CS-171, Intro to A.I. — Final Exam — Fall Quarter, 2014

YOUR NAME: ____

YOUR ID: ________ ID TO RIGHT:________ ROW:_______ SEAT NO.: ________

The exam will begin on the next page. Please, do not turn the page until told.

When you are told to begin the exam, please check first to make sure that you have all 12 pages, as numbered 1-12 in the bottom-left corner of each page.

The exam is closed-notes, closed-book. No calculators, cell phones, electronics.

Please clear your desk entirely, except for pen, pencil, eraser, an optional blank piece of paper (for optional scratch pad use), and an optional water bottle.

<u>Please turn off all cell phones now.</u>

This page summarizes the points available for each question so you can plan your time.

1. (12 pts total, 2 pts each) k-NearestNeighbor (k-NN) and Cross-validation.
2. (4 pts total, 1 pt each) Task Environment.
3. (10 pts total, 1/2 pt each) Search Properties.
4. (10 pts total, 1 pt each) Probability Rules and Independence.
5. (14 pts total, 2 pts each) Knowledge Representation in FOPL.
6. (5 pts total) Hierarchical Agglomerative Clustering.
7. (5 pts total) k-Means Clustering.
8. (10 points total, 2 pts each) Constraint Satisfaction Problems.
9. (10 pts total, 1 pt each) State-Space Search.
10. (5 pts total, -1 for each wrong answer, but not negative) Mini-Max, Alpha-Beta Pruning.
11. (5 pts total) Bayesian Networks.
12. (10 pts total) Christmas Angel Resolution Theorem Proving in Propositional Logic.

The Exam is printed on both sides to save trees! Work both sides of each page!

1. (12 pts total, 2 pts each) k-NearestNeighbor (k-NN) and Cross-validation. Consider this training data set. Examples are A-E, the single attribute is X, and class labels are 0 or 1.

Example	A	B	C	D	E
Attribute Value (X)	0.1	0.6	0.8	2.0	3.0
Class Label	0	0	0	1	1

1.a. (2 pts) Using 1-NearestNeighbor, what class label would be assigned to unseen example F, which has attribute value $X_F = 0.3$? (Write 0 or 1) ________

1.b. (2 pts) Using 3-NearestNeighbor, what class label would be assigned to unseen example F, which has attribute value $X_F = 0.3$? (Write 0 or 1) ________

1.c. (2 pts) Using 1-NearestNeighbor, what class label would be assigned to unseen example G, which has attribute value $X_G = 1.5$? (Write 0 or 1) ________

1.d. (2 pts) Using 3-NearestNeighbor, what class label would be assigned to unseen example G, which has attribute value $X_G = 1.5$? (Write 0 or 1) ________

1.e. (2 pts) Using 1-NearestNeighbor and 5-fold Cross-Validation, what is the cross-validated accuracy of 1-NearestNeighbor on this data set? (Write a fraction, as N/5) _____________

1.f. (2 pts) Using 3-NearestNeighbor and 5-fold Cross-Validation, what is the cross-validated accuracy of 3-NearestNeighbor on this data set? (Write a fraction, as N/5) _____________

2. (4 pts total, 1 pt each) Task Environment. Your book defines a task environment as a set of four things, with acronym PEAS. Fill in the blanks with the names of the PEAS components.
P___________ E___________ A___________ S___________

3. (10 pts total, 1/2 pt each) Search Properties. Fill in the values of the four evaluation criteria for each search strategy shown. Assume a tree search where b is the finite branching factor; d is the depth to the shallowest goal node; m is the maximum depth of the search tree; C* is the cost of the optimal solution; step costs are identical and equal to some positive ε; and in Bidirectional search both directions using breadth-first search.

Note that these conditions satisfy all of the footnotes of Fig. 3.21 in your book.

Criterion	Complete?	Time complexity	Space complexity	Optimal?
Breadth-First				
Uniform-Cost				
Depth-First				
Iterative Deepening				
Bidirectional (if applicable)				

****** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ******

4. (10 pts total, 1 pt each) Probability Rules and Independence.
Consider the following full joint distribution for Boolean variables A, B, and C:

A	***B***	***C***	***P(a,b,c)***
t	*t*	*t*	*0.03*
t	*t*	*f*	*0.12*
t	*f*	*t*	*0.17*
t	*f*	*f*	*0.18*
f	*t*	*t*	*0.03*
f	*t*	*f*	*0.12*
f	*f*	*t*	*0.24*
f	*f*	*f*	*0.11*

Calculate the following probabilities (write a number from the interval [0,1]):

4.a. (1 pt) *P(A = f) =* _____________

4.b. (1 pt) *P(B = t) =* _____________

4.c. (1 pt) *P(B = t, C = t) =* ______________

4.d. (1 pt) *P(A = f, C = t) =* ______________

4.e. (1 pt) *P(A = t | B = t) =* ______________

4.f. (1 pt) *P(C = f | B = t) =* ______________

4.g. (1 pt) Are A and B independent of each other? (Y=Yes, N=No): ___________

4.h. (1 pt) Are B and C independent of each other? (Y=Yes, N=No): ___________

4.i (1 pt) Are B and C conditionally independent given A? (Y=Yes, N=No): ___________

4.j. (1 pt) Are A and C conditionally independent given B? (Y=Yes, N=No): __________

5. (14 pts total, 2 pts each) Knowledge Representation in FOPL. Consider a vocabulary with the following symbols:

Occupation(p, o) : Predicate. Person *p* has occupation *o*.
Customer(p1, p2) : Predicate. Person *p1* is a customer of person *p2*.
Boss(p1, p2) : Predicate. Person *p1* is a boss of person *p2*.
Doctor, Surgeon, Lawyer, Actor : Constants denoting occupations.
Emily, Joe : Constants denoting people.

Use these symbols to write the following assertions in first-order logic:

5.a. (2 pts) Emily is either a surgeon or a lawyer.

5.b. (2 pts) Joe is an actor, but he holds another job.

5.c. (2 pts) All surgeons are doctors.

5.d. (2 pts) Joe does not have a lawyer (i.e., Joe is not a customer of any lawyer).

5.e. (2 pts) Emily has a boss who is a lawyer.

5.f. (2 pts) There exists a lawyer all of whose clients are doctors (i.e., all of whose customers are doctors).

5.g. (2 pts) Every surgeon has a lawyer (i.e., every surgeon is a customer of a lawyer).

****** TURN PAGE OVER. EXAM CONTINUES ON THE REVERSE ******

6. (5 pts total) Hierarchical Agglomerative Clustering. Consider this training data set (it is the same as in problem 1). Examples are A-E, the single attribute is X, and class labels are 0 or 1.

Example	A	B	C	D	E
Attribute Value (X)	0.1	0.6	0.8	2.0	3.0
Class Label	0	0	0	1	1

Draw the dendogram (clustering tree) that results from applying hierarchical agglomerative clustering to this data. When two clusters are merged, replace them with their cluster centroid, i.e., the statistical mean of all cluster members. This rule means, (1) each cluster is represented by its cluster centroid which is the numerical mean (average) of all of its cluster members; and (2) dissimilarity between clusters is computed as the distance between their cluster centroids using Euclidean distance. **(Note: A better measure of dissimilarity is the root-mean-squared-deviation [RMSD] of each cluster member from its cluster centroid; but that is infeasible in an exam like this.)** Label the cluster centroids by drawing an oval around the data points that are included in that cluster centroid. The first one is done for you as an example.

You are only obliged draw the clustering tree (dendogram) that results. You do not need to write in the Cluster Centroid and Dissimilarity information shown in the square box below, which is provided only for your information about how to work the problem.

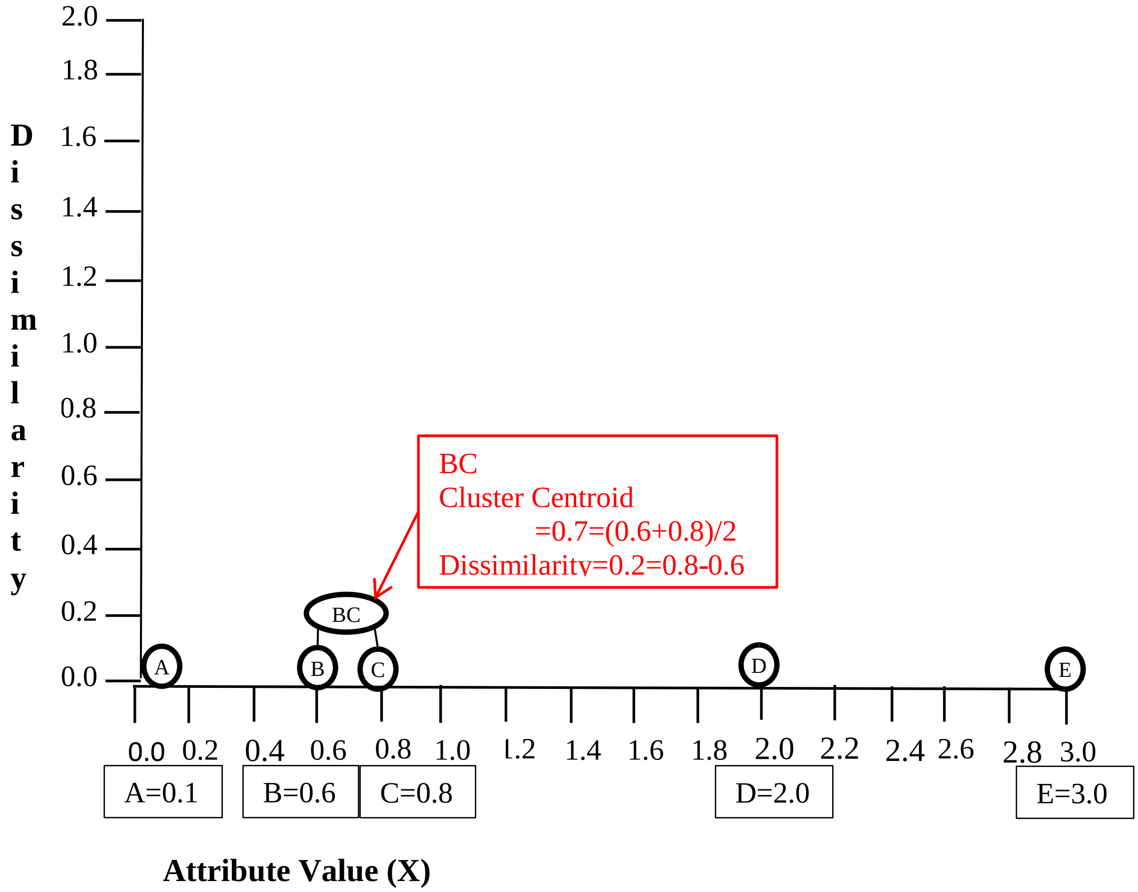

7. (5 pts total) k-Means Clustering. Consider this training data set (it is the same data set as in problems 1 and 6). Examples are A-E, the single attribute is X, and class labels are 0 or 1.

Example	A	B	C	D	E
Attribute Value (X)	0.1	0.6	0.8	2.0	3.0
Class Label	0	0	0	1	1

Apply k-Means Clustering to this data set for k=2, i.e., you will produce two data clusters.

7.a. (1 pt) You have randomly chosen two data points with which to initialize your two clusters. Randomly, you chose example A to initialize cluster #1 and example B to initialize cluster #2.

Write down the cluster assignments that result. Write C, D, and E in the blanks below according to which cluster they are assigned (A and B are already assigned).

cluster #1: ___A___________________ cluster #2: ___B___________________________

7.b. (1 pt) After assigning examples to clusters in 7.a, you recompute the cluster centroids (means) to be the mean (average) of the examples currently assigned to each cluster.

For each cluster, write the number that is the new cluster centroid (mean) of the examples that were assigned to that cluster in 7.a above.

cluster #1: ______________________ cluster #2: ____________________________

7.c. (1 pt) After recomputing the cluster centroids (means) in 7.b, you reassign the examples to the clusters to which they are closest (i.e., the example is assigned to the closest cluster centroid).

Write down the cluster assignments that result. Write A, B, C, D, and E in the blanks below according to which cluster they are assigned.

cluster #1: ______________________ cluster #2: ____________________________

7.d. (1 pt) After assigning examples to clusters in 7.c, you recompute the cluster centroids (means) to be the mean (average) of the examples currently assigned to each cluster.

For each cluster, write the number that is the new cluster centroid (mean) of the examples that were assigned to that cluster in 7.c above.

cluster #1: ______________________ cluster #2: ____________________________

7.e. (1 pt) After recomputing the cluster centroids (means) in 7.d, you reassign the examples to the clusters to which they are closest (i.e., the example is assigned to the closest cluster centroid).

Write down the cluster assignments that result. Write A, B, C, D, and E in the blanks below according to which cluster they are assigned.

cluster #1: ______________________ cluster #2: ____________________________

****** TURN PAGE OVER. EXAM CONTINUES ON THE REVERSE ******

8. (10 points total, 2 pts each) CONSTRAINT SATISFACTION PROBLEMS.

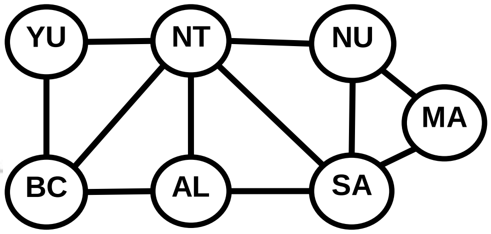

AL = Alberta
BC = British Columbia
MA = Manitoba
NT = Northwest Territories
NU = Nunavut
SA = Saskatchewan
YU = Yukon

You are a map-coloring robot assigned to color this western Canada map. Adjacent regions must be colored a different color (R=Red, B=Blue, G=Green). The constraint graph is shown.

8.a. (2 pts) FORWARD CHECKING. Cross out all values that would be eliminated by Forward Checking, after variable NT has just been assigned value G, as shown:

AL	BC	MA	NT	NU	SA	YU
R G B	R G B	R G B	G	R G B	R G B	R G B

8.b. (2 pts) ARC CONSISTENCY.
AL and MA have been assigned values, but no constraint propagation has been done. Cross out all values that would be eliminated by Arc Consistency (AC-3 in your book).

AL	BC	MA	NT	NU	SA	YU
B	R G B	R	R G B	R G B	R G B	R G B

8.c. (2 pts) MINIMUM-REMAINING-VALUES HEURISTIC. Consider the assignment below. YU is assigned and constraint propagation has been done. List all unassigned variables that might be selected by the Minimum-Remaining-Values (MRV) Heuristic:

____________________.

AL	BC	MA	NT	NU	SA	YU
R G B	G B	R G B	G B	R G B	R G B	R

8.d. (2 pts) DEGREE HEURISTIC. Consider the assignment below. (It is the same assignment as in problem 8.c. above.) YU is assigned and constraint propagation has been done. List all unassigned variables that might be selected by the Degree Heuristic:__________________.

AL	BC	MA	NT	NU	SA	YU
R G B	G B	R G B	G B	R G B	R G B	R

8.e. (2 pts) MIN-CONFLICTS HEURISTIC. Consider the complete but inconsistent assignment below. AL has just been selected to be assigned a new value during local search for a complete and consistent assignment. What new value would be chosen below for AL by the Min-Conflicts Heuristic?_____________________________.

AL	BC	MA	NT	NU	SA	YU
?	B	G	R	G	B	B

9. (10 pts total, 1 pt each) State-Space Search. Execute Tree Search through this graph (do not remember visited nodes, so repeated nodes are possible). It is not a tree, but you don't know that. Step costs are given next to each arc, and heuristic values are given next to each node (as h=x). The successors of each node are indicated by the arrows out of that node. As always, by convention, successors are returned in left-to-right order (successors of S are A,B,C; and successors of B are G,C; in that order).

The start node is S and the goal node is G. For each search strategy below, indicate **(1) the order** in which nodes are expanded, and **(2) the path** to the goal that was found, if any. Write "None" for the path if the goal was not found. The first one is done for you, as an example.

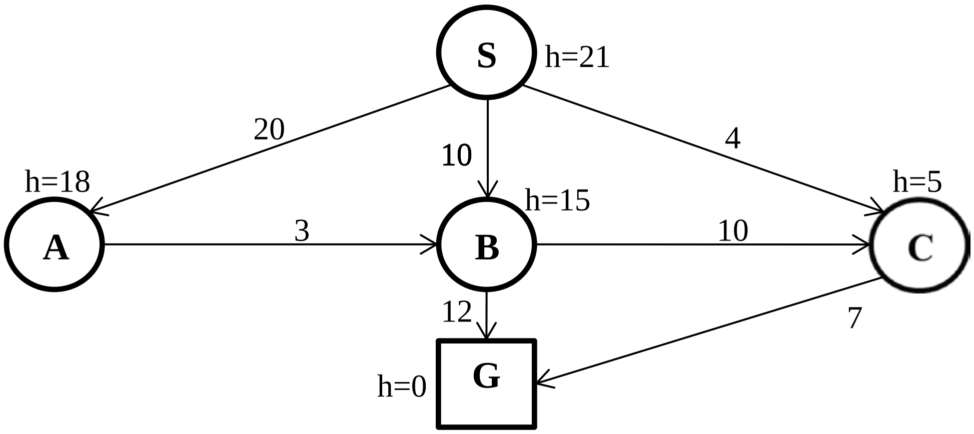

9.a. DEPTH-FIRST SEARCH:

9.a.(1) Order of expansion: S A B [G] ______________________________

9.a.(2) Path to goal found: S A B G ______________________________ Cost of path to goal: 35

9.b. BREADTH-FIRST SEARCH:

9.b.(1) Order of expansion: ______________________________

9.b.(2) Path to goal found: ______________________________ Cost of path to goal: __________

9.c. ITERATIVE DEEPENING SEARCH:

9.c.(1) Order of expansion: ______________________________

9.c.(2) Path to goal found: ______________________________ Cost of path to goal: __________

9.d. UNIFORM COST SEARCH:

9.d.(1) Order of expansion: ______________________________

9.d.(2) Path to goal found: ______________________________ Cost of path to goal: __________

9.e. GREEDY BEST FIRST SEARCH:

9.e.(1) Order of expansion: ______________________________

9.e.(2) Path to goal found: ______________________________ Cost of path to goal: __________

9.f. A* SEARCH:

9.f.(1) Order of expansion: ______________________________

9.f.(2) Path to goal found: ______________________________ Cost of path to goal: __________

****** TURN PAGE OVER. EXAM CONTINUES ON THE REVERSE ******

10. (5 pts total, -1 for each wrong answer, but not negative) Mini-Max, Alpha-Beta Pruning. In the game tree below it is **Max**'s turn to move. At each leaf node is the estimated score of that resulting position as returned by the heuristic static evaluator.

(1) Perform Mini-Max search and label each branch node with its value.

(2) What is Max's best move (A, B, or C)? ______________

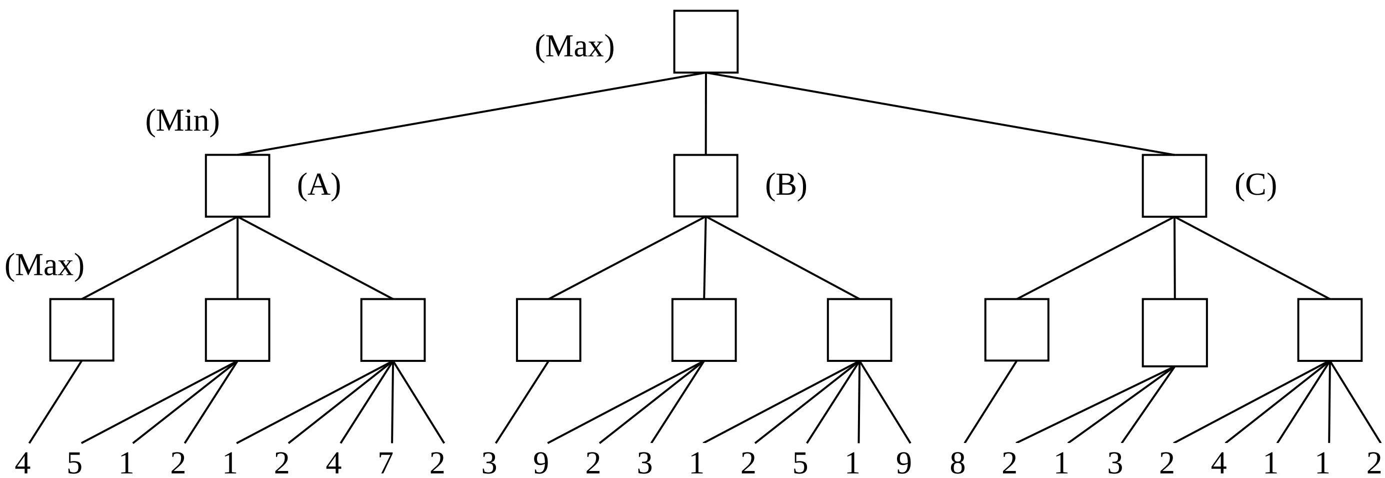

(3) Cross out each leaf node that would be pruned by alpha-beta pruning.

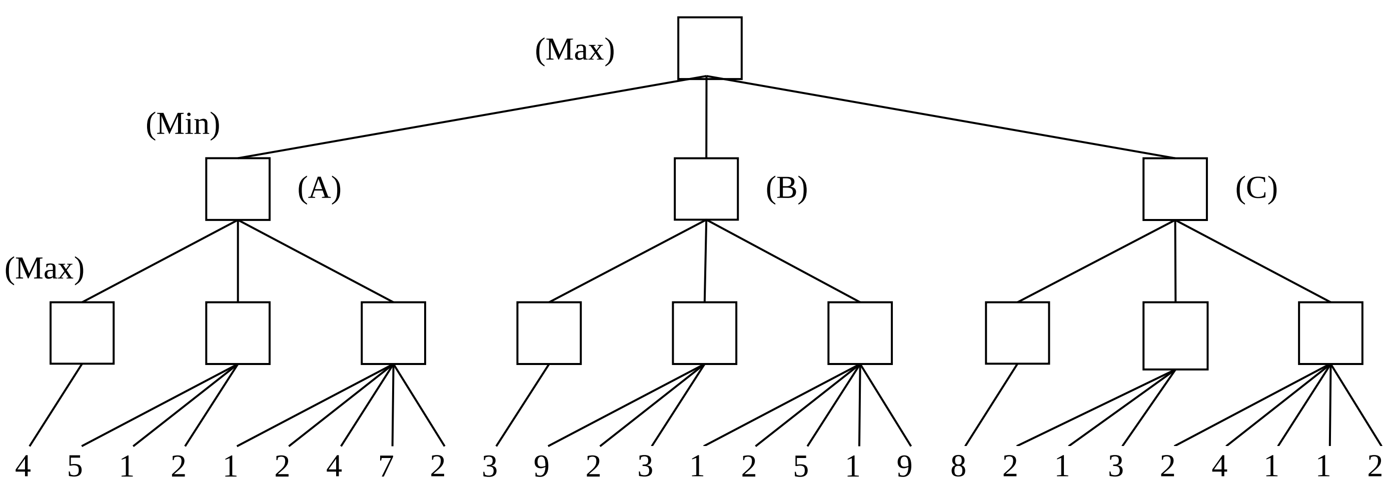

11. (5 pts total) Bayesian Networks.

11a. (1 pt) Write down the factored conditional probability expression that corresponds to the graphical Bayesian Network shown.

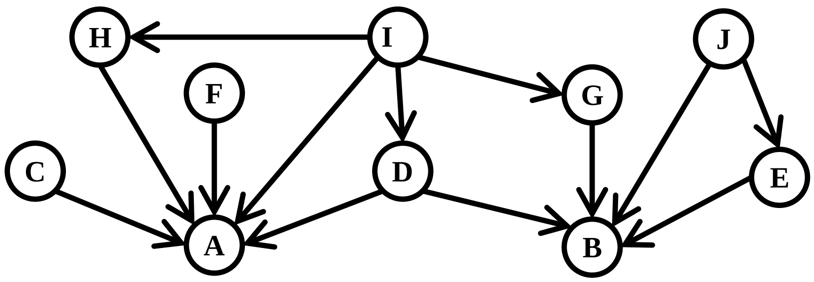

11b. (1 pt) Draw the Bayesian Network that corresponds to this conditional probability:

P(A | C,D,F,H) P(B | D,E,J) P(C | H) P(D | G,J) P(E) P(F | G,I) P(G | I, J) P(H) P(I) P(J)

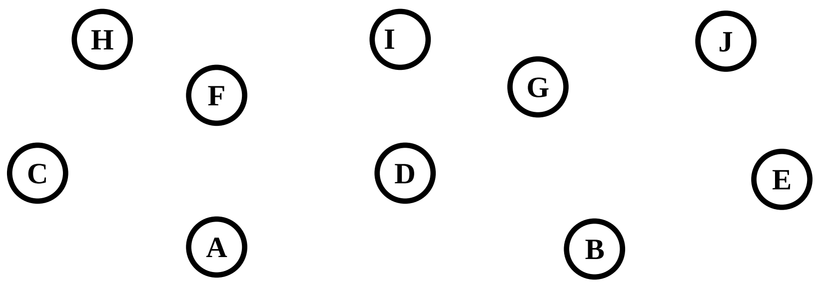

11.c. (3 pts) Below is the Bayesian network for the WetGrass problem [Fig. 14.12(a) in R&N].

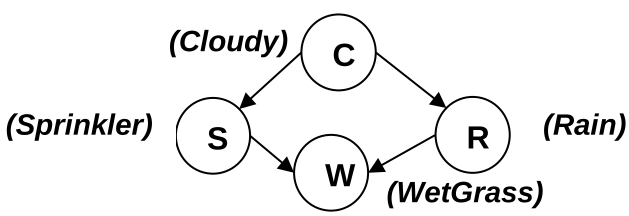

P(C)
.5

C	P(S)
t	.1
f	.5

C	P(R)
t	.8
f	.2

S	R	P(W)
t	*t*	.99
t	*f*	.90
f	*t*	.90
f	*f*	.00

Write down an expression that will evaluate to P(C=f ∧ R=f ∧ S=t ∧ W = t).

The probability tables show the probability that variable is True, e.g., P(M) means P(M=t). Express your answer as a series of numbers (numerical probabilities) separated by multiplication symbols. You do not need to carry out the multiplication to produce a single number (probability). **SHOW YOUR WORK.**

****** TURN PAGE OVER. EXAM CONTINUES ON THE REVERSE. ******

12. (10 pts total) Christmas Angel Resolution Theorem Proving in Propositional Logic.
(adapted from http://brainden.com/logic-puzzles.htm)

Four angels sat on the Christmas tree amidst other ornaments. Two had blue halos and two had gold halos (because we are University of California, of course our colors are blue and gold). You translate this fact into Propositional Logic (in prefix form) as:

/* Bi means angel i has a blue halo. */
(or (and B1 B2 (¬ B3) (¬ B4)) (and B1 (¬ B2) B3 (¬ B4))
(and B1 (¬ B2) (¬ B3) B4) (and (¬ B1) B2 B3 (¬ B4))
(and (¬ B1) B2 (¬ B3) B4) (and (¬ B1) (¬ B2) B3 B4)))

However, none of them could see above their head, and their views were obscured by branches.

Angel 1 reported, "Angels 2 & 3 have a blue halo and a gold halo, but I can't tell which."

Angel 2 reported, "Angels 3 & 4 have a blue halo and a gold halo, but I can't tell which."

Angel 3 reported, "Angel 4 has a blue halo."

You translate these facts into Propositional Logic (in prefix form) as:

(or (and B2 (¬ B3)) (and (¬ B2) B3)) (or (and B3 (¬ B4)) (and (¬ B3) B4)) B4

Angel 1 asks, "Is it true that I have a gold halo?"

You translate this query into Propositional Logic as "(¬ B1)" and form the negated goal as "B1." Your knowledge base (KB) in CNF plus negated goal (in clausal form) is:

(B1 B2 B3)	((¬ B1) (¬ B2) (¬ B3))
(B1 B2 B4)	((¬ B1) (¬ B2) (¬ B4))
(B1 B3 B4)	((¬ B1) (¬ B3) (¬ B4))
(B2 B3 B4)	((¬ B2) (¬ B3) (¬ B4))
(B2 B3)	((¬ B2) (¬ B3))
(B3 B4)	((¬ B3) (¬ B4))
B4	B1

Write a resolution proof that Angel 1 has a gold halo.

For each step of the proof, fill in the first two blanks with CNF sentences from KB that will resolve to produce the CNF result that you write in the third (resolvent) blank. The resolvent is the result of resolving the first two sentences. Add the resolvent to KB, and repeat. Use as many steps as necessary, ending with the empty clause.

The shortest proof I know of is only five lines long. (A Bonus Point for a shorter proof.)

Resolve __________________ with ___________________________ to produce: _______________

Resolve __________________ with ___________________________ to produce: _______________

Resolve __________________ with ___________________________ to produce: _______________

Resolve __________________ with ___________________________ to produce: _______________

Resolve __________________ with ___________________________ to produce: _______________

Resolve __________________ with ___________________________ to produce: _______________

Resolve __________________ with ___________________________ to produce: _______________

****** THIS IS THE END OF THE FINAL EXAM. HAPPY HOLIDAYS!! ******