

CS-171, Intro to A.I. — Final Exam — Summer Session 1, 2018

YOUR NAME: _____

YOUR ID: _____ ID TO RIGHT: _____ ROW: _____ SEAT: _____

Please turn off all cell phones now.

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

When told to begin, check first to ensure that your copy has all the pages, as numbered 1-18 in the bottom-right corner of each page. We will supply a new exam for any copy problems.

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

Clear your desk except for pen, pencil, eraser, & water bottle. Put backpacks under your seat.
Please do not detach the provided scratch paper from the exam.

After you first stand up from your seat, your exam is over and must be turned in immediately.
You may turn in your exam early and leave class when you are finished. IDs will be checked.

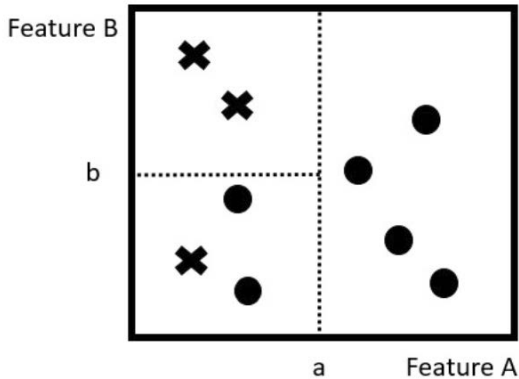
This page summarizes the points for each question, so you can plan your time.

1. (10 pts total) DECISION TREE LEARNING.
2. (10 pts total, 2 pts each) SEARCH STRATEGIES.
3. (10 pts total, -1 pt for each error, but not negative) MINI-MAX SEARCH/ALPHA-BETA PRUNING.
4. (5 pts total, 1 pt each) UNIFIERS AND UNIFICATION.
5. (5 points if correct, else partial credit up to 3 pts as 1 pt for each correct non-trivial derivation)
PROPOSITIONAL LOGIC CONVERSION TO CNF.
6. (5 points if correct, else partial credit up to 3 pts as 1 pt for each correct non-trivial derivation)
FIRST ORDER PREDICATE LOGIC CONVERSION TO CNF.
7. (10 pts total) PROBLEM SOLVING BY SEARCH/MINIATURE TOWER OF HANOI.
8. (10 pts total) BAYESIAN NETWORKS.
9. (7 pts total, 1 pt each) PROBABILITY FORMULAS.
10. (10 points total, 2 pts each) CONSTRAINT SATISFACTION PROBLEMS.
11. (3 pts if correct, else partial credit up to 2 pts as 1 pt for each correct non-trivial resolution)
FOPL: PROVE THAT TWEETY IS A BIRD.
12. (10 pts total, 1 pt each) MACHINE LEARNING CONCEPTS.
13. (5 pts if correct, else partial credit up to 3 pts as 1 pt for each correct non-trivial resolution)
RESOLUTION PROOF THAT CHARLIE DID IT.

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

1. (10 pts total) DECISION TREE LEARNING.


1.a. (4 pts) Draw a decision tree from the decision boundaries shown at the left hand side. Label each branch node with the feature selected for the decision at that node. Label each arc with either “Feature X < x” or “x ≤ Feature X” as appropriate. Label each leaf node with either a circle or an X.



1.b. (3 pts) This problem is unrelated to 1.a. above. You are constructing a decision tree and have reached a branch node at which there are four circles and four Xs. You are considering which feature to select as the next root of your decision tree sub-tree. You have three features available as choices: Feature1, Feature2, and Feature3. Each of these three features separates the initial four circles and four Xs as shown below. You will be asked below to order these three features in the order of largest information gain, i.e., in the order of their desirability to be the next root of your decision sub-tree at this branch node. (You should not need calculation to answer this question; the answer should be visually apparent based upon the material in your textbook and the lecture slides. However, for your convenience: $\log_2 1 = 0$; $\log_2 2 = 1$; $\log_2 3 \approx 1.7$; $\log_2 4 = 2$; $\log_2 5 \approx 2.3$.)

Fill in the blanks with Feature1, Feature2, and Feature3 in the order of largest information gain, i.e., in the order of their desirability to be the next root of your decision sub-tree at this branch node.

_____ > _____ > _____



1.c. (3 pts, 1 pt each) Decision Tree concepts. Fill in the blanks below with T (= True) or F (= False).

_____ Decision tree classifiers can avoid overfitting if a sufficiently large training data set is provided.

_____ The prediction error on the test data set of a decision tree classifier is always less than a perceptron classifier if both are provided with the same training data set.

_____ A decision tree can learn and represent any Boolean function.


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

2. (10 pts total, 2 pts each) SEARCH STRATEGIES. Execute Tree Search through this graph (i.e., do not remember visited nodes, so repeated nodes are possible). It is not a tree, but pretend that you don't know that it is not a tree. S is the start node and G1, G2 are the two goal nodes.

Step costs are given next to each arc. 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. **Successors are returned in left-to-right order. Successors of S are A, B, C; successors of A are A, B; successors of B are G1, G2; successors of C are B, C; in those orders.**

For each search strategy below, show the order in which nodes are expanded (i.e., to expand a node means that its children are generated) . If stuck in a loop, indicate the looping sequence of states. Show the path from start to goal, or write "None". Your "Path found" will be scored based on your "Order of node expansion" even if that answer was wrong.

The first one is done for you as an example.



2.a. (example) DEPTH FIRST SEARCH.

Order of node expansion: S A A A ...

Path found: None

2.b. (2 pts total) BREADTH FIRST SEARCH.

(1 pt) Order of node expansion: _____

(1 pt) Path found: _____

2.c. (2 pts total) UNIFORM COST SEARCH.

(1 pt) Order of node expansion: _____

(1 pt) Path found: _____

2.d. (2 pts total) GREEDY (BEST-FIRST) SEARCH.

(1 pt) Order of node expansion: _____

(1 pt) Path found: _____

2.e. (2 pts total) ITERATED DEEPENING SEARCH.

(1 pt) Order of node expansion: _____

(1 pt) Path found: _____

2.f. (2 pts total) A* SEARCH.

(1 pt) Order of node expansion: _____

(1 pt) Path found: _____

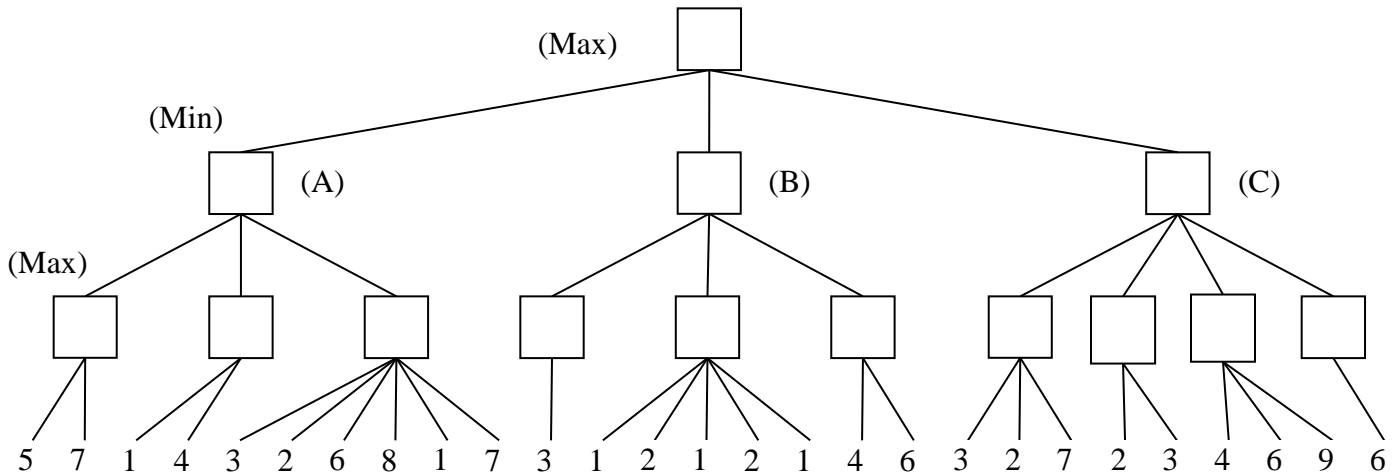
3. (10 pts total, -1 pt for each error, but not negative) MINI-MAX SEARCH/ALPHA-BETA PRUNING.
 The game tree below illustrates a position reached in the game. Process the tree left-to-right. It is **Max**'s turn to move. At each leaf node (number the at bottom) is the estimated score returned by the heuristic static evaluator.

3.a. Fill in each blank square with the proper mini-max search value.

3.b. What is the best move for Max? (write A, B, or C) _____

3.c. What score does Max expect to achieve? _____

3.d. Draw X over each leaf node (number at the bottom) that will be pruned by Alpha-Beta Pruning.



4. (5 pts total, 1 pt each) Unifiers and Unification.

Write the **most general unifier** (or MGU) of the two terms given, or "None" if no unification is possible. Write your answer in the form of a substitution as given in your book, e.g., the substitution $\{x / John, y / Mary, z / Bill\}$ means substitute x by $John$, substitute y by $Mary$, and substitute z by $Bill$.

The first one is done for you as an example.

4.a. (example) UNIFY(*Knows*(*John*, x), *Knows*(*John*, *Jane*)) $\{x / Jane\}$

4.b. (1 pt) UNIFY(*Knows*(*John*, x), *Knows*(y , *Jane*)) _____

4.c. (1 pt) UNIFY(*Knows*(*John*, x), *Knows*(y , *Father*(y))) _____

4.d. (1 pt) UNIFY(*Knows*(*John*, $F(x)$), *Knows*(y , $F(F(z))$)) _____

4.e. (1 pt) UNIFY(*Knows*(*John*, $F(x)$), *Knows*(y , $G(z)$)) _____

4.f. (1 pt) UNIFY(*Knows*(*John*, $F(x)$), *Knows*(y , $F(G(y))$)) _____

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

Problem 5 was canceled. Everyone gets it right, regardless of what you answered.

5. **Propositional Logic Conversion to CNF.** Convert this Propositional Logic wff (well-formed formula) to Conjunctive Normal Form and simplify. Show your work (correct result without work = 0 pts). **(4 points)**

$$(P \Leftrightarrow Q) \Rightarrow (\neg Q \wedge R)$$

(one valid solution; if you failed to produce the correct answer, partial credit will be awarded up to 3 points, at 1 point per useful conversion.)

Problem 6 was canceled. Everyone gets it right, regardless of what you answered.

6. **Order Predicate Logic Conversion to CNF.** Convert this FOPL wff (well-formed formula) to Conjunctive Normal Form and simplify. Show your work (correct result without work = 0 pts). **(4 points)**

“Everyone who loves all animals is loved by someone.”


$$\forall x [\forall y \text{ Animal}(y) \Rightarrow \text{Loves}(x,y)] \Rightarrow [\exists y \text{ Loves}(y,x)]$$

(one valid solution; if you failed to produce the correct answer, partial credit will be awarded up to 3 points, at 1 point per useful conversion.)

7. (10 pts total) PROBLEM SOLVING BY SEARCH/MINIATURE TOWER OF HANOI.

You are a robot tasked with solving the **MINIATURE TOWER OF HANOI** problem. There are 3 towers (A, B, C), and 2 disks (small one and large one). The goal is to move both disks from tower A to tower C (as illustrated in the figure below), subject to following three conditions:

- You can move only one disk at a time.
- You can move only the top disk in a stack of disks.
- You cannot put the large disk on top of the small disk.



The possible states can be denoted as follows:

$x:(y, z)$ where x is the state number (below), y is the tower letter for the large disk, and z is the tower letter for the small disk. E.g., 3:(A C) says that in state 3 the large disk is on tower A and the small disk is on tower C.

If we use this notation, the following nine states are possible.

1:(A A), 2:(A B), 3:(A C), 4:(B A), 5:(B B), 6:(B C), 7:(C A), 8:(C B), 9:(C C)

7.a. (2 pts total, 1 pt each) Which states are the initial state and goal state?

7.a.1. (1 pt) Initial state: _____

7.a.2. (1pt) Goal state: _____

7.b. (4 pts total, -1 for each error, but not negative) Enumerate all one-move transitions between states.

For example, 1->2, 3 means that it is possible to make a transition from state 1 to state 2 or state 3 in one move. Include “back transitions;” e.g., since you are given the example 1->2, 3, necessarily states 2 and 3 must have a transition back to state 1.

The first one is done for you as an example.

1-> 2, 3

2-> _____

3-> _____

4-> _____

5-> _____

6-> _____

7-> _____

8-> _____

9-> _____

[PROBLEM CONTINUES ON THE NEXT PAGE]

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

7.c. (4 pts total, -1 for each error, but not negative) PROBLEM SOLVING BY SEARCH (continued).


Find a solution using breath-first search and draw the search tree. Stop when you reach goal state 9.

Use Tree Search, i.e., do not remember visited nodes. Children of a node are returned in increasing numerical order by state; i.e., the children of a node are ordered left-to-right by increasing state number.

Assume that cycles are detected and eliminated by never expanding a node containing a state that is repeated on the path back to the root. That is, if a newly-generated child node state also occurs on the path from the parent back to the root, then that child is pruned immediately. You do not need to write down nodes that are pruned in this way; you may assume that they are simply discarded and the search back-tracks immediately.


This pruning strategy is a “light-weight” way to avoid many cycles and repeated states in Tree Search without the full memory burden of remembering all visited nodes. However, it is an approximate strategy and does not fully solve the repeated node problem. (See the lecture notes on this topic.)

The first expansion is done for you as an example.




8. (10 pts total) BAYESIAN NETWORKS.

8.a. (3 pts total, -1 for each error, but not negative) Write down the factored conditional probability expression corresponding to this Bayesian Network:

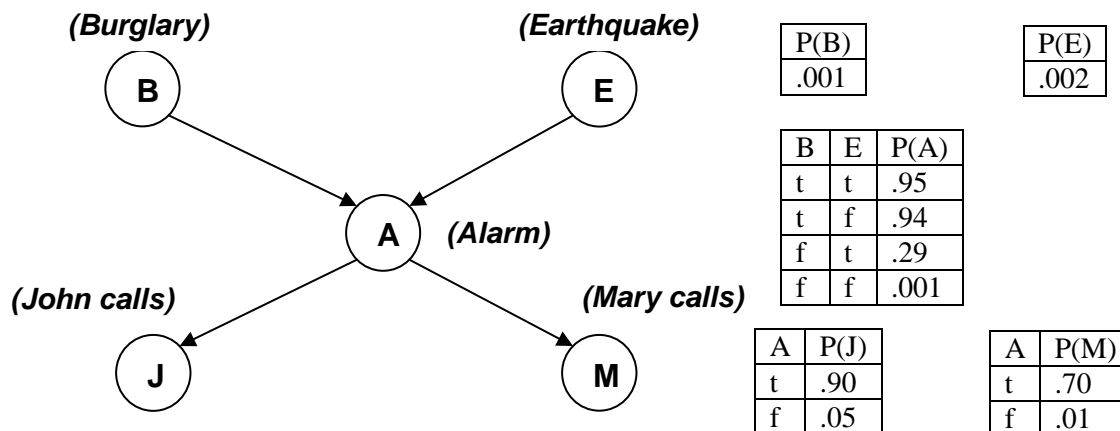


8.b. (3 pts, -1 for each error, but not negative) Draw the Bayesian Network corresponding to this factored conditional probability expression:

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



8.c. (4 pts, -1 for each error, but not negative) Shown below is the Bayesian network corresponding to the Burglar Alarm problem, i.e., $P(J, M, A, B, E) = P(J | A) P(M | A) P(A | B, E) P(B) P(E)$. This is Fig. 14.2 in your R&N textbook. The probability tables show the probability that a boolean random variable is true. For example, $P(J=t | A=t)$ is 0.9 and so $P(J=f | A=t)$ is $1 - 0.9 = 0.1$.



Write down an expression that will evaluate to $P(J=f \wedge M=t \wedge A=t \wedge B=t \wedge E=f)$. 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).

$$P(J=f \wedge M=t \wedge A=t \wedge B=t \wedge E=f) =$$

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

9. (7 pts total, 1 pt each) PROBABILITY FORMULAS. Write out the following probability formulas. Below, “in terms of X” means X should appear in your answer. All answers should be formulas, not text.

9.a. (1 pt) Write the formula for $P(A \wedge B)$ in terms of $P(A \vee B)$ and possibly other terms.

$$P(A \wedge B) =$$

9.b. (1 pt) Write the formula for the conditional probability $P(A | B)$.

$$P(A | B) =$$

9.c. (1 pt) Factor $P(A \wedge B \wedge C)$ completely using the Product Rule (or Chain Rule). You may use any variable ordering you wish.

$$P(A \wedge B \wedge C) =$$

9.d. (1 pt) Given a joint probability distribution $P(A \wedge B \wedge C)$, use the Sum Rule (or Law of Total Probability) to write the marginal probability of $P(A)$.

$$P(A) =$$

9.e. (1 pt) Write Bayes' Rule (or Bayes' Theorem).

$$P(A | B) =$$


9.f. (1 pt) Assume that A and B are independent. Write $P(A \wedge B)$ in terms of $P(A)$ and $P(B)$ and possibly other terms.

$$P(A \wedge B) =$$

9.g. (1 pt) Assume that A and B are conditionally independent given C. Write $P(A \wedge B | C)$ in terms of $P(A | C)$ and $P(B | C)$ and possibly other terms.

$$P(A \wedge B | C) =$$

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



AL = Algeria
CH = Chad
LI = Libya
ML = Mali
MR = Mauretania
NI = Niger
SE = Senegal
SS = Spanish Sahara

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

10.a. (2 pts total, -1 each wrong answer, but not negative) FORWARD CHECKING. NI has been assigned the value B, as shown; but no constraint propagation has been done. Cross out all values that would be eliminated by Forward Checking (FC):

AL	CH	LI	ML	MR	NI	SE	SS
R G B	R G B	R G B	R G B	R G B	B	R G B	R G B

10.b. (2 pts total, -1 each wrong answer, but not negative) ARC CONSISTENCY.

NI has been assigned B and AL has been assigned R, as shown; but no constraint propagation has been done. Cross out all values that would be eliminated by Arc Consistency (AC-3 in your book).

AL	CH	LI	ML	MR	NI	SE	SS
R	R G B	R G B	R G B	R G B	B	R G B	R G B

10.c. (2 pts total, -1 each wrong answer, but not negative) MINIMUM-REMAINING-VALUES (MRV) HEURISTIC. Consider the assignment below. NI has been assigned B and constraint propagation has been done, as shown. List all unassigned variables (in any order) that might be selected now by the Minimum-Remaining-Values (MRV) Heuristic: _____.

AL	CH	LI	ML	MR	NI	SE	SS
R G	R G	R G	R G	R G B	B	R G B	R G B

10.d. (2 pts total, -1 each wrong answer, but not negative) DEGREE HEURISTIC (DH). Consider the assignment below. (It is the same assignment as in problem 10.c above.) NI has been assigned B and constraint propagation has been done, as shown. Ignoring the MRV heuristic, list all unassigned variables (in any order) that might be selected now by the Degree Heuristic (DH): _____.

AL	CH	LI	ML	MR	NI	SE	SS
R G	R G	R G	R G	R G B	B	R G B	R G B

10.e. (2 pts total, -1 each wrong answer, but not negative) LEAST-CONSTRAINING-VALUE (LCV) HEURISTIC. Consider the assignment below. (It is the same assignment as in problem 10.c above.) NI has been assigned B and constraint propagation has been done, as shown. MR has been chosen as the next variable to explore. List all values for MR that could be explored first by the Least-Constraining-Value Heuristic (LCV). _____.

AL	CH	LI	ML	MR	NI	SE	SS
R G	R G	R G	R G	R G B	B	R G B	R G B

11. (3 pts if correct, else partial credit up to 2 pts as 1 pt for each correct non-trivial resolution)

FOPL: Prove that Tweety is a bird. (Adapted from

[//www.cs.cornell.edu/courses/cs472/2007fa/lectures/17-kb-systems_fol.pdf](http://www.cs.cornell.edu/courses/cs472/2007fa/lectures/17-kb-systems_fol.pdf))

You are a logic robot given the Knowledge Base below. Produce a FOPL resolution proof that Tweety is a bird.

Category	FOPL wff	CNF clausal form
General Knowledge about the domain (Laws of Physics)	$\forall x \text{ Feathers}(x) \Rightarrow \text{Bird}(x)$	$\neg \text{Feathers}(x) \vee \text{Bird}(x)$
Description of a specific problem instance	$\text{Feathers}(\text{Tweety})$	$\text{Feathers}(\text{Tweety})$
Goal sentence	$\text{Bird}(\text{Tweety})$	$\text{Bird}(\text{Tweety})$
Negated Goal sentence	$\neg \text{Bird}(\text{Tweety})$	$\neg \text{Bird}(\text{Tweety})$

The KB sentences have been labeled S_i in order to make them easy to refer to in your proof. Your resulting KB plus the negated goal sentence is as shown:

$S_1: \neg \text{Feathers}(x) \vee \text{Bird}(x)$

$S_2: \text{Feathers}(\text{Tweety})$

$S_3: \neg \text{Bird}(\text{Tweety})$

Produce a FOPL resolution proof that Tweety is a bird. On each line below, enter the two clauses in your KB that you wish to resolve (denoted as S_i for ease in identification), the Most General Unifier (MGU) of those two clauses, and the resolvent clause that results from the unification and resolution. Insert the resolvent clause that results back into your KB, and repeat. The shortest proof that I know of is only two lines long. Longer proofs are OK as long as they are mathematically correct. Add extra lines if needed.

Below, fill in the blanks of “Resolve _____ with _____” with S_i to designate the sentences to be resolved.

Fill in the blanks of “using MGU _____” with the substitution that makes the two resolved literals syntactically identical. Recall that an MGU substitution is written as $\{x/y, z/A\}$, which means that y is substituted for (= replaces) x and A is substituted for (= replaces) z . Write “None” if no substitution is needed.

End with “{ }” to indicate that you have produced the empty clause (a contradiction, proving the goal sentence).

Resolve _____ with _____ using MGU _____ to yield $S_4 =$ _____

Resolve _____ with _____ using MGU _____ to yield $S_5 =$ _____

Resolve _____ with _____ using MGU _____ to yield $S_6 =$ _____

Resolve _____ with _____ using MGU _____ to yield $S_7 =$ _____

Resolve _____ with _____ using MGU _____ to yield $S_8 =$ _____

Resolve _____ with _____ using MGU _____ to yield $S_9 =$ _____

Resolve _____ with _____ using MGU _____ to yield $S_{10} =$ _____

12. (10 pts total, 1 pt each) Machine Learning concepts.

For each of the following items on the left, write in the letter corresponding to the best answer or the correct definition on the right. The first one is done for you as an example.

A .	Learning	A	Improves performance of future tasks after observing the world
	Information Gain	B	Fixed set, list, or vector of features/attributes paired with a value
	Decision Boundary	C	Tests $P(C) \prod_i P(X_i C)$, where C is a class label and X_i are features
	Cross-validation	D	Tests $\mathbf{w} \cdot \mathbf{f} > 0$, where \mathbf{w} is a weight vector and \mathbf{f} is a feature vector
	Linear Classifier (Perceptron)	E	Example input-output pairs, from which to discover a hypothesis
	Factored Representation (Feature Vector)	F	Examples distinct from training set, used to estimate accuracy
	Overfitting	G	Randomly split the data into a training set and a test set
	Test Set	H	Surface in a high-dimensional space that separates the classes
	Naïve Bayes Classifier	I	Internal nodes test a value of an attribute, leaf nodes=class labels
	Training Set	J	Expected reduction in entropy from testing an attribute value
	Decision Tree	K	Choose an over-complex model based on irrelevant data patterns

13. (5 pts if correct, else partial credit up to 3 pts as 1 pt for each correct non-trivial resolution)

RESOLUTION PROOF THAT CHARLIE DID IT. (Adapted from <http://www.braingle.com>)

Detective Dorothy interviewed four local burglars to identify who stole Lady Diva's teapot.

It was well known that each burglar told exactly one lie:

Arnold: I didn't do it. Brian did it.

Brian: I didn't do it. Derek did it.

Charlie: I didn't do it. Brian is lying when he says Derek did it.

Derek: I didn't do it. If Arnold didn't do it, then Brian did it.

Use these propositional variables:

A = Arnold did it. B = Brian did it. C = Charlie did it. D = Derek did it.

You translate the evidence into propositional logic (recall that each suspect told exactly one lie):

Arnold: $(A \wedge B) \vee (\neg A \wedge \neg B)$

Brian: $(B \wedge D) \vee (\neg B \wedge \neg D)$

Charlie: $(C \wedge \neg D) \vee (\neg C \wedge D)$

Derek: $(D \wedge (\neg A \Rightarrow B)) \vee (\neg D \wedge \neg(\neg A \Rightarrow B))$

At most one burglar stole the teapot:

$(A \Rightarrow \neg B \wedge \neg C \wedge \neg D) \wedge (B \Rightarrow \neg A \wedge \neg C \wedge \neg D) \wedge (C \Rightarrow \neg A \wedge \neg B \wedge \neg D) \wedge (D \Rightarrow \neg A \wedge \neg B \wedge \neg C)$

After converting to Conjunctive Normal Form, your Knowledge Base (KB) consists of:

Arnold: $(A \vee \neg B) \quad (\neg A \vee B)$ **Brian:** $(B \vee \neg D) \quad (\neg B \vee D)$

Charlie: $(C \vee D) \quad (\neg C \vee \neg D)$ **Derek:** $(\neg A \vee D) \quad (\neg B \vee D) \quad (A \vee B \vee \neg D)$

At most one: $(\neg A \vee \neg B) \quad (\neg A \vee \neg C) \quad (\neg A \vee \neg D) \quad (\neg B \vee \neg C) \quad (\neg B \vee \neg D) \quad (\neg C \vee \neg D)$

(Side note: Normally, you would start four proofs, one for each goal sentence: A, B, C, D. Only the proof of C would succeed, and you would know Charlie did it. For this timed test, you will do only one proof.)

You will be asked to prove, "Charlie did it." The goal is (C). You adjoin the negated goal to your KB:

$(\neg C)$

Produce a resolution proof, using KB and the negated goal, that "Charlie did it."

Repeatedly choose two clauses, write one clause in the first blank space on a line, and the other clause in the second. Apply resolution to them. Write the resulting clause in the third blank space, and insert it into the knowledge base. Continue until you produce (). If you cannot produce (), then you have made a mistake. The shortest proof I know is only three lines. It is OK to use more lines, if your proof is correct. It is OK to use abbreviated CNF, i.e., $(\neg A \neg B)$ instead of $(\neg A \vee \neg B)$. It is OK to omit the parentheses.

Resolve _____ and _____ to give _____.

Resolve _____ and _____ to give _____.

Resolve _____ and _____ to give _____.

Resolve _____ and _____ to give _____.

Resolve _____ and _____ to give _____.

Resolve _____ and _____ to give _____.

Resolve _____ and _____ to give _____.

Resolve _____ and _____ to give _____.

**** THIS IS THE END OF THE FINAL EXAM ****

Scratch Paper (1) Please Do Not Detach From Test

Scratch Paper (2) Please Do Not Detach From Test

Scratch Paper (3) Please Do Not Detach From Test

Scratch Paper (4) Please Do Not Detach From Test