Name: ______________________________ Andrew Id (Block Capitals): ____________________

ANSWERS TO 15-381 Final, Spring 2004

Friday May 7, 2004

1. Place your name and your andrew email address on the front page.
2. You may use any and all notes, as well as the class textbook. Keep in mind, however, that this final was designed in full awareness of such. You may NOT use the Internet, but you can use a calculator.
3. We only require that you provide the answers. We don't need to see your work.
4. The maximum possible score on this exam is 100. You have 180 minutes.
5. Good luck!

Question	Score
1	20
2	6
3	6
4	9
5	14
6	9
7	9
8	9
9	9
10	9
Total	100

1 Short Questions

(a) When you run Waltz algorithm on the following drawing, which of the following statements is true? Circle the correct answer.

(i) The algorithm will label all edges uniquely.

(ii) The algorithm will report that some edges are ambiguous.

(iii) The algorithm will report that the image cannot be labeled consistently.

ANSWER: (ii)

(b) How many degrees of freedom does a rigid 3-d object have if it moves in a 3-d space?

ANSWER: 6

(c) How does randomized hill-climbing choose the next move each time? Circle the correct answer.

(i) It generates a random move from the moveset, and accepts this move.

(ii) It generates a random move from the whole state space, and accepts this move.

(iii) It generates a random move from the moveset, and accepts this move only if this move improves the evaluation function.

(iv) It generates a random move from the whole state space, and accepts this move only if this move improves the evaluation function.

ANSWER: (iii)

(d) Suppose you are using a genetic algorithm. Show the children of the following two strings if single point crossover is performed with a cross-point between the 4th and the 5th digits:
1 4 6 2 5 7 2 3 and 8 5 3 4 6 7 6 1

ANSWER: 1 4 6 2 6 7 6 1 and 8 5 3 4 5 7 2 3

(e) What is the entropy of these examples: 1 3 2 3 1 3 3 2

ANSWER: 1.5

(f) Which of the following is the main reason of pruning a decision tree? Circle the correct answer.

(i) to save computational cost

(ii) to avoid over-fitting

(iii) to make the training error smaller

ANSWER: (ii)

(g) Which of the following does the Naive Bayes classifier assume? Circle the correct answer.

(i) All the attributes are independent.

(ii) All the attributes are conditionally independent given the output label.

(iii) All the attributes are jointly dependent to each other.

ANSWER: (ii)

(h) By which of the following networks can XOR function be learned? Circle the correct answer.

(i) linear perceptron
(ii) single layer Neural Network
(iii) 1-hidden layer Neural Network
(iv) none of the above

ANSWER: (iii)

(i) If we use K-means on a finite set of samples, which of the following statement is true? Circle the correct answer.

(i) K-means is not guaranteed to terminate.
(ii) K-means is guaranteed to terminate, but is not guaranteed to find the optimal clustering.
(iii) K-means is guaranteed to terminate and find the optimal clustering.

ANSWER: (ii)

(j) In the worst case, what is the number of nodes that will be visited by Breadth-First Search in a (non-looping) tree with depth d and branching factor b?

ANSWER: $O(b^d)$

(k) **True** or **False** : If a search tree has cycles, A* Search with an **in**admissible heuristic might never converge when run on that tree.

ANSWER: False

(l) Circle the Nash Equilibria in the following matrix-form game:

ANSWER:

Player 2

Player 1		D	E	F
	A	0, 1	3, 5	2, 1
	B	6, 3	1, 3	5, 2
	C	4, 2	3, 4	7, 7

(m) Assume the following zero-sum game, where player 1 is the maximizer:

ANSWER:

Player 2

Player 1	C	D
A	2	0
B	0	1

If Player 1 chooses strategy A with probability p, and if Player 2 always plays strategy C, what is the expected value of the game?

ANSWER: $2p + 0(1 - p) = 2p$

(n) In the mixed strategy Nash equilibrium for the above game, with what probability does Player 1 use strategy A?

$2p = 1(1 - p)$
$3p = 1$
ANSWER: $p = 1/3$

(o) **True** or **False** : In a second-price, sealed bid auction, it is optimal to bid your true value. There is no advantage to bluffing.

ANSWER: True

(p) How many values does it take to represent the joint distribution of 4 boolean variables?

ANSWER: 16

(q) If P(A) = 0.3, P(B) = 0.4, and P(A|B) = 0.6

(a) What is P(A∧B)?
ANSWER: $P(A \wedge B) = P(A|B) * P(B) = 0.24$

(b) What is P(B|A)?
ANSWER: $P(B|A) = \frac{P(A|B)P(B)}{P(A)} = 0.8$

(c) Are A and B independent?
ANSWER: No

(r) For the following questions, use the diagram below. If you do not have enough information to answer a question, answer False.

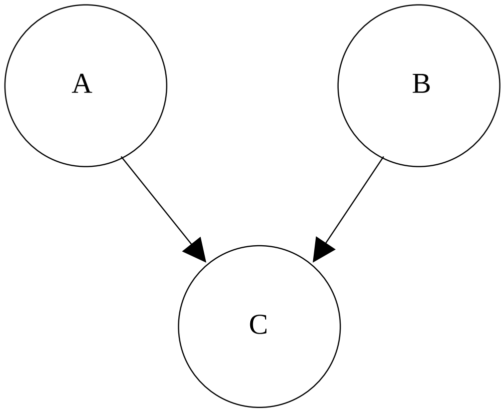

(a) **True** or **False** : $A \perp B$
ANSWER: True

(b) **True** or **False** : $A \perp C$
ANSWER: False

(c) **True** or **False** : I$< A, \{C\}, B >$
ANSWER: False

(d) **True** or **False** : I$< C, \{A\}, B >$
ANSWER: False

(s) **True** or **False** : Policy iteration will usually converge to a better policy than value iteration.

ANSWER: False

(t) **True** or **False** : For a densely connected MDP with many actions, policy iteration will generally converge faster than value iteration.

ANSWER: True

2 Hill Climbing, Simulated Annealing and Genetic Algorithm

The N-queens problem requires you to place N queens on an N-by-N chessboard such that no queen attacks another queen. (A queen attacks any piece in the same row, column or diagonal.) Here are some important facts:

- We define the states to be any configuration where the N queens are on the board, one per column.
- The moveset includes all possible states generated by moving a single queen to another square in the same column. The function to obtain these states is called the successor function.
- The evaluation function $Eval(state)$ is the number of non-attacking pairs of queens in this $state$. (Please note it is the number of NON-attacking pairs.)

In the following questions, we deal with the 6-queens problem (N=6).

1. How many possible states are there totally?

   ANSWER: 6^6

2. For each state, how many successor states are there in the moveset?

   ANSWER: 30

3. What value will the evaluation function $Eval()$ return for the current state shown below?

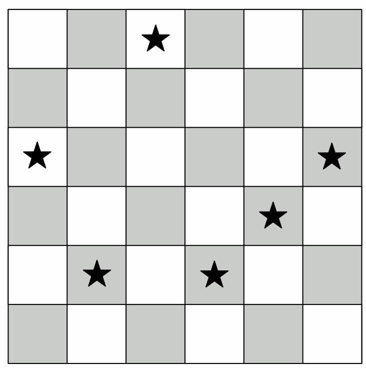

   ANSWER: 9

4. If you use Simulated Annealing (currently T=3), and the current state and the random next state are shown below, will you accept this random next state immediately? or accept it with some probability? If it is the latter case, what is the probablity?

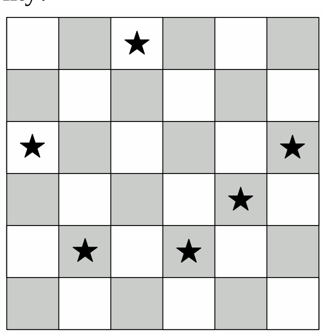

current state

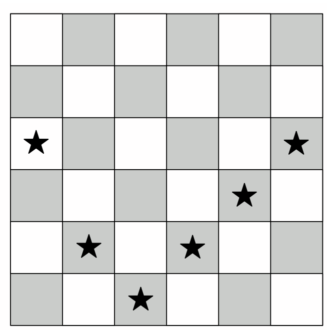

random next state

   ANSWER: For the current state, $E1 = 9$. For the next state, $E2 = 6$. So $E2 < E1$.
   $P = exp(-(E1 - E2)/T) = exp(-(9 - 6)/3)) = 1/e$
   We will accept the next state with probability $1/e$.

5. Suppose you use a Genetic Algorithm. The current generation includes four states, $S1$ through $S4$. The evaluation values for each of the four states are: $Eval(S1) = 9$, $Eval(S2) = 12$, $Eval(S3) = 11$, $Eval(S4) = 8$. Calculate the probability that each of them would be chosen in the ”selection” step (also called ”reproduction” step).

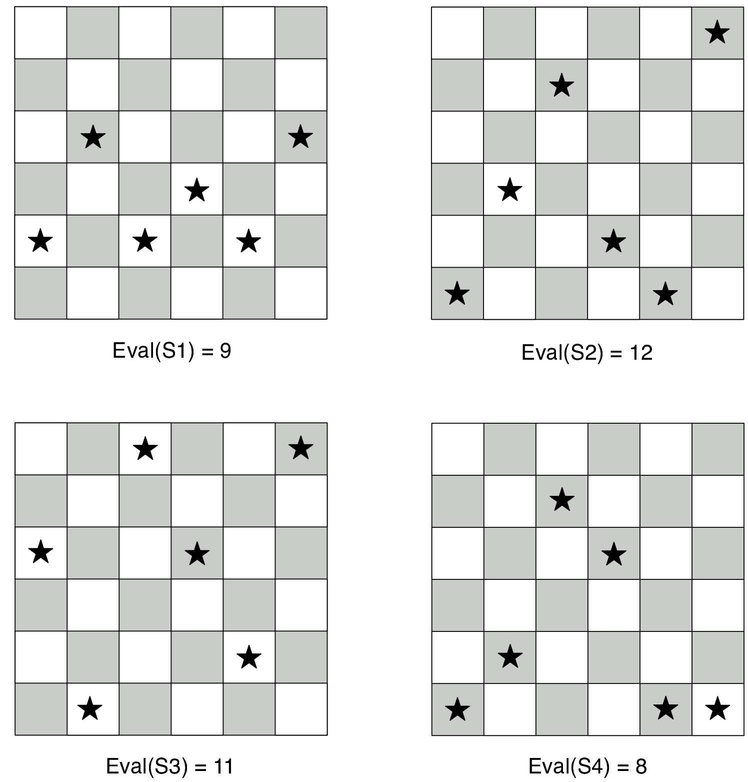

ANSWER: The probabilities are 9/40, 12/40, 11/40, 8/40

6. In a Genetic Algorithm, each state of 6-queens can be represented as 6 digits, each indicating the position of the queen in that column. Which action in genetic algorithm (among {selection, cross-over, mutation}) is most similar to the successor function described in previous page?

   ANSWER: Mutation

3 Cross Validation

Suppose you are running a majority classifier on the following training set. The training set is shown below. It consists of 10 data points. Each data point has a class label of either 0 or 1. A majority classifier is defined to output the class label that is in the majority in the training set, regardless of the input. If there is a tie in the training set, then always output class label 1.

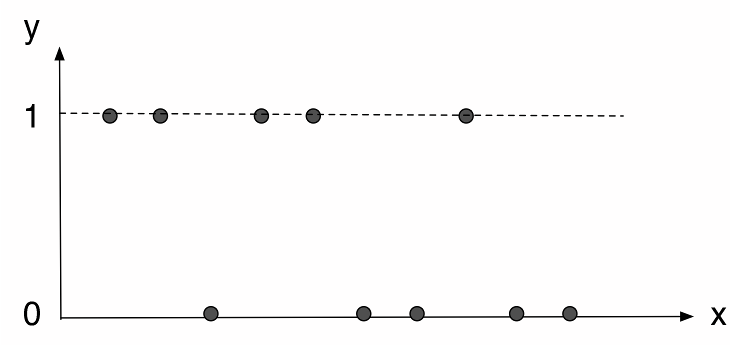

1. What is the training error? (report the error as a ratio)

   ANSWER: 5/10

2. What is the leave-one-out Cross-Validation error? (report the error as a ratio)

   ANSWER: 10/10

3. What is the two-fold Cross-Validation error? Assume the left 5 points belong to one partition while the right 5 points belong to the other partition. (report the error as a ratio)

   ANSWER: 8/10

4 Probabilistic Reasoning/Bayes Nets

1. If A and B are independent then $\sim A$ is independent of $\sim B$. True or False?
   Show the work supporting your answer. You might find the following statements useful:

$$\begin{aligned} p(A \vee B) &= p(A) + p(B) - p(A \wedge B) \\ p(\sim A) &= 1 - p(A) \\ \sim A \wedge \sim B &= \sim (A \vee B) \\ p(A|X) &= 1 - p(\sim A|X) \\ p(B|X) &= 1 - p(\sim B|X) \end{aligned}$$

   Answer: True

$$\begin{aligned} P(\sim A \wedge \sim B) &= P(\sim (A \vee B)) = 1 - P(A \vee B) = 1 - P(A) - P(B) + p(A \wedge B) \\ &= P(\sim A) - P(B) + p(A)p(B) = P(\sim A) - P(B)(1 - P(A)) \\ &= P(\sim A)(1 - P(B)) = P(\sim A)P(\sim B) \end{aligned}$$

   Other solutions such as showing that $P(\sim A | \sim B) = P(\sim A)$ were also possible.

2. Two students A and B are both registered for a certain course. Student A attends the class 80% of the time. Student B attends the class 60% of the time. Suppose their absences are independent.

   (a) What is the probability that neither show up to class on any given day?

   ANSWER: $P(\sim A \wedge \sim B) = P(\sim A)P(\sim B) = .2 * .4 = .08$

   (b) What is the probability that at least one of them is in class on any given day?
   $P(A \vee B) = P(\sim (\sim A \wedge \sim B)) = 1 - P(\sim A \wedge \sim B) = 1 - .08 = .92$
   Also, $P(A \vee B) = P(A) + P(B) - P(A \wedge B) = .8 + .6 - .8 * .6 = .92$
   ANSWER: 0.92

   Suppose there is also a student C who always comes to class if and only if student A or student B (or both) show up.

   (c) is the absence of A still independent of the absence of B? (yes, no)

   ANSWER: Yes, the knowledge about B still doesn't tell us anything about A (unless we know C).

   (d) construct a Bayes Net to show the relationships of A, B and C. Indicate the necessary CPTs (Conditional Probability Tables)

   ANSWER:

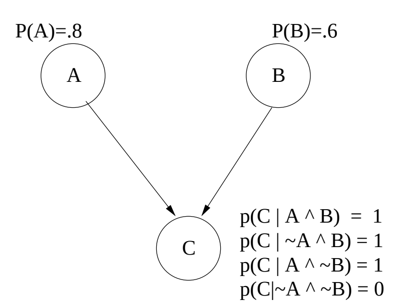

(e) is A conditionally independent of B given C? (yes, no)

ANSWER: No, A and B are dependent given C since C unblocks the path between A and B.

(f) suppose you know that C came to class, what is the probability of A coming if you know that B showed up too?

ANSWER: 0.8

Since B coming to class could fully explain the appearence of C, the probability of $P(A|B \wedge C) = P(A) = .8$. The result can also be obtained from the probabilities:

$$
\begin{aligned}
P(A|B \wedge C) &= \frac{P(A \wedge B \wedge C)}{P(B \wedge C)} = \frac{P(A)P(B)P(C|A \wedge B)}{P(A \wedge B \wedge C) + P(\sim A \wedge B \wedge C)} \\
&= \frac{P(A)P(B)}{P(C|A \wedge B)P(A)P(B) + P(C| \sim A \wedge B)P(\sim A)P(B)} \\
&= \frac{P(A)}{P(A) + P(\sim A)} = P(A)
\end{aligned}
$$

5 Neural Networks

1. Draw a linear perceptron network and calculate corresponding weights to correctly classify 4 points below. The output node returns 1 if the weighted sum is greater than or equal to the threshold (.5). If it looks too complicated you are probably wrong. You are allowed to make use of a ”constant 1” unit input.

x	y	out
0	0	1
0	1	1
1	0	0
1	1	1

ANSWER:

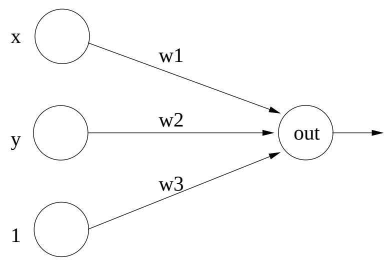

The dataset is linearly separable as shown below. Any set of weights such that $w_1 < 0$, $w_3 \in (0.5, .5 - w_1]$ and $w_2 \geq \max\{0, .5 - w_1 - w_3\}$ would have been a correct solution to the problem. A common solution was $w_1 = -1, w_2 = 1, w3 = 1$.

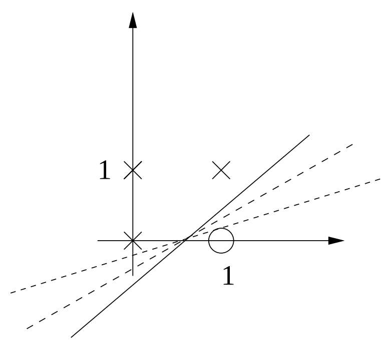

2. Is it possible to modify the network so that it will classify both - the dataset above **and** the one below with 100% accuracy?

x	y	out
0	0	0
0	1	1
1	0	1
1	1	1

   (a) by changing weights? (yes,no)
   ANSWER: NO, no neural network is able to perfectly classify a dataset that has conflicting labels (for example, record 0 0 is labeled 1 in the first set and 0 in the second)

   (b) by adding more layers? (yes,no)
   ANSWER: NO, the same reason as above.

6 Naive Bayes

Assume we have a data set with three binary input attributes, **A**, **B**, **C**, and one binary outcome attribute **Y**. The three input attributes, **A**, **B**, **C** take values in the set $\{0,1\}$ while the **Y** attribute takes values in the set $\{$True, False$\}$.

A	**B**	**C**	**Y**
0	1	1	True
1	1	0	True
1	0	1	False
1	1	1	False
0	1	1	True
0	0	0	True
0	1	1	False
1	0	1	False
0	1	0	True
1	1	1	True

Statistics about the Data Set
NOTE: We believe that some (but not all) of these statistics will be useful to you

	Y=T	**Y**=F
The fraction with **A**=0 is	4/6	1/4
The fraction with **A**=1 is	2/6	3/4
The fraction with **B**=0 is	1/6	2/4
The fraction with **B**=1 is	5/6	2/4
The fraction with **C**=0 is	3/6	0/4
The fraction with **C**=1 is	3/6	4/4

	Y=True	**Y**=False
Among records with **A**=0	4/5	1/5
Among records with **A**=1	2/5	3/5
Among records with **B**=0	1/3	2/3
Among records with **B**=1	5/7	2/7
Among records with **C**=0	3/3	0/3
Among records with **C**=1	3/7	4/7

If we are using a Naive Bayes Classifier with one binary valued output variable Y, the following theorem is true:

Theorem: A non-impossible set of input values, S, (i.e. a set of input values with $P(S) > 0$) will have an unambiguous predicted classification of $Y = \text{True} \Leftrightarrow P(Y = \text{True} \wedge S) > P(Y = \text{False} \wedge S)$

1. How would a Naive Bayes classifier classify the record (**A**=1,**B**=1,**C**=0)? (True/False)

   ANSWER: True

$$\begin{aligned} P(A=1|Y=True)P(B=1|Y=True)P(C=0|Y=True)P(Y=True) &= \frac{2}{6}*\frac{5}{6}*\frac{3}{6}*\frac{6}{10} = \frac{1}{12} \\ P(A=1|Y=False)P(B=1|Y=False)P(C=0|Y=False)P(Y=False) &= \frac{3}{4}*\frac{2}{4}*\frac{0}{4}*\frac{4}{10} = 0 \end{aligned}$$

   and $1/12 > 0$.

2. How would a Naive Bayes classifier classify the record (**A**=0,**B**=0,**C**=1)? (True/False)

   ANSWER: FALSE

$$\begin{aligned} P(A=0|Y=True)P(B=0|Y=True)P(C=1|Y=True)P(Y=True) &= \frac{4}{6}*\frac{1}{6}*\frac{3}{6}*\frac{6}{10} = 0.0333 \\ P(A=0|Y=False)P(B=0|Y=False)P(C=1|Y=False)P(Y=False) &= \frac{1}{4}*\frac{2}{4}*\frac{4}{4}*\frac{4}{10} = 0.05 \end{aligned}$$

   and $.03(3) < 0.05$.

3. How would a Naive Bayes classifier classify the record (**A**=0,**B**=0,**C**=0)? (True/False)

   ANSWER: TRUE

$$
\begin{aligned}
P(A=0|Y=True)P(B=0|Y=True)P(C=0|Y=True)P(Y=True) &= \frac{4}{6} * \frac{1}{6} * \frac{3}{6} * \frac{6}{10} = \frac{1}{30} \\
P(A=0|Y=False)P(B=0|Y=False)P(C=0|Y=False)P(Y=False) &= \frac{1}{4} * \frac{2}{4} * \frac{0}{4} * \frac{4}{10} = 0
\end{aligned}
$$

   and $\frac{1}{30} > 0$.

4. Would it be possible to add just one record to the data set that would result in a Naive Bayes classifier changing its classification of the record (**A**=1,**B**=0,**C**=1)? (Yes/No)

   ANSWER: No

   With the current data set the record (**A**=1,**B**=0,**C**=1) classifies to False since:
   $P(A=1|Y=True)P(B=0|Y=True)P(C=1|Y=True)P(Y=True) = \frac{2}{6} * \frac{1}{6}\frac{3}{6} * \frac{6}{10} = 0.0166$
   $< P(A=1|Y=False)P(B=0|Y=False)P(C=1|Y=False)P(Y=False) = \frac{3}{4} * \frac{2}{4}\frac{4}{4} * \frac{4}{10} = 0.15$

   If the record we add has **Y** $= True$ then the estimate of
   $P(A=1|Y=True)P(B=0|Y=True)P(C=1|Y=True)P(Y=True)$ would increase the most if the added record also has **A**=1, **B**=0, and **C**=1 in which case the estimate of it would become $\frac{3}{7} * \frac{2}{7} * \frac{4}{7} * \frac{7}{11} = 0.0445$. However this value is still less than, 0.15, the unchanged estimate of
   $P(A=1|Y=False)P(B=0|Y=False)P(C=1|Y=False)P(Y=False)$.

   If the record we add has **Y** $= False$ then the estimate of
   $P(A=1|Y=False)P(B=0|Y=False)P(C=1|Y=False)P(Y=False)$ would decrease the most if the added record has **A**=0, **B**=1, and **C**=0 in which case the estimate of it would become $\frac{3}{5} * \frac{2}{5} * \frac{4}{5} * \frac{5}{11} = 0.0873$ which would still be greater than 0.0166 the unchanged estimate of $P(A=1|Y=True)P(B=0|Y=True)P(C=1|Y=True)P(Y=True)$.

7 Decision Tree

For this problem we will use the same data set below as in the Naive Bayes question. Again assume we have three binary input attributes, **A**, **B**, **C**, and one binary outcome attribute **Y**. The three input attributes, **A**, **B**, **C** take values in the set $\{0,1\}$ while the **Y** attribute takes values in the set $\{$True, False$\}$.

A	**B**	**C**	**Y**
0	1	1	True
1	1	0	True
1	0	1	False
1	1	1	False
0	1	1	True
0	0	0	True
0	1	1	False
1	0	1	False
0	1	0	True
1	1	1	True

Specific Conditional Entropies		
H(**Y**|**A**=0)=0.72	H(**Y**|**A**=0,**B**=0)=0.00	H(**Y**|**A**=1,**C**=0)=0.00
H(**Y**|**A**=1)=0.97	H(**Y**|**A**=0,**B**=1)=0.81	H(**Y**|**A**=1,**C**=1)=0.81
H(**Y**|**B**=0)=0.92	H(**Y**|**A**=1,**B**=0)=0.00	H(**Y**|**B**=0,**C**=0)=0.00
H(**Y**|**B**=1)=0.86	H(**Y**|**A**=1,**B**=1)=0.92	H(**Y**|**B**=0,**C**=1)=0.00
H(**Y**|**C**=0)=0.00	H(**Y**|**A**=0,**C**=0)=0.00	H(**Y**|**B**=1,**C**=0)=0.00
H(**Y**|**C**=1)=0.99	H(**Y**|**A**=0,**C**=1)=0.92	H(**Y**|**B**=1,**C**=1)=0.97

1. Which attribute (**A**,**B**, or **C**) has the highest information gain?
   ANSWER: Attribute **C** since
   $H(Y|A) = P(A=0)*H(Y|A=0) + P(A=1)*H(Y|A=1) = \frac{5}{10}*0.72 + \frac{5}{10}*0.97 = 0.845$
   $H(Y|B) = P(B=0)*H(Y|B=0) + P(B=1)*H(Y|B=1) = \frac{3}{10}*0.92 + \frac{7}{10}*0.86 = 0.878$
   $H(Y|C) = P(C=0)*H(Y|C=0) + P(C=1)*H(Y|C=1) = \frac{3}{10}*0.00 + \frac{7}{10}*0.99 = 0.693$

   thus $H(Y) - H(Y|X)$ will be greatest when X=**C**.

2. Construct the full decision tree for this problem, without doing any pruning. If at any point you have a choice between splitting on two equally desirable attributes, choose the one that comes first alphabetically. If there is a tie on how to label a leaf, then choose True.

   ANSWER:

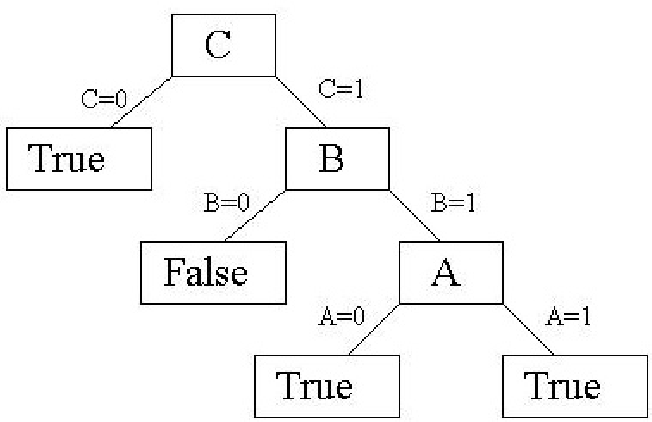

Note the first attribute split on in the case **C**=1 is **B** instead of **A** since
$H(Y|A,C=1) = P(A=0|C=1)H(Y|A=0,C=1) + P(A=1|C=1)H(Y|A=1,C=1) = \frac{3}{7}*0.92 + \frac{4}{7}*0.81 = .857$
$> H(Y|B,C=1) = P(B=0|C=1)H(Y|B=0,C=1) + P(B=1|C=1)H(Y|B=1,C=1) = \frac{2}{7}*0.00 + \frac{5}{7}*.97 = .693$

3. How would your decision tree classify the record (**A**=0,**B**=0,**C**=1)? (True/False)

   ANSWER: FALSE

4. How would your decision tree classify the record (**A**=1,**B**=0,**C**=0)? (True/False)

   ANSWER: TRUE

5. If you pruned all nodes from your decision tree except the root node, now how would your decision tree classify the record (**A**=0,**B**=0,**C**=1)? (True/False) Again, assume any ties are broken by choosing True. (NOTE: This was clarified during the exam to mean that the tree would split on one attribute and then classify)

   ANSWER: FALSE

6. If you pruned all nodes from your decision tree except the root node, now how would your decision tree classify the record (**A**=1,**B**=0,**C**=0)? (True/False) Again, assume any ties are broken by choosing True.

   ANSWER: TRUE

8 K-Means

The circles in the numbered boxes below represent the data points. In the first numbered box there are three squares, representing the initial location of cluster centers of the k-means algorithm. Trace through the first nine iterations of the k-means algorithm or until convergence is reached, whichever comes first. For each iteration draw three squares corresponding to the location of the cluster centers during that iteration. (NOTE: It is not necessary to draw the exact location of the squares, but it should be clear from your placement of the squares that you understand how k-means performs quantitatively)

ANSWER:

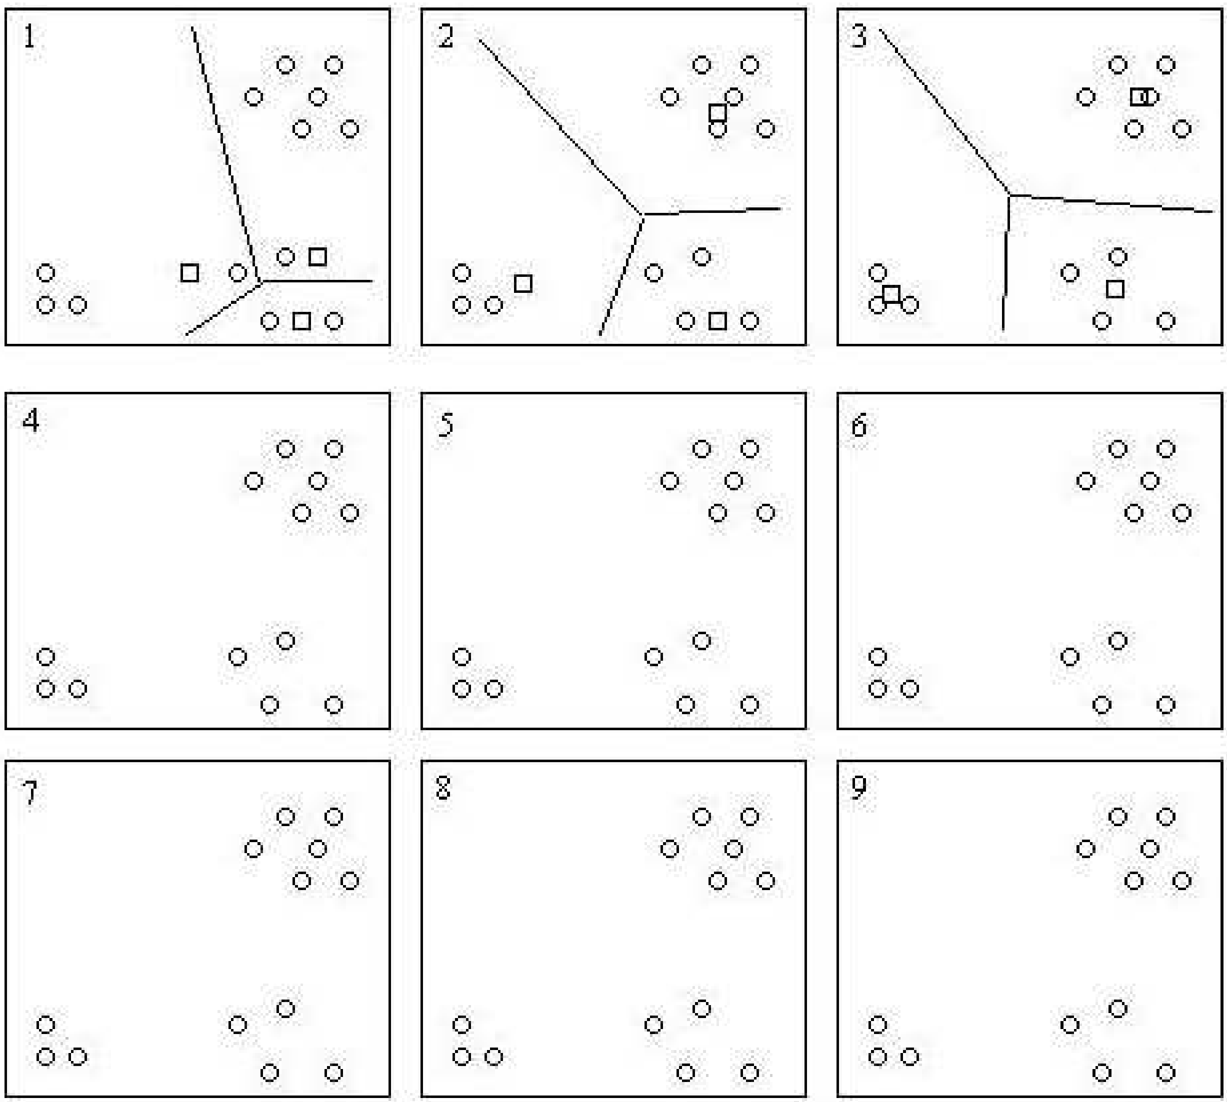

9 Reinforcement Learning

9.1 Q-Learning

Perform Q-learning for a system with two states and two actions, given the following training examples. The discount factor is $\gamma = 0.5$ and the learning rate is $\alpha = 0.5$. Assume that your Q-table is initialized to 0.0 for all values.

(Start = S_1, Action = a_1, Reward = 10, End = S_2)

	S_1	S_2
a_1	5.0	0.0
a_2	0.0	0.0

(Start = S_2, Action = a_2, Reward = -10, End = S_1)

	S_1	S_2
a_1	5.0	0.0
a_2	0.0	-3.75

(Start = S_1, Action = a_2, Reward = 10, End = S_1)

	S_1	S_2
a_1	5.0	0.0
a_2	6.25	-3.75

(Start = S_1, Action = a_1, Reward = 10, End = S_1)

	S_1	S_2
a_1	9.0625	0.0
a_2	6.25	-3.75

What is the policy that Q-learning has learned?

$$\pi(1) = a_1 \qquad \pi(2) = a_1$$

9.2 Certainty Equivalent Learning

In the diagram below, draw the state transitions and label them according to the values that would be discovered by Certainty Equivalent learning given follwing training examples.

(Start = S_1, Action = a_1, Reward = 10, End = S_2)
(Start = S_2, Action = a_2, Reward = -10, End = S_1)
(Start = S_1, Action = a_2, Reward = 10, End = S_1)
(Start = S_1, Action = a_1, Reward = 10, End = S_1)
(Start = S_1, Action = a_2, Reward = 10, End = S_1)
(Start = S_1, Action = a_1, Reward = 10, End = S_2)
(Start = S_2, Action = a_1, Reward = -10, End = S_2)
(Start = S_2, Action = a_2, Reward = -10, End = S_2)
(Start = S_2, Action = a_2, Reward = -10, End = S_1)

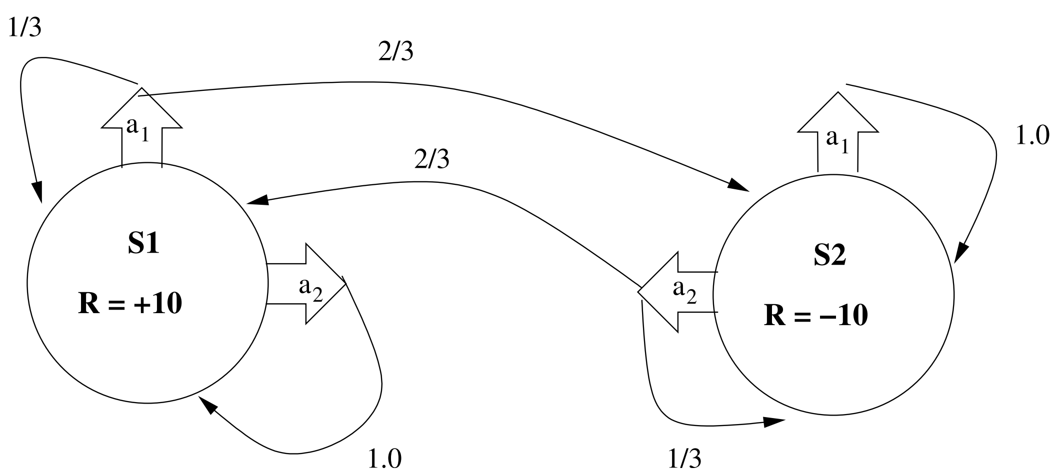

What is the policy that CE has learned?

$$\pi(1) = a_2 \qquad \pi(2) = a_2$$

10 Markov Decision Processes

You are a wildy implausible robot who wanders among the four areas depicted below. You hate rain and get a reward of -30 on any move that starts in the deck and -40 on any move that starts in the Garden. You like parties, and you are indifferent to kitchens.

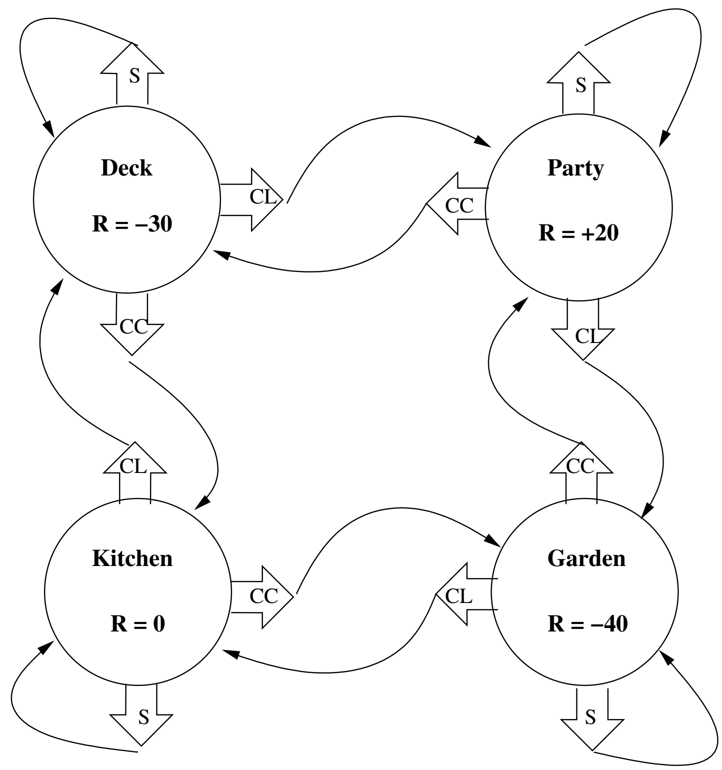

Actions: All states have three actions: **Clockwise (CL), Counter-Clockwise (CC), Stay (S)**. Clockwise and Counter-Clockwise move you through a door into another room, and Stay keeps you in the same location. All transitions have are deterministic (probability 1.0).

1. How many distinct policies are there for this MDP?
   $3^4 = 81$

2. Let $J^\star(\text{Room})$ = expected discounted sum of future rewards assuming you start in "Room" and subsequently act optimally. Assuming a discount factor $\gamma = 0.5$, give the $J^\star$ values for each room.

   ```
   By eyeballing the problem and quickly checking the CL and S options in the
   Kitchen, you can quickly determine that the optimal policy is:
   ```

   π(Deck) = CL, π(Party) = S, π(Kitchen) = S, π(Garden) = CC

   ```
   and write
   ```

   $J^\star(\text{Deck}) = -30 + 0.5 J^\star(\text{Party})$
   $J^\star(\text{Party}) = +20 + 0.5 J^\star(\text{Party})$
   $J^\star(\text{Kitchen}) = 0 + 0.5 J^\star(\text{Kitchen})$
   $J^\star(\text{Garden}) = -40 + 0.5 J^\star(\text{Party})$

   ```
   Solving for the Js gives:
   ```

$J^\star(\text{Deck}) = -10$
$J^\star(\text{Party}) = 40$
$J^\star(\text{Kitchen}) = 0$
$J^\star(\text{Garden}) = -20$

3. The optimal policy when the discount factor, γ, is small but non-zero (e.g. $\gamma = 0.1$) is different from the optimal policy when γ is large (e.g. $\gamma = 0.9$).

   If we began with $\gamma = 0.1$, and then gradually increased γ, what would be the threshold value of γ above which the optimal policy would change?

   As the discount factor increases,the policy changes from S in the Kitchen to CL. This change occurs at the point where the value of taking action S in the Kitchen is equal to the value of taking action CL:

   $J^S(\text{Kitchen}) = 0 + \gamma J^S(\text{Kitchen}) = 0$
   $J^{CL}(\text{Kitchen}) = 0 + \gamma J^\star(\text{Deck})$

   We already know that the optimal policy in the Deck is CL, regardless of the discount factor:

   $J^\star(\text{Deck}) = -30 + \gamma J^\star(\text{Party})$

   and we know

   $J^\star(\text{Party}) = +20 + \gamma J^\star(\text{Party})$
   $J^\star(\text{Party}) = \frac{20}{1-\gamma}$

   So we want to solve

   $\gamma(-30 + \gamma(\frac{20}{1-\gamma})) = 0$
   $-30 + \gamma(\frac{20}{1-\gamma}) = 0$
   $20\gamma = 30(1-\gamma)$
   $\gamma = 3/5 = 0.6$