

CS 188: Practice Midterm

Spring 2006

March 4, 2006

1 *True/False*

T/F A rational agent may choose to play a game of chance with negative expected payoff.

T/F If a tree search algorithm is complete, the corresponding graph search algorithm will be, too.

T/F For all random variables X, Y and Z and distributions over them,

$$\frac{P(Y|X)P(Z|X)P(X)}{P(Y,Z)} = P(X|Y,Z)$$

T/F The perceptron algorithm always converges for inseparable data

T/F If a CSP is known to have a tree-structured graph, then forward checking and arc consistency are equivalent operations.

2 *Search*

Consider the following search problem:

S

G

Which path will each search algorithm return, assuming all successor functions work out in such a way that nodes are explored in alphabetical order whenever possible?

(a) Breadth-first search

(b) Depth-first search

(c) Uniform-cost search

(d) A* search

(e) Greedy search

(f) Name a node that A* will expand, but uniform-cost will not.

Describe a reasonably general case in which each of the following will occur, or state that the scenario is impossible. Correct answers should not take more than one or two sentences.

(g) Depth-first search never terminates, despite a finite goal

(h) Breadth-first search never terminates, despite a finite goal

(i) Uniform cost search and breadth-first search expand the same nodes and return the same goal

3 *CSPs*

You must arrange three statues in an exhibit hall: an ice carving of a swan (i), a gold lion (g), and a marble abstract piece (m). There are three tables, 1, 2, and 3, arranged in a row, with 1 closest to the door and 3 farthest into the exhibit hall. It is a hot day and so the ice carving cannot be nearest the door. Your manager also informs you that it will look bad to have to animal sculptures on adjacent tables. Reality tells you that each table must have a different sculpture.

If we formulate this problem as a binary CSP with variables X_1 , X_2 , and X_3 , each with domain $\{i, g, m\}$:

(a) What are the unary constraint(s) (list them explicitly)

(b) What are the binary constraint(s) (list them explicitly)

Assume we enforce the unary constraint(s) in pre-processing for the remaining parts:

(c) Which variable will be assigned first by the MRV heuristic?

(d) If we assign $X_3 = i$, show the domains of the remaining variables after forward checking.

(e) If no variables are assigned, show the initial domains after running arc consistency.

(f) If it's a cool day, and we drop the requirement that the ice swan cannot be nearest the door, what are the initial domains after running arc consistency?

4 *Probability*

You are hired by a casino to help detect cheaters who manage to sneak loaded dice into a game which involves rolling a die 4 times in a row. With CS188 under your belt, you've decided to use a Naive Bayes model in order to detect cheaters. Assume that the prior probability of a player cheating is $P(\textit{cheat})$. Non-cheaters all use fair dice, and cheaters all use a single type of loaded die.

(a) The casino has kept a record of the rolls of the last three known cheaters: they were $[1, 2, 6, 6]$, $[6, 3, 5, 4]$, and $[6, 6, 6, 6]$. Using this data, estimate the distribution of the loaded die.

(b) What is the posterior probability that a gambler is cheating given a roll sequence of $[6, 2, 3, 6]$ according to your estimate of the distribution from (a)?

(c) If the utility of falsely accusing a non-cheater is -10, what is the minimum utility for catching a cheater which will make accusing the gambler in (b) the rational action?

5 *Classification*

Imagine we have features f_1, f_2, f_3, f_4 and three classes, $\{x, y, z\}$. Assume we are training a multi-class perceptron and a given point in the training, it has the following weight vectors:

$$w_x = \langle 0, 0, 0, 0 \rangle$$

$$w_y = \langle 0, 2, 0, 0 \rangle$$

$$w_z = \langle 2, 0, 1, 0 \rangle$$

(a) If we next encounter the instance $\langle 1, 0, 0, 1 \rangle$ with true label x , write the resulting weights after processing this new instance:

(b) If we next encounter the instance $\langle 0, 1, 1 \rangle$ with true label y , write the resulting weights after processing this new instance:

Consider a domain with three Boolean attributes, $\{X, Y, Z\}$ and the target function $f(x, y, z) = x \text{ xor } z$. Let H be the space of decision trees over these attributes.

(c) Is f realizable? If so, draw a decision tree which proves it; if not, argue why it is not.

Consider the following data set:

X	Y	Z	f
1	0	1	1
1	1	0	0
0	0	0	0
0	1	1	1
1	0	1	1
0	0	1	0
0	1	1	1
1	1	1	0

(d) Draw the decision tree which would be learned from this data using the recursive splitting algorithm presented in class. Assume that splits are chosen using information gain, and gain ties are broken to prefer splits by alphabetical order.