THE BCS PROFESSIONAL EXAMINATION
Professional Graduate Diploma

April 2004

EXAMINERS' REPORT

Knowledge Based Systems

Question 1

1. Data mining is the activity of finding patterns and relationships in data. Describe in detail one AI technique that could help you predict who is likely to respond to a particular marketing campaign (i.e. make a prediction) and one AI technique that identifies similar groups of customers (i.e. classification). (**15 marks)**

Individual techniques are rarely used in isolation. Why? **(5 marks)**

Answer Pointers

Typical predictive techniques include neural networks and rule induction techniques such as C5.0 and CART. Typical classification techniques include neural network techniques such as Kohonen networks. The answer expects a good overview of the algorithm – its structure and the mechanism used to generate a result. [15 marks]

It is rare to use just one technique primarily because there is no way of accurately determining which technique will work best with any particular dataset. By using multiple methods, it is also possible to use modelling algorithms together to gain the best result; for example only taking results where more than one algorithm agrees.
[5 marks]

Examiners' Comments

The question was generally well answered. Typical predictive techniques include neural networks and rule induction techniques such as C5.0 and CART. Typical classification techniques include neural network techniques such as Kohonen networks or K-means. Most candidates chose neural networks as the predictive model, and a tree-based algorithm as the classification technique.

Many people did not answer the second part of the question well. In a business environment, where people who are not necessarily skilled at using these algorithms, it is preferable to rely on the results of more than one model. This helps reduce the overall level of error by, for example, only selecting candidates that are common between two or more modelling techniques.

Question 2

2. You have been asked to manage the development project of a system that will automatically monitor a busy seaport. The remit for the project is to raise an alarm if the level of activity in the harbour becomes too great, thus risking damage to the vessels using the harbour.

You have decided to use a vision-based system to take pictures of the harbour at 5 minute intervals. Write a document describing how you will use Neural Network techniques to process the images that are taken. Your document should outline how you will train the Neural Networks and should also provide a list of potential problems. **(20 marks)**

Answer Pointers

The first part of the question should describe the standard activities associated with building a NN model. The team would need to take pictures at different times of the day, with different activity levels. This information then needs to be converted in to a form that can be read by the NN. Separately, each image would need to be flagged as to whether the harbour is too busy or not. Once the model has been trained, it should then be tested on unseen scenarios. [12 marks]

The real problem with this system is in the learning phase. For a NN to make accurate to make accurate predictions, it will be important for the system to learn on a broad range of test cases. Factors that could impact the quality of the end results include:

- How often is the harbour too busy? You may not be able to generate enough examples, thus massively reducing the quality of the result.
- Light – and reflections on the water may affect the images which will have an adverse effect on the data.
- …and what about fog?

Applications such as this have been tried, generally with limited success. [8 marks]

Examiners' Comments

This question was either answered very well or very poorly. The poor answers simply addressed the question of "what is a neural network?" – often a repeat of their answer from question 1. The good answers were able to describe a design to solve the problem using neural networks.

The first part of the question should describe the standard activities associated with building a NN model. The team would need to take pictures at different times of the day, with different activity levels. This information then needs to be converted into a form that can be read by the NN. Separately, each image would need to be flagged as to whether the harbour is too busy or not.

Question 3

3. "Historically, building an expert system has been a higher risk project than building a more traditional system such as payroll applications. However, as expert systems become more widely used in research and business, by 2005 this statement will no longer be true". Discuss**.** **(20 marks)**

Answer Pointers

Given the complexity of many modern software applications, there is an argument that the inclusion of expert system technology does not add significantly to overall levels of complexity. The real question is "What type of expert system?" Business applications using Neural Net technology can now be built in a small number of days. Vision or robotics based systems still prove to be much more complex than traditional systems. [20 marks]

Examiners' Comments

Given the complexity of many modern software applications, there is an argument that the inclusion of expert system technology does not add significantly to overall levels of complexity. The real question is, "What type of expert system?" Business applications using Neural Net or rule based technology can now be built in a small number of days. Vision or robotics based systems still prove to be much more complex than traditional systems. Many of the answers felt a little out of date – perhaps relevant in the early 1990's, but not reflecting the reality of applications such as:

- Intelligent monitoring of alarms
- Identifying target audiences for direct marketing campaigns
- Real time scoring of individuals in a call centre, identifying their likelihood of moving to a competitor.

All of which are being used today in industry to help drive business decisions and all of which include a neural network or a decision tree.

Question 4

4. Explain and illustrate how a genetic algorithm works by giving an example. You can choose your own example. You may like to answer this question by finding the maximum value of a simple function with a single parameter x in the range of 0 to 31. **(20 marks)**

Answer Pointers

Find the maximum value of a simple function with a single parameter x in the range of 0 to 31. The function for this purpose is $31x\text{-}x^2$ where x varies between 0 and 31. The genetic material is called a genome which in this case contains only a single 5-bit gene for the parameter x (5 bits are needed to represent the number from 0 to 31). Genetic algorithms work by evolving successive generations of genomes that that get progressively more and more fit. The goal is to maximise the fitness of the genomes in the population.

In nature the fitness is simple whether an organism survives to reproduce. On a computer there is much more flexibility. We choose the fitness function to solve the problem. For this example, the appropriate fitness function is $31x\text{-}x^2$. The technique consists of applying the following steps:

- Identify the genome and fitness function
- Create an initial generation of genomes
- Modify the initial population by applying selection, crossover and mutation.
- Repeat above step until the fitness of the population no longer improves.

[8 marks]

Genome	x	Fitness	%
10110	22	198	40.4

00011	3	84	17.2
00010	2	58	11.8
11001	25	150	30.6

Selection is similar to natural selection where only the fittest individuals in the population survive to pass their genetic material on to the next generation. Unlike nature, though, the size of the population usually remains constant from one generation to the next. So there is no chance of the population becoming extinct. The chance of a genome surviving to the next is proportional to its fitness value, the higher the value relative to other genomes, the more copies that survive to the next generation.

After selection randomly, the genome which has the highest fitness would have the highest chance to be selected. This becomes

Genome	x	Fitness
10110	22	198
00011	3	84
10110	22	198
11001	25	150

[4 marks]

For two of the genomes,
101| 10
000|11

After crossover, it becomes
101|11
000|10

The resulting genomes, called children, have a piece inherited from each of their parents. Applying crossover to the population proceeds by selecting pairs of genomes and flipping a coin to determine whether they cross over. This probability is the cross over probability. [4 marks]

[From Data Mining by Michael Berry and Gordon Linoff]

Examiners' Comments

This question assesses the students understanding on how a genetic algorithm is working. They are advised to use an example to demonstrate their understanding. Most of the students have answered correctly the steps of the algorithm but few of them run to the end the example they have given to themselves. This question is useful for those students with less expressive power in their narrative to gain marks by answering practical questions.

Question 5

5. Describe and explain the five basic components in Natural Language Processing with examples in each component. **(20 marks)**

Answer Pointers

4 marks per each of the components
out of the 4 marks, 1 mark for the example

Morphological Analysis
Individual words are analysed in to their components, and on-word tokens, such as punctuation, are separated from the words.

"I want to print Bill's .init file"
Morphological analysis would pull apart the word "Bill's" in to the proper noun "Bill" and the possessive suffix "s" and would recognise the sequence ".init" as a file extension that is functioning as an objective in the sentence.

Syntactic Analysis
Linear sequences of words are transformed in to structures that show the words relate to each other. Some word sequences may be rejected if they violate the language's rules for how words may be combined. For example, an English syntactic analyser would reject the sentence "boy the go the to store".

Semantics Analysis
The structures created by the syntactic analyser are assigned meanings. In other words a mapping is made between the syntactic structures and objects in the task domain. Structure for which no such mapping is possible may be rejected. For example, in most universes, the sentence "colourless green ideas sleep furiously" would be rejected as semantically anomalous.

Discourse Integration
The meaning of an individual sentence may depend on the sentences that precede it and may influence the meanings of the sentences that follow it. For example, the word "it" in the sentence, "John wanted it" depends on the prior discourse context, while the word "John" may influence the meaning of the later sentences (such as "He always had.")

Pragmatic Analysis
The structure representing what was said is reinterpreted to determine what was actually meant. For example, the sentence "do you now what time it is?" should be interpreted as a request to be told the time.

Examiners' Comments

This question is about the description and the explanation of the five basic components in Natural Language Processing with examples in each component. As it is a memorisation question, most of the students have answered correctly the question getting in some cases the full mark. It is useful question as it gives an opportunity to those students who have good memorisation to improve their marks.