

# Artificial Intelligence (CS 436)

## Homework #2—Due Monday, October 27, 2008

All work must be individual effort, and use of the web is prohibited (unless otherwise indicated). Any external sources must be cited. All papers are due at the start of class on the due date.

- [5 points] Suppose we use *alpha-beta* pruning to search a game tree and decide that some successors of node  $n$  can be pruned. Later in the traversal of the game, we encounter the same state (i.e., board position and player whose turn it is to move) at a different node  $m$ , which is at the same level as  $n$ . Can we immediately prune the same successors of  $m$  from our search tree? Explain.
- Consider the  $3 \times 3$  game in Figure 1. Assume each player is attempting to maximize their return.
  - [5 points] Identify all pairs of strategies where one strategy weakly dominates the other.
  - [5 points] Assume you are allowed to remove a weakly dominated strategy of some player. Do so, and repeat the process (of iterated elimination of weakly dominated strategies) until you find a single strategy pair of the original game. (Remember that two strategies with identical payoffs do not weakly dominate each other.)
  - [5 points] Find such an iterated elimination of weakly dominated strategies that results in a strategy pair other than the one found in (b) where *both* strategies, and the payoffs to the players, are different.
  - [5 points] What are the Nash equilibria (in pure strategies) of the game?
- Given the following predicates, translate the first-order logic sentences into English.

$Damsel(x) = x$  is a damsel  
 $HasSaved(x, y) = x$  has saved  $y$   
 $Horse(x) = x$  is a horse  
 $InDistress(x) = x$  is in distress  
 $Knight(x) = x$  is a knight  
 $Owens(x, y) = x$  owns  $y$   
 $Strong(x) = x$  is strong

Figure 1:  $3 \times 3$  Non-Zero Sum Game Payoff Matrix

		Player 2		
		A	B	C
Player 1	X	0 1	1 3	1 1
	Y	1 1	0 3	1 0
	Z	2 2	3 3	2 0

- (a) [2 points]  $\neg(\forall x \text{Damsel}(x) \Rightarrow \text{InDistress}(x))$
- (b) [2 points]  $\forall x \text{Knight}(x) \Rightarrow (\exists y \text{Horse}(y) \wedge \text{Owns}(x, y))$
- (c) [2 points]  $\exists y \text{Horse}(y) \wedge (\forall x \text{Knight}(x) \Rightarrow \text{Owns}(x, y))$
- (d) [2 points]  $\exists x \text{Knight}(x) \wedge (\exists y \text{Owns}(x, y) \wedge \text{Horse}(y) \wedge \text{Strong}(y) \wedge (\forall z \text{Owns}(x, z) \wedge \text{Horse}(z) \wedge \text{Strong}(z) \Rightarrow (z = y)))$  [Use at most 12 English words.]
- (e) [2 points]  $\neg\exists x \text{Knight}(x) \wedge (\forall y \text{Damsel}(y) \Rightarrow \text{HasSaved}(x, y))$
4. Given the same predicates as problem 5, translate the following English sentences into first-order logic.
- (a) [2 points] No damsel in distress has saved a knight.
- (b) [2 points] Some knights have been saved by their horses.
- (c) [2 points] Only strong damsels have saved knights.
- (d) [2 points] Exactly one knight has saved his own horse.
- (e) [2 points] If a damsel is in distress, then she has not been saved by a knight; otherwise, she has.
5. [10 points] The law says that it is a crime to sell an unregistered gun. Reese has several unregistered guns, and all of them were purchased from Laslo. Using the rules of inference described in class and the text book (other than resolution), derive the conclusion that Laslo is a criminal. Make sure you set this problem up as a symbolic logic problem.
6. [10 points] We can use resolution as a process for solving constraint satisfaction problems. In this case, the value assignments occur through unification as clauses are resolved. For these types of problems, the query is posed as a conjunction of positive literals containing some number of variables. In addition, the database of clauses consists entirely of positive ground literals. The task is to find bindings for the variables in the query such that each of the conjuncts is identical to a literal in the database and the conjunction is true. In other words, we are seeking a variable assignment that “satisfies” the query.

Suppose we have the following literals in a database:

$\text{Parent}(\text{Alex}, \text{Jim})$	$\text{Parent}(\text{Bill}, \text{Kelli})$	$\text{Parent}(\text{Chuck}, \text{Leah})$
$\text{Parent}(\text{Amy}, \text{Jim})$	$\text{Parent}(\text{Bea}, \text{Kelli})$	$\text{Parent}(\text{Cindy}, \text{Leah})$
$\text{CPA}(\text{Amy})$	$\text{Senator}(\text{Jim})$	
$\text{CPA}(\text{Chuck})$	$\text{Senator}(\text{Kelli})$	

In addition, suppose we have the following query:

$$(\text{Parent}(x, y) \wedge \text{CPA}(x) \wedge \text{Senator}(y))$$

In other words, we want to find the pair,  $x$  and  $y$  such that  $x$  is the parent of  $y$ ,  $x$  is a CPA, and  $y$  is a senator. Use resolution to solve this constraint satisfaction problem. (*Hint:* Include a special literal  $\text{Answer}(x, y)$  to capture the variable bindings when done.)

7. Consider the training examples shown in Table 1 for a binary classification problem.
- (a) [5 points] What is the entropy of this collection of training examples with respect to the positive class?
- (b) [5 points] What are the information gains of  $a_1$  and  $a_2$  relative to these training examples?
- (c) [10 points] For  $a_3$ , which is a continuous variable, compute the information gain for every possible split and determine where the best split for this attribute would occur.
8. [15 points] Consider the following approach for testing whether classifier A is better than another classifier B. Let  $n$  be the size of a given data set,  $p_A$  be the accuracy of classifier A,  $p_B$  be the accuracy

Table 1: Sample Data Set

Instance	$a_1$	$a_2$	$a_3$	Target Class
1	T	T	1.0	+
2	T	T	6.0	+
3	T	F	5.0	-
4	F	F	4.0	+
5	F	T	7.0	-
6	F	T	3.0	-
7	F	F	8.0	-
8	T	F	7.0	+
9	F	T	5.0	-

of classifier B, and  $p = (p_A + p_B)/2$  be the average accuracy for the two classifiers. To test whether classifier A is significantly better than classifier B, the following Z-statistic is used:

$$Z = \frac{p_A - p_B}{\sqrt{\frac{2p(1-p)}{n}}}$$

Classifier A is assumed to be better than classifier B if  $Z > 1.96$ .

Table 2 compares the accuracies of three different classifiers, denoted A, B, and C. Compare these classifiers using the Z-statistic and evaluate/discuss their relative performances.

Table 2: Comparing the Accuracy on Three Classifiers

Data Set	Size ( $n$ )	$acc_A$	$acc_B$	$acc_C$
Anneal	898	92.09	79.62	87.19
Australia	690	85.51	76.81	84.78
Auto	205	81.95	58.05	70.73
Breast	699	95.14	95.99	96.42
Cleve	303	76.24	83.50	84.49
Credit	690	85.80	77.54	85.07
Diabetes	768	72.40	75.91	76.82
German	1000	70.90	74.70	74.40
Glass	214	67.29	48.59	59.81
Heart	270	80.00	84.07	83.70
Hepatitis	155	81.94	83.23	87.10
Horse	368	85.33	78.80	82.61
Ionosphere	351	89.17	82.34	88.89
Iris	150	94.67	95.33	96.00
Labor	57	78.95	94.74	92.98
Led7	3200	73.34	73.16	73.56
Lymphography	148	77.03	83.11	86.49
Pima	768	74.35	76.04	76.59
Sonar	208	78.85	69.71	76.92
Tic-tac-toe	958	83.72	70.04	98.33
Vehicle	846	71.04	45.04	74.94
Wine	178	94.38	96.63	98.88
Zoo	101	93.07	93.07	96.04