Name: _____ ID # : _____

CSE 473 FALL 2000 FINAL EXAM

Due date: Monday, December 11, 2000, 12:00 pm

Turn-in: Return the completed final to David's office (Sieg 222) between **10am and 12pm on December 11**, or email him to arrange otherwise.

Instructions:

This is an open-book final. That means you can use any source of information, short of asking other people. Do not discuss the questions with anyone except Pedro and David.

This exam has 8 problems on 14 pages for a total of 40 points.

Make sure you write your name and ID number on every page.

Justify all your answers. This will allow you to get partial credit even if you do not complete the question or the final answer is incorrect. Please write clearly and concisely. Use the back of the page and additional pages if you need to.

Good luck!

Name: _____ ID # : _____

Queston #1: Consider the following search tree, where A is the initial state, G is the goal state, each node/state points to its successors, each branch is labeled with the cost of traversing it, and each node/state is labeled with a heuristic estimate of the cost of reaching the goal from it:

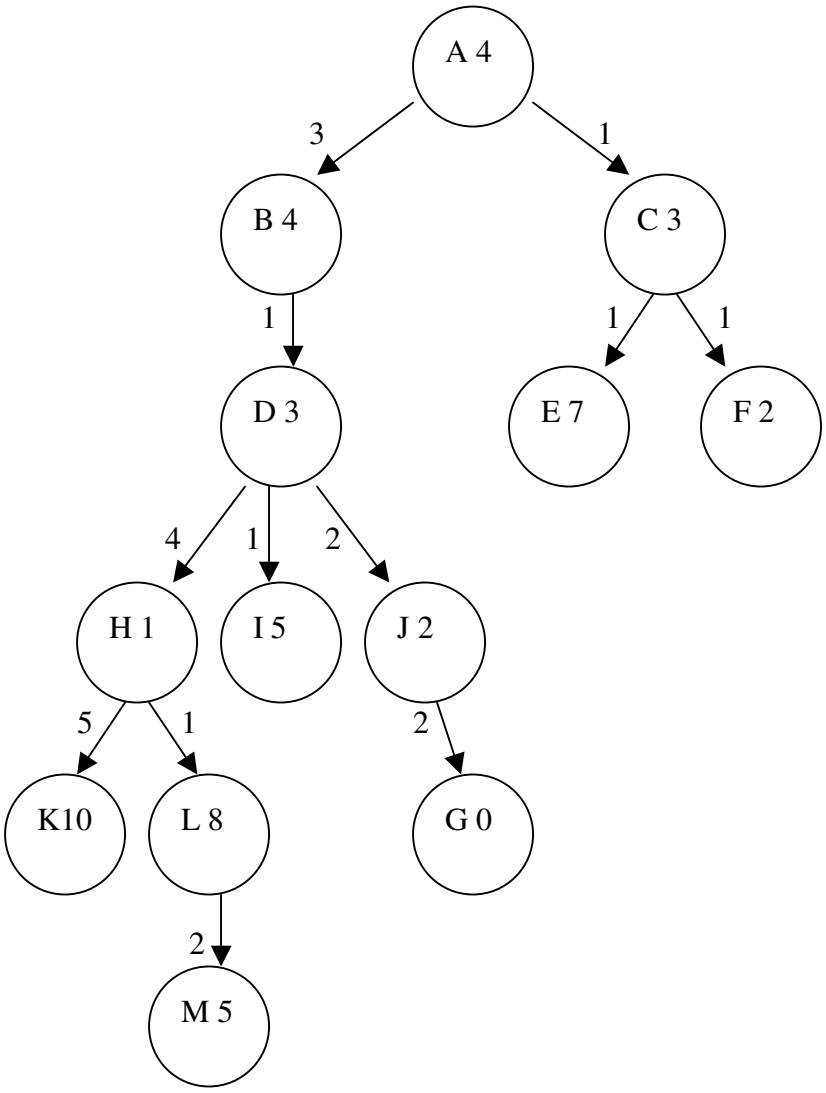

Question #1, continued:

a) For each search method below, show the order in which nodes are visited by each of the following search methods:

i) (1 point) Breadth-first

ii) (1 point) Depth-first

iii) (2 points) Uniform-cost

iv) (2 points) Greedy

v) (2 points) A*

b) (2 points) Is the heuristic admissible?

Name: _____ ID # : _____

Question #2: Consider a world in which there are only six propositions, A, B, C, D, E and F. How many models are there for the following sentences?

a) (2 points) D => F

b) (2 points) A ^ (C v E)

Question #3: Consider the following statements:

All farm animals are mammals or birds.
All mammals produce milk.
All birds lay eggs.
If something produces milk, then it's a source of revenue.
If something lays eggs, then it's a source of revenue.

a) (2 points) Represent them in first-order logic.

Question #3, continued:

b) (4 points) Use these statements to prove by resolution that if Daisy is a farm animal then she is a source of revenue.

Name: _____ ID # : _____

Question #4: Consider the following Bayesian network structure, where A, B, C and D are Boolean variables:

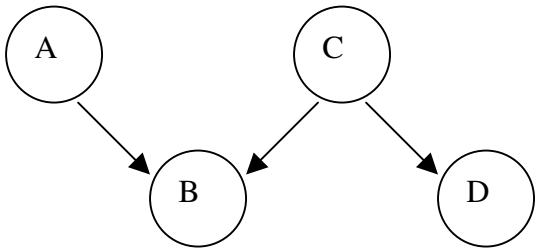

a) (1 point) Is A independent of D?

b) (1 point) Is A independent of D given B?

c) (1 point) Is A independent of D given C?

Question 4, continued:

d) (2 points) Suppose you are given the following set of training examples:

A	**B**	**C**	**D**
0	1	0	1
0	0	1	1
1	0	0	0
1	0	1	1

Complete the Bayesian network by computing the probability tables induced by this training set.

Question #4, continued:

e) (3 points) Using this network, compute P(C = 1 | B = 0).

Question #5: (2 points) Consider the following events:

A = You receive a million dollars.
B = You receive a utility of 0.2.
C = You receive a utility of 0.7.

If you are indifferent between A and a lottery between B and C where your chances of winning B are 0.15 and your chances of winning C are 0.85, what is the utility of a million dollars to you?

Name: _____ ID # : _____

Question #6: Let the instance space be the region $0 < X < 40$, $0 < Y < 30$ on the XY plane, and let + and - represent positive and negative training examples, respectively. Suppose you are given the following training set:

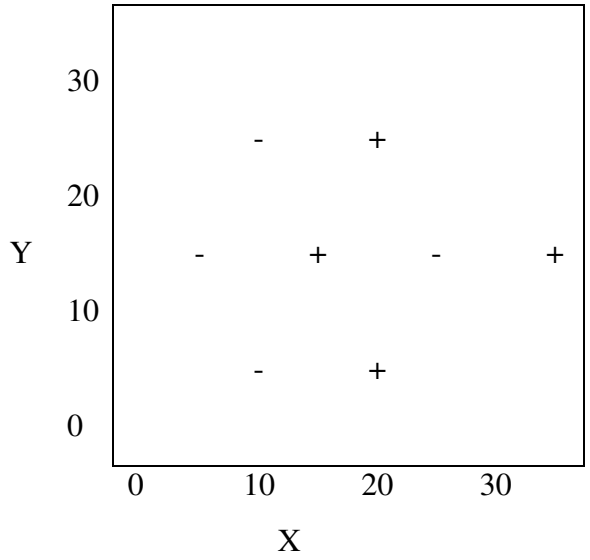

a) (1 point) Draw the Voronoi diagram corresponding to this training set, using Euclidean distance.

Question #6, continued:

b) (1 point) Draw the smallest decision tree that represents the same concept as the nearest-neighbor algorithm with this training set and Euclidean distance.

Question #6, continued:

c) (1 point) Is this concept learnable using a perceptron?

d) (1 point) Suppose the hypothesis space is the set of hypotheses of the form "If $x > a$, then +; otherwise -" or "If $y > b$, then +; otherwise -", where a and b are integers and $0 < a < 40$, $0 < b < 30$. Suppose the classes of the examples at (15, 15) and (25, 15) are switched. What are the S and G frontiers produced by the version space candidate elimination algorithm?

Question #7: (3 points) Let S be a set of training examples, and let A be one of their attributes. Show that, if A divides S into two subsets of equal size, then its information gain on S cannot be negative.

Question #8: (3 points) Consider the EM algorithm for learning a mixture of Gaussians. The algorithm we saw in class assumes that an example is generated in two steps:

1. Choose with uniform probability which Gaussian the example belongs to.
2. Generate the example's value according to that Gaussian.

How would you modify the algorithm to handle the case where the Gaussians are not chosen with uniform probability (i.e., some Gaussians are more likely to be chosen than others)? (You can continue to assume that the number of Gaussians is known, and that the variance is the same for all Gaussians and known.)