have n incoming edge, in addition to a threshold edge. The weights would be $W_0 = 0.5$ and else, $W_i = \frac{1}{n}$.

3. (5 points) Which model – a neural network, or a decision tree - is more suitable for this problem (for the general case of n experts, where n may be large)? Explain, in terms of representation space.

   Majority voting over multiple inputs cannot be described in a logical form. Therefore, a decision tree, while can be tailored to represent the relevant combinations that represent a majority, is less suitable for this problem in terms of representation space.

4. (5 points) Suppose that every individual expert uses completely different information and they all have the same accuracy level (say, 85%) on the given set of examples. Can majority voting improve accuracy in this case on the same set of examples? Explain your answer in one sentence.

   Yes, a multiple classifier system can only improve the performance when the members in the system are diverse from each other. If the experts are not completeley dependent, and don't do exactly the same mistakes, then majority voting would improve performance.

Question 7 Probability (12 Points)

Let A and B be two binary random variables independent events with probabilities $P(A = 1) = 0.1$ and $P(B = 1) = 0.4$. Let C denote the event that at least one of the events A and B is on, *i.e.*, C=A OR B, and let D be the event that exactly one of the events A and B occurs, *i.e.*, D = A XOR B.

1. (3 points) Compute $P(C = 1)$.

$$P(C = 1) = P(A{\cap}B) = P(A) + P(B) - P(A{\cap}B) = 0.1 + 0.4 - 0.04 = 0.46$$

2. (3 points) Compute $P(D = 1)$.

$$P(D) = P(A{\cup}B) - P(A{\cap}B) = 0.46 - 0.04 = 0.42$$

3. (3 points) Compute $P(D|A)$.

$$P(D|A) = \frac{P(A{\cap}D)}{P(A)}$$

$$P(A{\cap}D) = P(A) - P(A{\cap}B) = 0.1 - 0.04 = 0.06$$

$$P(D|A) = \frac{0.06}{0.42} = 0.143$$

4. (3 points) Prove that A and D are not independent, using the results from the previous questions.

   Simply note that P(D—A) is not equal to P(D). Therefore D and A are not independent.

Question 8 Decision Trees (10 Points)

A candy manufacturer interviews a customer on his willingness to eat a candy of a particular color or flavor. The following table shows the collected responses:

Color	Flavor	Edibility
Red	Grape	Yes
Red	Cherry	Yes
Green	Grape	Yes
Green	Cherry	No
Blue	Grape	No
Blue	Cherry	No

1. (2 points) What is H(edibility)? It is not necessary to show your work.

   $-\left(\frac{1}{2}log_2\frac{1}{2} + \frac{1}{2}log_2\frac{1}{2}\right) = 1$

2. (2 points) What is H(edibility | color)? It is not necessary to show your work.

   $-\left(\frac{1}{3}(1log_2 1 + 0log_2 0) + \frac{1}{3}(\frac{1}{2}log_2\frac{1}{2} + \frac{1}{2}log_2\frac{1}{2}) + \frac{1}{3}(1log_2 1 + 0log_2 0)\right) = \frac{1}{3}$

3. (2 points) Which feature (color or flavor) has the larger mutual information with edibility?

   `>From inspection, H(edibility | color) < H(edibility | flavor), so color has the larger mutual information with edibility.`

4. (2 points) Draw the decision tree for predicting edibility that maximizes the information gain.

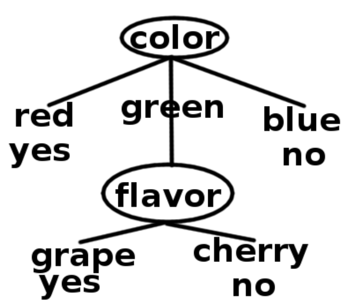

5. (2 points) Using your decision tree, what would you predict for the edibility of a blue, blueberry-flavored candy?

   `Edibility: No`

Question 9 k-Nearest Neighbors (10 Points)

In the following diagram, A through Y are all data points. Each data point has features (x,y) corresponding to its coordinate in the grid.

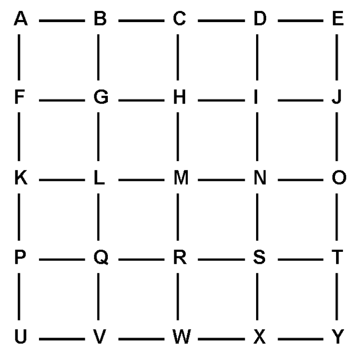

1. (3 points) What are the $k = 5$ nearest neighbors of data point M using Euclidean distance? Break any ties with alphabetical ordering.

   `G,H,L,N,R`

2. (3 points) If data points A through L belong to class 1 and data points N through Y belong to class 2, what is the classification of M when using the $k = 5$ nearest neighbors?

   `Class 1`

3. (4 points) Suppose if in addition to the (x, y) position features, we also had an additional "weight" w attached to each data point. Assume that the standard deviation of x, y, and w are as follows: $\sigma_x = 2$, $\sigma_y = 2$, $\sigma_w = 200$. Would Euclidean distance still be an appropriate distance metric? Briefly describe why or why not.

   No. In order to use Euclidean distance, we would want to *standardize* each feature by its standard deviation before combining them. Otherwise, the ''weight'' feature would overwhelm the ''x-position'' and ''y-position'' features in the distance calculations.

Question 10 Reinforcement Learning (20 Points)

1. (16 points) Consider a system with two states and two actions. You perform actions and observe the rewards and transitions listed below. Each step lists the current state, reward, action and resulting transition as $S_i; R = r; a_k : S_i \rightarrow S_j$. Perform Q-learning using a learning rate of $\alpha = 0.5$ and a discount factor of $\gamma = 0.5$ for each step. The Q-table entries are initialized to zero.

$S_1 \quad R = -10 \quad a_1 : S_1 \rightarrow S_1$

Q	S_1	S_2
a_1	-5	0
a_2	0	0

$$Q(a,s) \leftarrow Q(a,s) \quad +\alpha\big(R(s) \quad +\gamma \max_{a'}[Q(a',s')] \quad -Q(a,s)\big)$$

$$Q(a_1,S_1) \leftarrow 0 \quad +0.5\big(-10 \quad +0.5 \max_{a'(s'=S_1)}[0,0] \quad -0\big)$$

$$= -5$$

$S_1 \quad R = -10 \quad a_2 : S_1 \rightarrow S_2$

Q	S_1	S_2
a_1	-5	0
a_2	-5	0

$$Q(a_2,S_1) \leftarrow 0 \quad +0.5\big(-10 \quad +0.5 \max_{a'(s'=S_2)}[0,0] \quad -0\big)$$

$$= -5$$

$S_2 \quad R = +20 \quad a_1 : S_2 \rightarrow S_1$

Q	S_1	S_2
a_1	-5	8.75
a_2	-5	0

$$Q(a_1,S_2) \leftarrow 0 \quad +0.5\big(+20 \quad +0.5 \max_{a'(s'=S_1)}[-5,-5] \quad -0\big)$$

$$= 8.75$$

$S_1 \quad R = -10 \quad a_2 : S_1 \rightarrow S_2$

Q	S_1	S_2
a_1	-5	8.75
a_2	-5.3124	0

$$Q(a_2,S_1) \leftarrow -5 \quad +0.5\big(-10 \quad +0.5 \max_{a'(s'=S_2)}[8.75,0] \quad -(-5)\big)$$

$$= -5.3125$$

2. (4 points) What is the optimal policy at this point?

$$\pi(S_1) = a_1$$
$$\pi(S_2) = a_1$$