CS-171, Intro to A.I. — Final Exam — Fall Quarter, 2012

NAME AND EMAIL ADDRESS: _____

YOUR ID: _____ ID TO RIGHT:_____ ROW:_____ NO. FROM RIGHT:_____

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 10 pages, as numbered 1-10 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.
Please turn off all cell phones now.**

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

1. (10 pts total) Decision Tree Classifier Learning.

2. (5 pts total, -1 pt each wrong answer, but not negative) Search Properties.

3. (10 pts total) Naïve Bayes Classifier Learning.

4. (15 pts total, 5 pts each, -1 each error, but not negative) Bayesian Networks.

5. (10 points total, 2 pts each) Constraint Satisfaction Problems.

6. (10 pts total, -1 for each error, but not negative) Alpha-Beta Pruning.

7. (10 pts total, -2 for each error, but not negative) Conversion to CNF.

8. (10 pts total, -2 for each error, but not negative) Resolution Theorem Proving.

9. (10 pts total, 1 pt each) State-Space Search.

10. (10 pts total, 2 pts each) English to FOL Conversion.

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

1. (10 pts total) Decision Tree Classifier Learning. You are a robot in a lumber yard, and must learn to discriminate Oak wood from Pine wood. You choose to learn a Decision Tree classifier. You are given the following examples:

Example	Density	Grain	Hardness	Class
Example #1	Heavy	Small	Hard	Oak
Example #2	Heavy	Large	Hard	Oak
Example #3	Heavy	Small	Hard	Oak
Example #4	Light	Large	Soft	Oak
Example #5	Light	Large	Hard	Pine
Example #6	Heavy	Small	Soft	Pine
Example #7	Heavy	Large	Soft	Pine
Example #8	Heavy	Small	Soft	Pine

1a. (2 pts) Which attribute would information gain choose as the root of the tree?

_____ .

1b. (4 pts) Draw the decision tree that would be constructed by recursively applying information gain to select roots of sub-trees, as in the Decision-Tree-Learning algorithm.

Classify these new examples as Oak or Pine using your decision tree above.
1c. (2 pts) What class is [Density=Light, Grain=Small, Hardness=Hard]? _____
1d. (2 pts) What class is [Density=Light, Grain=Small, Hardness=Soft]? _____

2. (5 pts total, -1 pt each wrong answer, but not negative) 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 use 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)				

3. (10 pts total) Naïve Bayes Classifier Learning. You are a robot in a lumber yard, and must learn to discriminate Oak wood from Pine wood. You choose to learn a Naïve Bayes classifier. You are given the following (noisy) examples:

Example	Density	Grain	Hardness	Class
Example #1	Light	Small	Hard	Oak
Example #2	Heavy	Large	Hard	Oak
Example #3	Heavy	Small	Soft	Oak
Example #4	Heavy	Small	Soft	Oak
Example #5	Light	Large	Hard	Pine
Example #6	Light	Small	Soft	Pine
Example #7	Heavy	Large	Soft	Pine
Example #8	Light	Large	Hard	Pine

Recall that Baye's rule allows you to rewrite the conditional probability of the class given the attributes as the conditional probability of the attributes given the class. As usual, α is a normalizing constant that makes the probabilities sum to one.

P(Class | Density, Grain, Hardness) = α P(Density, Grain, Hardness | Class) P(Class)

3a. (2 pts) Now assume that the attributes (Density, Grain, and Hardness) are conditionally independent given the Class. Rewrite the expression above, using this assumption of conditional independence (i.e., rewrite it as a Naïve Bayes Classifier expression).

3b. (4 pts total; -1 for each wrong answer, but not negative) Fill in numerical values for the following expressions. Leave your answers as common fractions (e.g., 1/4, 3/5).

P(Oak)= _____   P(Pine)= _____

P(Density=Light | Class=Oak)=_____   P(Density=Light | Class=Pine)=_____

P(Density=Heavy | Class=Oak)=_____   P(Density=Heavy | Class=Pine)=_____

P(Grain=Small | Class=Oak)=_____   P(Grain=Small | Class=Pine)=_____

P(Grain=Large | Class=Oak)=_____   P(Grain=Large | Class=Pine)=_____

P(Hardness=Hard | Class=Oak)=_____   P(Hardness=Hard | Class=Pine)=_____

P(Hardness=Soft | Class=Oak)=_____   P(Hardness=Soft | Class=Pine)= _____

3c. (2 pt each) Consider a new example (Density=Heavy ^ Grain=Small ^ Hardness=Hard). Write these class probabilities as the product of α and common fractions from above.

P(Class=Oak | Density=Heavy ^ Grain=Small ^ Hardness=Hard) = _____

P(Class=Pine | Density=Heavy ^ Grain=Small ^ Hardness=Hard) = _____

**4. (15 pts total, 5 pts each, -1 each error, but not negative) Bayesian Networks.**
4a. (5 pts) Draw the Bayesian Network that corresponds to this conditional probability:

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

4b. (5 pts) Write down the factored conditional probability expression that corresponds to the graphical Bayesian Network shown.
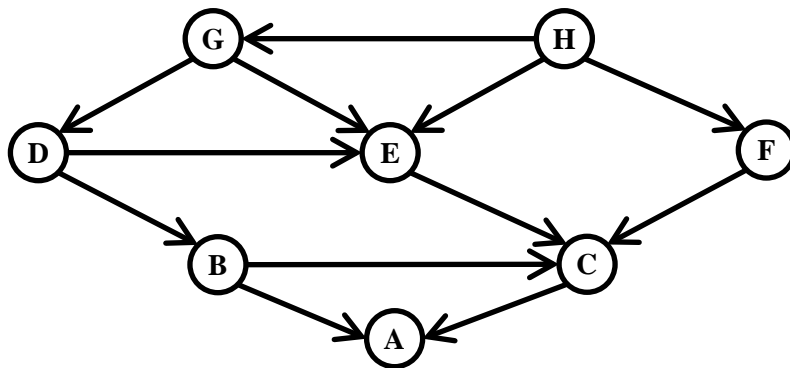
4.c. (5 pts) Shown below is the Bayesian network corresponding to the Burglar Alarm problem, P(J | A) P(M | A) P(A | B, E) P(B) P(E).
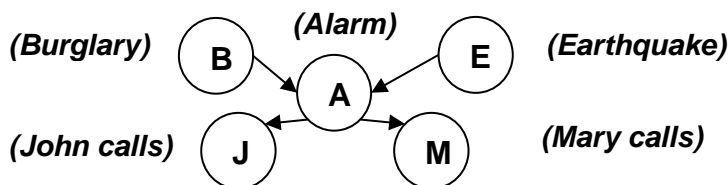
P(E)
.002

| A	P(M)
t	.70
f	.01

B	E	P(A)
t	t	.95
t	f	.94
f	t	.29
f	f	.001

A	P(J)
t	.90
f	.05

P(B)
.001

The probability tables show the probability that variable is True, e.g., P(M) means P(M=t).
Write down an expression that will evaluate to P(j=t ∧ m=f ∧ a=f ∧ b = f ∧ e = 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.**

5. (10 points total, 2 pts each) Constraint Satisfaction Problems.
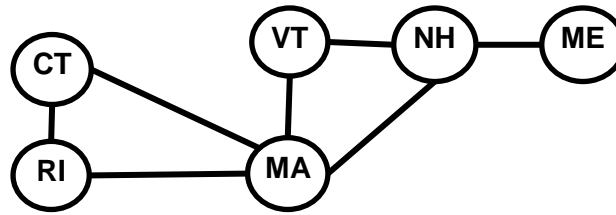
You are a map-coloring robot assigned to color this New England USA map. Adjacent regions must be colored a different color (R=Red, B=Blue, G=Green). The constraint graph is shown.

5a. (2pts total, -1 each wrong answer, but not negative) FORWARD CHECKING.
Cross out all values that would be eliminated by Forward Checking, after variable MA has just been assigned value R as shown:

CT	RI	MA	VT	NH	ME
R G B	R G B	R	R G B	R G B	R G B

5b. (2pts total, -1 each wrong answer, but not negative) ARC CONSISTENCY.
CT and RI 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).

CT	RI	MA	VT	NH	ME
R	G	R G B	R G B	R G B	R G B

5c. (2pts total, -1 each wrong answer, but not negative) MINIMUM-REMAINING-VALUES HEURISTIC. Consider the assignment below. RI is assigned and constraint propagation has been done. List all unassigned variables that might be selected by the Minimum-Remaining-Values (MRV) Heuristic:_____.

CT	RI	MA	VT	NH	ME
R B	G	R B	R G B	R G B	R G B

5d. (2pts total, -1 each wrong answer, but not negative) DEGREE HEURISTIC.
Consider the assignment below. (It is the same assignment as in problem 5c above.) RI is assigned and constraint propagation has been done. List all unassigned variables that might be selected by the Degree Heuristic:_____.

CT	RI	MA	VT	NH	ME
R B	G	R B	R B	R G B	R G B

5e. (2pts total) MIN-CONFLICTS HEURISTIC. Consider the complete but inconsistent assignment below. MA 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 MA by the Min-Conflicts Heuristic?_____.

CT	RI	MA	VT	NH	ME
B	G	?	G	G	B

6. (10 pts total, -1 for each error, but not negative) 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) Cross out each leaf node that would be pruned by alpha-beta pruning.
(3) What is Max's best move (A, B, or C)? _____

(Max)

(Min)

(A)   (B)   (C)

(Max)

4 3 8   2 2 2 4 1 2   8 5 2 3 4 9 5 1 8   3 9 4 3 2 4 1 1 2
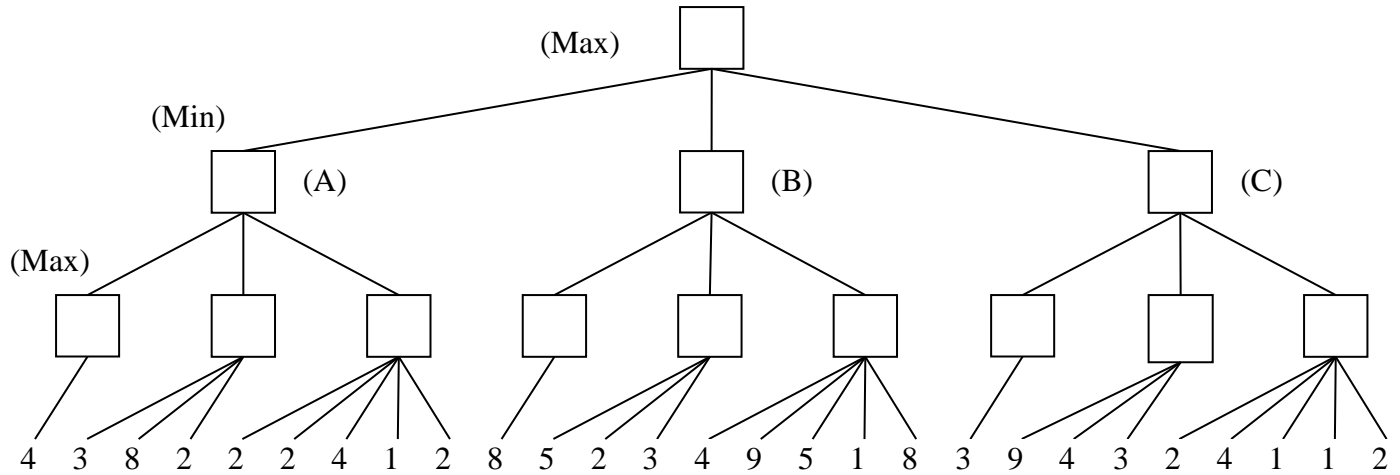
7. (10 pts total, -2 for each error, but not negative) Conversion to CNF. Convert this Propositional Logic wff (well-formed formula) to Conjunctive Normal Form and simplify.
Show your work (correct result, 0 pts; correct work, 10 pts).

$$[\neg (Q \Rightarrow P)] \Leftrightarrow P$$

8. (10 pts total, -2 for each error, but not negative) Resolution Theorem Proving. You are a robot in a logic-based question answering system, and must decide whether or not an input goal sentence is entailed by your Knowledge Base (KB). Your current KB in CNF is:

S1: (P Q)
S2: (¬P Q)
S3: (P ¬Q)
S4: (¬P R)

Your input goal sentence is: (P $\land$ Q $\land$ R).

8a. (2 pts) Write the negated goal sentence in CNF.

S5: _____

8b. (8 pts total, -2 for each error, but not negative) Use resolution to prove that the goal sentence is entailed by KB, or else explain why no such proof is possible. For each step of the proof, fill in Si and Sj with the sentence numbers of previous CNF sentences that resolve to produce the CNF result that you write in the resolvent blank. The resolvent is the result of resolving the two sentences Si and Sj. Use as many steps as necessary, ending by producing the empty clause; or else explain why no such proof is possible.

Resolve Si _____ with Sj _____ to produce resolvent **S6:** _____

Resolve Si _____ with Sj _____ to produce resolvent **S7:** _____

Resolve Si _____ with Sj _____ to produce resolvent **S8:** _____

Resolve Si _____ with Sj _____ to produce resolvent **S9:** _____

Resolve Si _____ with Sj _____ to produce resolvent **S10:** _____

Resolve Si _____ with Sj _____ to produce resolvent **S10:** _____

Add additional lines below if needed; or, if no such resolution proof is possible, use the space below to explain why not:

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 pretend 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. **(Note: D is a successor of itself).** As usual, successors are returned in left-to-right 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.
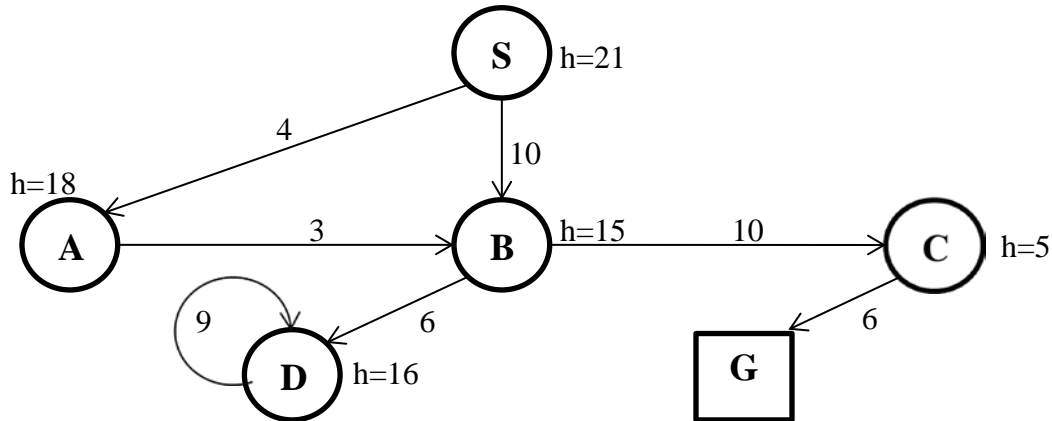
9.a. DEPTH-FIRST SEARCH:

9.a.(1) Order of expansion: S A B D D D D ...

9.a.(2) Path to goal found: None

9.b. BREADTH-FIRST SEARCH:

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

9.b.(2) Path to goal found: _____

9.c. ITERATIVE DEEPENING SEARCH:

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

9.c.(2) Path to goal found: _____

9.d. UNIFORM COST SEARCH:

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

9.d.(2) Path to goal found: _____

9.e. GREEDY BEST FIRST SEARCH:

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

9.e.(2) Path to goal found: _____

9.f. A* SEARCH:

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

9.f.(2) Path to goal found: _____

9

10. (10 pts total, 2 pts each) English to FOL Conversion. For each English sentence below, write the FOL sentence that best expresses its intended meaning. Use Dog(x) for "x is a dog," Bone(x) for "x is bone," and Likes(x, y) for "x likes y."
    The first one is done for you as an example.

10a. (2 pts) "Every dog likes every bone."

∀x ∀y [Dog(x) ∧ Bone(y)] ⇒ Likes(x, y)

10b. (2 pts) "Some dog likes some bone."

10c. (2 pts) "For every dog, there is a bone that the dog likes."

10d. (2 pts) "For every bone, there is a dog who likes that bone."

10e. (2 pts) "There is a bone that every dog likes."

10f. (2 pts) "There is a dog who likes every bone."