


## Final Review

1. Consider the following game tree.



- (a) Find the best move for the MAX player using the minimax procedure.
- (b) Perform a left-to-right alpha-beta pruning on the tree. Indicate where the cutoffs occur.
- (c) Perform a right-to-left alpha-beta pruning on the tree. Discuss why different pruning occurs.
2. For each of the following pairs of atomic sentences, give the most general unifier, if it exists.
- $P(A,B,B), P(x,y,z)$
  - $Q(y,G(A,B)), Q(G(x,x),y)$
  - $Older(Father(y),y), Older(Father(x),John)$
  - $Knows(Father(y),y), Knows(x,x)$
  - $Ancestor(x,y), Ancestor(John,Father(John))$
3. The function  $cons(x,y)$  denotes list formed by inserting the element  $x$  at the head of the list  $y$ . We denote the empty list by  $Nil$ ; the list (2) by  $cons(2, Nil)$ ; the list (1,2) by  $cons(1, cons(2, Nil))$ ; and so on. The formula  $Member(e,l)$  is intended to mean that  $e$  is a member of the list  $l$ . We have the following axioms:
- $(\forall x, y)[Member(x, cons(x, y))]$
  - $(\forall x, y, z)[Member(x, y) \supset Member(x, cons(z, y))]$

Prove  $Member(b, cons(a, cons(b, nil)))$  from these axioms by the method of resolution refutation.

4. Dr. Somebody, Dr. Anybody, and Dr. Nobody are computer scientists. We know the following facts about them:
  1. Dr. Somebody is an associate professor.
  2. Dr. Nobody is an assistant professor and has published papers with Dr. Anybody.
  3. Dr. Anybody is either an associate or an assistant professor (but not both) and has published papers with Dr. Somebody.

Use resolution refutation to prove that an assistant professor has published papers with an associate professor; that is, prove  $(\exists x, y)[Assistant(x) \wedge Associate(y) \wedge PPW(x, y)]$ .

5. We are given the following paragraph:

Tony, Mike, and John belong to the Alpine Club. Every member of the Alpine Club is either a skier or a mountain climber or both. No mountain climber likes rain, and all skiers like snow. Mike dislikes whatever Tony likes and likes whatever Tony dislikes. Tony dislikes rain and snow.

Represent this information by predicate-calculus sentences in such a way that you can represent the question “Is there a member of the alpine club who is a skier but not a mountain climber?” as a predicate-calculus expression. Use resolution refutation to answer it.

6. A simple version of the *nim* game is played as follows: Two players alternate in removing stones from three piles initially containing one, two, and three stones, respectively. The player who picks up the last stone wins. Each player can pick one or more stones from a single pile; at least one stone has to be picked every time. At every turn the players can pick from different piles. Show, by drawing the game tree, which player can always win. Is it necessary to generate the whole tree to find a winning strategy?
7. Consider the following statements:
  - (1) Whoever can read is literate  $(\forall x[R(x) \Rightarrow L(x)])$ .
  - (2) Horses are not literate  $(\forall x[H(x) \Rightarrow \neg L(x)])$ .
  - (3) Some horses are intelligent  $(\exists x[H(x) \wedge I(x)])$ .

Use these statements to prove the following statement by resolution:

- (4) Some who are intelligent cannot read  $(\exists x[I(x) \wedge \neg R(x)])$ .

8. You are given the following table:

Outlook	Temperature	Humidity (%)	Windy?	Class
sunny	warm	< 75	yes	Play
sunny	hot	> 75	yes	Don't Play
sunny	hot	> 75	no	Don't Play
sunny	warm	> 75	no	Don't Play
sunny	warm	< 75	no	Play
overcast	warm	> 75	yes	Play
overcast	hot	> 75	no	Play
overcast	cool	< 75	yes	Play
overcast	hot	< 75	no	Play
rain	warm	> 75	yes	Don't Play
rain	cool	< 75	yes	Don't Play
rain	warm	< 75	no	Play
rain	cool	> 75	no	Play

- (a) Build a decision tree to distinguish the two classes. Use the information theory to decide which attribute should be used at the root of the decision tree (10p for this step).
- (b) Indicate the 'Play' concept represented by the tree.
9. We represent the statement that everything is representable in the predicate calculus as  $\forall x, \text{represents}(pc, x)$ . General Problem Solver (GPS) is a system for automated problem solving. We represent the statement that all problems representable in predicate calculus are solvable using GPS as  $\forall x, (\text{problem}(x) \wedge \text{represents}(pc, x)) \Rightarrow \text{solves}(gps, x)$ . Now using these two statements and the fact that the Traveling Salesperson Problem (TSP) is a problem ( $\text{problem}(tsp)$ ), prove that GPS solves it.
10. Attempt to unify the following pairs of expressions. Either show their most general unifiers or explain why they will not unify.
- (a)  $p(X, Y)$  and  $p(a, Z)$
- (b)  $p(X, X)$  and  $p(a, b)$
- (c)  $\text{ancestor}(X, Y)$  and  $\text{ancestor}(\text{bill}, \text{father}(\text{bill}))$
- (d)  $\text{ancestor}(X, \text{father}(X))$  and  $\text{ancestor}(\text{david}, \text{george})$
11. Consider the sentence "Heads I win; tail you lose." We can represent this sentence plus associated domain knowledge in the propositional logic using the following proper axioms, where *Heads*, *Tails*, *WinMe*, and *LoseYou* are propositional variables:

$$\begin{aligned}
 \text{Heads} &\Rightarrow \text{WinMe} \\
 \text{Tails} &\Rightarrow \text{LoseYou} \\
 \neg \text{Heads} &\Rightarrow \text{Tails} \\
 \text{LoseYou} &\Rightarrow \text{WinMe}
 \end{aligned}$$

- a) Determine if it is possible to prove *WinMe* using the rule of inference modus ponens and these four axioms.
- b) Convert each of the four axioms to a disjunction of literals.
- c) For each of the resulting disjunctions, specify if it is a Horn clause.

- d) Determine if it is possible to prove WinMe using the resolution rule of inference and the four axioms written as disjunctions of literals.
12. We have defined four different binary logical connectives ( $\wedge$ ,  $\vee$ ,  $\Rightarrow$ ,  $\Leftrightarrow$ ).
- Are there any others that might be useful?
  - How many binary connectives can there be?
  - Why are some of them not very useful?
13. Consider the following problem: The rules in the National Zoo forbid visitors to feed animals. A prairie dog, an animal, has some candy, and all of its candy was given to it by BB, who is a Zoo visitor.
- Represent these facts as first-order logic clauses.
  - Use backward-chaining to prove that BB broke Zoo rules.
  - Use forward-chaining to prove that BB broke Zoo rules.
14. Build a decision tree to distinguish two classes ('reads' and 'skips') from the following table:

Example	Action	Author	Thread	Length
e1	skips	known	new	long
e2	skips	known	follow_up	long
e3	skips	unknown	follow_up	long
e4	reads	known	follow_up	short
e5	reads	unknown	new	short
e6	reads	known	new	short
e7	skips	unknown	new	long
e8	skips	unknown	follow_up	short

Use the information theory to decide which attribute to use at each non-terminal node of the tree.

15. Sam, Clyde, and Oscar are rabbits. We know the following facts about them:
- Sam is pink.
  - Clyde is gray and likes Oscar.
  - Oscar is either pink or gray (but not both) and likes Sam.

Use resolution refutation to prove that a gray rabbit likes a pink rabbit; that is, prove  $(\exists x, y)[Gray(x) \wedge Pink(y) \wedge Likes(x, y)]$ .

16. The function  $cons(x, y)$  denotes list formed by inserting the element  $x$  at the head of the list  $y$ . We denote the empty list by  $nil$ ; the list (2) by  $cons(2, nil)$ ; the list (1, 2) by  $cons(1, cons(2, nil))$ ; and so on. The formula  $Last(l, e)$  is intended to mean that  $e$  is the last element of the list  $l$ . We have the following axioms:

- $(\forall u)[Last(cons(u, nil), u)]$
- $(\forall x, y, z)[Last(y, z) \supset Last(cons(x, y), z)]$

1. Prove the following theorem from these axioms by the method of resolution refutation:

$$(\exists v)[Last(cons(2, cons(1, nil)), v)]$$

2. Use answer extraction to find  $v$ , the last element of the list  $(2, 1)$ .
17. The set of inputs  $x_1, x_2$  define a 2-dimensional space. A boolean function defining a concept is said to be *linearly separable* if there exists a hyperplane (line) dividing the space into inputs for which the function produces a 1 and inputs for which the function produces a 0.

- a) Which of these four boolean functions are linearly separable?

*or and xor equal*

Justify your answer.

- b) What needs to be done to make concepts which are not linearly separable learnable by the perceptron learning rule. Illustrate your answer using a boolean function that is not linearly separable.
18. A simple version of a “Kayles game” is played as follows: Two players have in front of them a single contiguous sequence of objects, say 5 pennies, which are placed next to each other so that the first penny touches the second, the second touches the third, the third touches the fourth, and the fourth touches the fifth penny. The first player removes 1 or 2 pennies whose sides that are touching (for example, the first player can remove any single penny, or pennies 1 and 2, or pennies 2 and 3, or pennies 3 and 4, or pennies 4 and 5). Each player alternatively thereafter removes 1 penny or 2 pennies whose sides are touching. (Note: If the first player removes penny 2, the pennies 1 and 3 will have a gap between them and cannot be removed at the same time.) The last player to pick a penny (or two) loses. Show, by drawing a game tree, whether any of the players can always win. You can ignore symmetries and obvious losing plays if there is a winning play available, such as picking both pennies when there are 2 pennies left.
19. Consider a vocabulary with only four propositions,  $A$ ,  $B$ ,  $C$ , and  $D$ . How many models are there for the following sentences?

- a.  $(A \wedge B) \vee (B \wedge C)$

- b.  $A \vee B$

- c.  $A \Leftrightarrow B \Leftrightarrow C$

Hint: In each sentence all four propositions have to be considered.

20. Build a decision tree to distinguish two classes ('+' and '-') from the following table:

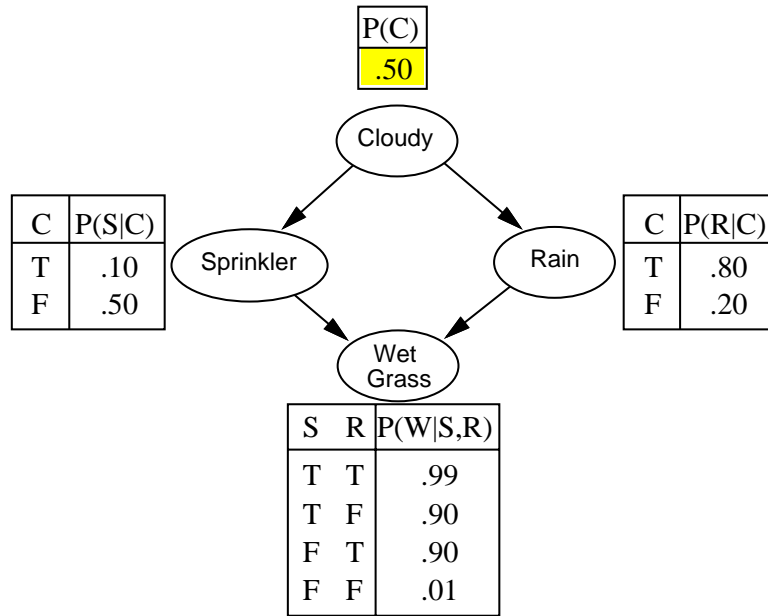
food	medium	type	class		description
herbivore	land	harmless	mammal	+	deer (e1)
carnivore	land	harmful	mammal	-	lion (c1)
omnivorous	water	harmless	fish	+	goldfish (e2)
herbivore	amphibious	harmless	amphibian	-	frog (c2)
omnivorous	air	harmless	bird	-	parrot (c3)
carnivore	land	harmful	reptile	+	cobra (e3)
carnivore	land	harmless	reptile	-	lizard (c4)
omnivorous	land	moody	mammal	+	bear (e4)

Use information theory to decide which attribute should be used at each nonleaf node of your decision tree.

21. A simple version of the *nim* game is played as follows: Two players alternate in removing stones from two piles initially containing several stones each. The player who picks up the last stone wins. At any given turn a player can pick one or more stones from a single pile; at least one stone has to be picked every time.
  - a) Show, by drawing a game tree, which player can always win if the piles have two and two stones.
  - b) Show, by drawing a game tree, which player can always win if the piles have two and three stones. You can use your result from a).
  - c) What can you say about a game in which the piles have  $m > 0$  and  $n > 0$  stones?
22. You are building a decision tree which tells us whether we should *go* to a restaurant based on various attributes: how many *patrons* it has, whether or not the food is *cheap*, and what *type* of food is served. The information is:

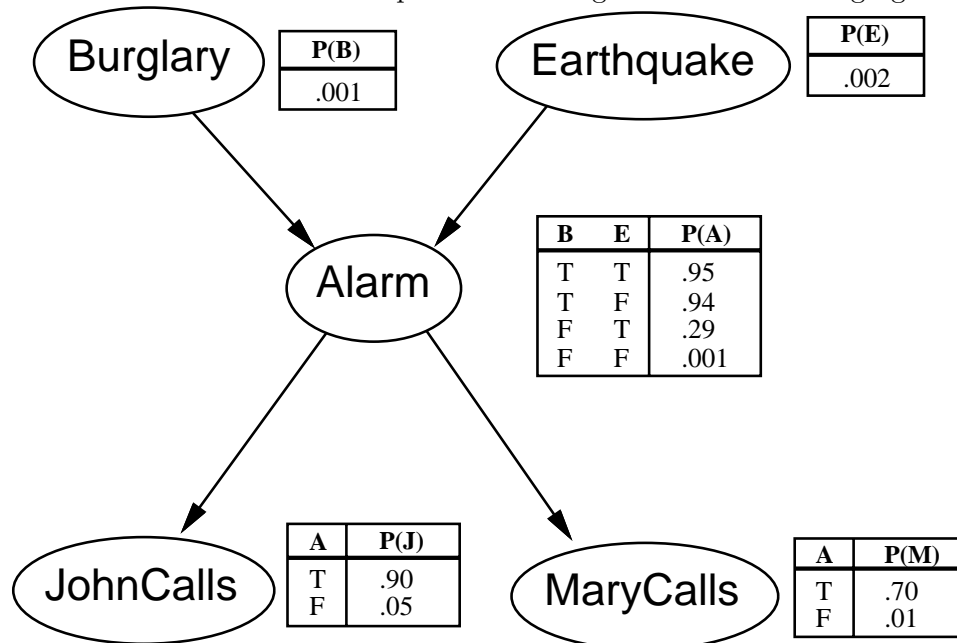
Example	<i>patrons</i>	<i>cheap</i>	<i>type</i>	<i>go</i>
1	some	yes	french	yes
2	empty	yes	thai	no
3	some	no	burger	yes
4	some	yes	thai	yes
5	full	no	french	no
6	empty	yes	italian	yes

- a) Show the tree after asking attribute questions in this order: *type?*, *cheap?*, *patrons?*
  - b) Express the concept represented by the decision tree as a first order logical expression.
  - c) What is the first attribute question you should ask and why? Show your work.
23. A simple version of the *nim* game is played as follows: Two players alternate in removing stones from three piles initially containing two stones each. The player who picks up the last stone wins. Each player can pick one or more stones from a single pile; at least one stone has to be picked every time. At any turn the player can pick from any one pile. Show, by drawing the game tree, whether a player can always win, and if so which one. Is it necessary to generate the whole tree to find a winning strategy?
  24. A typical belief network with conditional probabilities is given in the following figure:



The letters  $C, R, S$  and  $W$  stand for *Cloudy*, *Rain*, *Sprinkler*, and *Wet Grass*, respectively. All variables (nodes) are Boolean, so the probability of, say,  $\neg A$  in any row of its table is  $1 - P(A)$ .

- a) Assuming that the nodes are introduced in the following order *Wet Grass*, *Sprinkler*, *Rain* and *Cloudy* construct a corresponding belief network. Show which probabilities need to be specified.
  - b) Compute probabilities  $\mathbf{P}(W)$  and  $\mathbf{P}(S|W)$ .
25. A typical belief network with conditional probabilities is given in the following figure:



The letters  $B, E, A, J$ , and  $M$  stand for *Burglary*, *Earthquake*, *Alarm*, *JohnCalls*, and *MaryCalls*, respectively. All variables (nodes) are Boolean, so the probability of, say,  $\neg A$  in any row of its table is  $1 - P(A)$ .

- a) Assuming that the nodes are introduced in the following order *JohnCalls*, *MaryCalls*, *Alarm*, *Earthquake*, and *Burglary* construct a corresponding belief network. Show which probabilities need to be specified.
- b) Compute probabilities  $\mathbf{P}(A)$ ,  $\mathbf{P}(E|A)$ , and  $\mathbf{P}(B|A, E)$ .