Computer Science Department
Stanford University
Comprehensive Examination in Artificial Intelligence
Autumn 1996

October 15, 1996

PLEASE READ THIS FIRST

- You should write your answers for this part of the Comprehensive Examination in a BLUE BOOK. Be sure to write your MAGIC NUMBER on the cover of every blue book that you use.
- Be sure you have all the pages of this exam. There are three pages in addition to this.
- This exam is OPEN BOOK. You may use notes, articles, or books — but no help from people or computers.
- Show your work, since PARTIAL CREDIT will be given for incomplete answers. For example, you can get credit for making a reasonable start on a problem even if the idea does not work out. You can also get credit for realizing that certain approaches are incorrect. On a true/false question, you might get partial credit for explaining why you think something is true when we think it is false.
- Points in this exam add up to 100.

1 Blind Search (20 points)

Here is the description of breadth-first search from Ginsberg's book (page 24):

1. Set L to be a list of the initial nodes in the problem.
2. Let n be the first node on L. If L is empty, fail.
3. If n is a goal node, stop and return it and the path from the initial node to n.
4. Otherwise, remove n from L and add to the end of L all of n's children, labelling each with its path from the initial node. Return to step 2.

Once this algorithm generates a goal node, the length of the path from the initial state to the goal state will be as short as possible.

This property holds under the assumption that all arcs (edges) in the tree have the same length (or cost). Modify breadth-first search to ensure that the above property holds even when arc lengths are not necessarily equal (but the lengths are known).

Just describe the changes, do not rewrite the algorithm.

2 Adversary Search (20 points)

Consider the following Minimax tree. The first player is the maximizer.

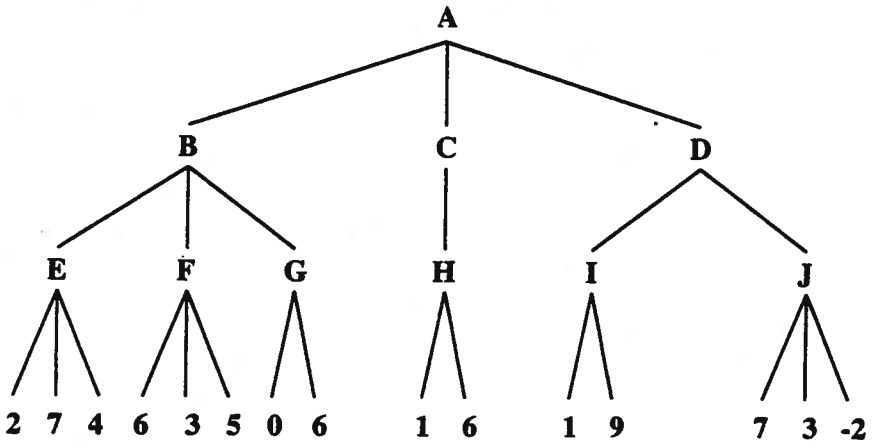

a. What is the solution? That is, which move should be made next and what is the expected value of that move?

b. Using the alpha-beta pruning method (and standard left-to-right evaluation of nodes), how many leaves of the tree get evaluated?

c. Using alpha-beta pruning, but right-to-left evaluation of nodes, how many leaves get evaluated?

(9)

3 Learning (20 points)

You are in charge of machine induction of classifiers at the Grand Kernel Bank. A request comes in from the financial officer Dr. Penny De L'Argent to induce a classifier from data on fourteen cases that concern the creditworthiness of applicants for a special Gold Credit Card, which grants a high credit limit and other privileges. The classifier should determine whether an applicant is creditworthy. Note that the Bank also grants regular cards, and some applicants for the Gold Card already hold the regular card. Past applicants who have been granted the Gold card are to be considered creditworthy.

Case	high income?	already holds regular card?	was granted gold card?
1	yes	yes	yes
2	yes	yes	yes
3	yes	yes	yes
4	yes	yes	no
5	yes	no	no
6	yes	no	yes
7	yes	no	no
8	no	no	no
9	no	no	no
10	no	no	no
11	no	no	no
12	no	no	yes
13	no	no	yes
14	no	yes	yes

What tree, including classifications, would be output by an ID3-style discrimination tree induction algorithm?

[Hints: the last column in the training table is used to determine whether a given training instance is positive or negative. This leaves two feature columns. You may want to number these 0 and 1. To answer this question, it is not necessary to compute the formula for information gain, but just to understand what it means.]

4 Probability (20 points)

Males are 47 percent of the population; 30 percent of all males smoke, and 25 percent of the entire population smokes. What fraction of all smokers are female? You may assume that people are either male or female. A numerical formula is sufficient.

(10)

5 Vision (20 points)

Ginsberg (page 329) draws a distinction between these two types of edges:

1. edges in the image, and

2. edges in the objects the image depicts.

Edges of type 1 are abrupt changes in image intensity, while edges of type 2 are discontinuities in the depth or in the tangent of a surface in the world.

Ginsberg also mentions image noise as one reason why the two types of edges may not coincide. Argue, by example or otherwise, that there are more fundamental reasons why these two types of edges do not in general coincide.

(11)