UNIVERSITY OF EDINBURGH

FACULTY OF SCIENCE

ARTIFICIAL INTELLIGENCE 4

Date: 1 June 1989 Time: 2.00 - 3.30

Examiners: Chairman - J.A.M. Howe
External - A.G. Cohn

EXPERT SYSTEMS

INSTRUCTIONS TO CANDIDATES

1. Candidates *must* answer question 1, and *either* question 2 *or* question 3.

2. Each question is marked out of 100%. The marks at the side of the questions show how these are apportioned.

3. If more than two questions are attempted, candidates should cross out the answer which is **NOT** to be marked. Otherwise, the examiners will mark only the first two answers which appear in the script.

Answer question 1, and either question 2 or 3.

Question 1.

Imagine a scenario in which we have a blocksworld description language, five positive instances of a concept and four negative instances. Assume that the positive instances are labelled 1, 5, 6, 8 and 9 while the negative instances are labelled 2, 3, 4 and 7. Running the Classification algorithm on these instances has produced the following decision tree.

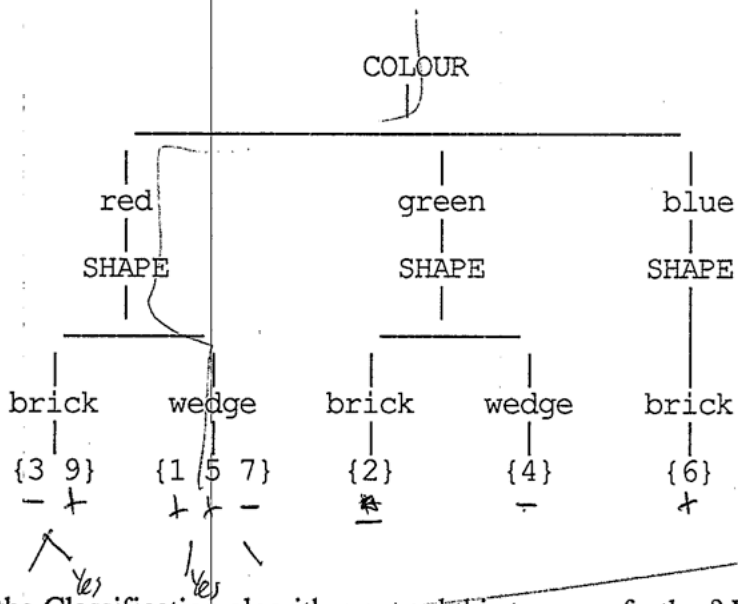

(a) Could the Classification algorithm extend this tree any further? Which nodes are "yes" nodes?

[15%]

(b) If this tree represents our total knowledge about the instances (i.e. if we don't have access to the labels shown above), how can we derive the expected information content of a message about instance no. 5?

[25%]

(c) What does it mean to say that ID3 maximises the information gain in each cycle?

[30%]

(d) Discuss the strengths and limitations of ID3 as a technique for automatic knowledge acquisition.

[30%]

Question 2.

(a) What are the principal subtasks in manual knowledge acquisition? List some of the techniques which can be used for each one.

[30%]

(b) What problems arise in manual knowledge acquisition and how can machine learning techniques be used to alleviate them?

[30%]

(c) In what circumstances could the AQ11 algorithm be used to automatically acquire inference rules? What are the main characteristics of rules generated by AQ11? Use an example to illustrate your answer.

[40%]

Question 3.

(a) What do we need to know in order to calculate the information content of an arbitrary event? What is the connection between information and uncertainty, and how can we measure our uncertainty in the case where we are confronted with N possible outcomes?

[30%]

(b) Construct an algorithm for computing the informational uncertainty with respect to a set of N possible outcomes. Assume that the probability of the i'th outcome is stored in the i'th field of a vector P.

[40%]

(c) What is Bayes formula? What problems might we encounter if we try to derive inference rules by applying Bayesian statistics to a database of conditional probabilities? What assumptions can be exploited to alleviate these problems?

[30%]