THE BCS PROFESSIONAL EXAMINATIONS
Professional Graduate Diploma

April 2007

EXAMINERS' REPORT

Knowledge Based Systems

General

Table 1 shows the statistics based on all questions answered, including where some students answered four questions. Averages are fairly consistent across questions, though, a little on the high side. Standard deviation results are fairly low across all questions. Thus, it may be concluded that the candidates appear to be fairly equated in their ability, but perhaps were not challenged enough by the examination. Questions 1, 2 and 3 were the most popular amongst candidates, whereas the more general questions (4 and 5) were attempted by only a minority of candidates. It appears that candidates prefer specific questions on KBS issues and technologies rather than general discursive type questions on broader issues.

TABLE 1: stats for all questions answered

	Q1	Q2	Q3	Q4	Q5	Total
Examiner (initials)						
Number Attempted	49	37	46	26	16	
% Attempted	87.50%	66.07%	82.14%	46.43%	28.57%	
Number Accepted	49	37	46	26	16	
% Accepted	87.50%	66.07%	82.14%	46.43%	28.57%	
Number Passed	46	34	44	25	15	56
% Passed	93.88%	91.89%	95.65%	96.15%	93.75%	100.00%
Max Mark	23	20	22	23	22	97.33333333
Min Mark	0	8	8	9	8	40
Average Mark	15.33	15.03	16.83	16.42	15.38	65.57
Standard Deviation	4.16	3.20	3.37	3.64	4.27	10.82

Six students answered four questions instead of three. When the top three questions are considered alone, the average is 64.17 and standard deviation is 9.92. See table 2.

TABLE 2: stats for top three questions answered

	Q1	Q2	Q3	Q4	Q5	Total
Examiner (initials)						
Number Attempted	48	36	44	25	15	
% Attempted	85.71%	64.29%	78.57%	44.64%	26.79%	
Number Accepted	48	36	44	25	15	
% Accepted	85.71%	64.29%	78.57%	44.64%	26.79%	
Number Passed	46	33	43	24	14	56
% Passed	95.83%	91.67%	97.73%	96.00%	93.33%	100.00%
Max Mark	23	20	22	23	22	81.33333333
Min	4	8	8	9	8	40

Mark						
Average Mark	15.65	15.00	17.16	16.60	15.60	64.17
Standard Deviation	3.56	3.24	3.05	3.60	4.32	9.92

Question 1

Knowledge based Systems (KBS) are developed to deal with particular application domains in which elements of human intelligence are essential in producing solutions.

a) Identify and discuss five aspects of human intelligence that could be used to characterise intelligent knowledge-based systems. **(10 marks)**

b) Take one aspect of human intelligence and explain how it would be difficult to emulate using rule based systems. **(5 marks)**

c) Describe when and how a KBS should be tested. **(5 marks)**

d) Suggest measures that could be applied to increase the reliability of decision making by a KBS. **(5 marks)**

Model Answer Pointers

Question 1.a: 10 marks. 2 marks for each aspect of human intelligence identified and related to KBS.

Some possibilities include:

- Ability to temporarily alter behaviour according to environmental stimulus (**adaptability**)
- Ability to permanently alter behaviour as a result of accumulated experience with environmental stimuli (**learning**)
- Ability to deal with ill-defined and ambiguous situations (**uncertainty**)
- Ability to prioritise and focus (**goal directedness**)
- Ability to bring to bear subjective insight (**judgement**)
- Ability to deal with complexity and recognise relevance (**abstraction**)

Application of these traits to a medical diagnostic KBS should consider the points outlined below.

Adaptability: an investigative procedure can be refined as new symptoms emerge and new test results become available.

Abstraction: the most significant symptoms will be considered, and those may be considered as special cases of more general categories of symptoms. Connections between symptoms and causes may be established at the general level, e.g. fever is one of the flu-like symptoms, and flu-like symptoms could point to influenza, pneumonia, viral infection, etc.

Goal directedness: the diagnostic system has a main goal: to identify the cause of the symptoms; and will plan procedures that progress towards achieving that goal.

Judgement: recommendations and decisions will be based on experience with previous patient cases; thus, different experts may have different judgement (c.f. *second opinion*).

Learning: procedures for selecting and conducting tests are based on treatment of previous patients (with similar symptoms). Results of the treatment fed back to the KBS will be used to influence future decisions.

Uncertainty: prognoses about illnesses are often predictions that are based on incomplete information about the progression of the illness in the patient. That is, there can be incomplete information about the disease and about the patient.

Question 1.b: 5 marks. General distribution of marks according to salient features.

Some possibilities include:

Creativity and imagination are currently not well understood, and are thus, difficult to describe in an executable computer model. An ability to recognise faces and differentiate between male and female faces is another area where clear rules are difficult to elicit because the relevant knowledge is not easily brought to conscious attention. In general, activities that could be better described as skills (opposed to knowledge) are not amenable to knowledge based systems.

Question 1.c: 5 marks. General distribution of marks according to salient features.

Testing should be conducted throughout the entire lifecycle of systems development. It should include ensuring that the knowledge models are an accurate depiction of the knowledge as perceived by the knowledge workers and knowledge engineers, as well as stage testing of both paper and software prototypes. Al stakeholders responsible for the design and construction of a KBS should be involved, e.g. the knowledge worker being "modelled" should run simulations on the knowledge base to ensure accuracy and consistency. The software programmer should test for functionality and usability of the software with the inclusion of end users.

Question 1.d: 5 marks. General distribution of marks according to salient features.

Reliability ensues from ensuring high quality, which is facilitated through a strict testing regime. Involvement of multiple "experts" during knowledge elicitation, and testing on multiple problems at the edge of the system's capabilities (stress testing) can be beneficial. Designing for easy extension and maintenance, as well as clear guidance on expectations of use, such as the environmental conditions and the kinds of problem to be tackled, will all lead to better reliability.

Examiner's comments

The majority of candidates (87.5%) answered question 1 and made a reasonable attempt (average 15.33/25; standard deviation 4.16). Parts a and b were answered well, but part c caused a problem in that students failed to adequately describe when testing should be conducted. Most assumed that testing could be confined to *after* the system was developed rather than considering *continuous* testing during the different phases and stages of the life-cycle. Methods for testing, however, were fairly well covered. Part d caused some problem with candidates misinterpreting the question by thinking that ways to measure reliability (i.e. metrics) was required, rather than measures (i.e. methods) of increasing reliability.

Question 2

Knowledge elicitation involves modelling the knowledge used by an expert to solve problems. Consider an example application domain and construct a knowledge base for the domain by completing the following tasks:

a) Describe briefly general methods that could be applied to elicit the knowledge needed to solve a small complex problem. **(5 marks)**

b) Identify a suitable task, and produce a decision tree of five levels to represent the decision process. **(5 marks)**

c) Illustrate the consistency of the decision tree with a worked example using sample data for the task (specify any assumptions made). **(10 marks)**

d) Explain why it would be difficult to develop a KBS for a domain in which there was a considerable reliance on tacit and implicit knowledge. **(5 marks)**

Model Answer Pointers

Question 2.a: 5 marks. General distribution of marks according to salient features.
Techniques such as interviewing, observation, and automated elicitation systems should be discussed.

Question 2.b: 5 marks. General distribution of marks according to salient features.

Any task that affords five levels of decision tree should be stated, e.g. diagnosing faults in a photocopier. A sample decision tree for deciding on the amount of discount to give in a transaction is given below for three levels.

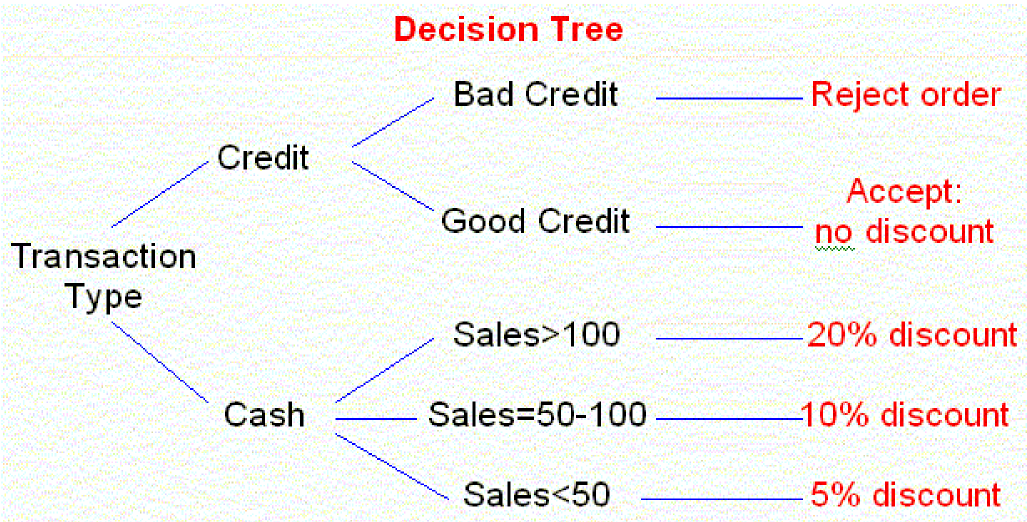

Question 2.c: 5 marks. General distribution of marks according to salient features.

Assume the following scenario: cash sale of £75.

The decision process would flow through the lower branch at level 1, then the centre branch (Sales=50-100), to conclude that a 10% discount is warranted.

The examinee must work through five levels.

Question 2.d: 5 marks. General distribution of marks according to salient features.
The nature of tacit and implicit knowledge is such that describing them in a tangible form is difficult. Such knowledge is demonstrated in use rather than being amenable to representation. KBS that attempt to emulate tacit and implicit knowledge need to overcome the problem of necessarily depending on an executable model of knowledge. A solution is to develop methods that allow tacit and implicit knowledge to emerge during execution of explicit knowledge – very challenging though.

Examiner's comments

About two thirds of candidates (66.0%) answered question 2 and made a fair attempt (average 15.03/25; standard deviation 3.20). Part a, identification of elicitation methods, was straightforward, though often an insufficient number of methods was discussed. Decision trees in part b were on the whole sufficient but some candidates failed to specify five levels. Worked examples in part c often did not provide a trace of the decision-making path, and were limited to a single scenario. Part d was more problematic since candidates failed to understand the characteristics of tacit and implicit knowledge, and therefore, could not relate those characteristics to problems with knowledge elicitation and acquisition.

Question 3

Problem solving involves both inference and searching for a solution.

a) Explain the principle of rule inference and describe both forward and backward chaining using illustrative examples. **(15 marks)**
b) Explain both brute-force and heuristic search methods and discuss their relative merits. **(10 marks)**

Model Answer Pointers

Question 3.a: 15 marks. General distribution of marks according to salient features.
Forward chaining starts with the available data and uses inference rules to extract more data (from an end user for example) until an optimal goal is reached. An inference engine using forward chaining searches the inference rules until it finds one where the **If** clause is known to be true. When found it can conclude, or infer, the **Then** clause, resulting in the addition of new information to its dataset. Inference engines will often cycle through this process until an optimal goal is reached. For example: I have a pet named Fritz, he's green and he hops, what is he? 1. **If** Fritz hops **- Then** Fritz is green 2. **If** Fritz is green **- Then** Fritz is a frog Forward-chaining inference is often called data driven — in contrast to backward-chaining inference, which is referred to as goal driven reasoning. **Backward chaining** starts with a list of goals (or a hypothesis) and works backwards to see if there are data available that will support any of these goals. An inference engine using backward chaining would search the inference rules until it finds one which has a **Then** clause that matches a desired goal. If the **If** clause of that inference rule is not known to be true, then it is added to the list of goals (in order for your goal to be confirmed you must also provide data that confirms this new rule). For example, suppose a rulebase contains two rules and that the goal is to conclude that Fritz is a frog, given that he hops: 1. **If** Fritz hops **- Then** Fritz is green 2. **If** Fritz is green **- Then** Fritz is a frog This rulebase would be searched and rule (2) would be selected, because its conclusion (the **Then** clause) matches the goal (that Fritz is a frog). It is not yet known that Fritz is green, so the **If** statement is added to the goal list (in order for Fritz to be a frog, he must be green). The rulebase is again searched and this time rule (1) is selected, because its **Then** clause matches the new goal just added to the list (that Fritz is green). The **If** clause (Fritz hops) is known to be true and therefore the goal that Fritz is a frog can be concluded (Fritz hops and therefore must be green, Fritz is green and therefore must be a frog). Because the list of goals determines which rules are selected and used, this method is called goal driven, in contrast to data-driven forward-chaining inference.

Question 3.b: 10 marks. General distribution of marks according to salient features.

A searching algorithm requires a target for which to search. The list is searched until either the target is located or the algorithm has determined that the target is not in the list. A comparison must be made to determine if the current element retrieved from the collection is the target one; therefore, a measure of similarity is needed. One of the fields, called the *key* field, serves as the measure on which comparison is performed.

The set of all possible solutions to a problem is called the search space. Brute-force search or uninformed search algorithms use the simplest, most intuitive method of searching through the search space, whereas informed search algorithms use heuristics to apply knowledge about the structure of the search space to try to reduce the amount of time spent searching.

Uninformed search

An uninformed search algorithm is one that does not take into account the specific nature of the problem. As such, they can be implemented in general, and then the same implementation can be used in a wide range of problems thanks to abstraction. The drawback is that most search spaces are extremely large, and an uninformed search (especially of a tree) will take a reasonable amount of time only for small examples. As such, to speed up the process, sometimes only an informed search will do.

Informed search

In an informed search, a heuristic that is specific to the problem is used as a guide. A good heuristic will make an informed search dramatically out-perform any uninformed search.

There are few prominent informed list-search algorithms. A possible member of that category is a hash table with a hashing function that is a heuristic based on the problem at hand. Most informed search algorithms explore trees, such as the Best-first search, which is a search with a heuristic that attempts to predict how close the end of a path is to a solution, so that paths which are judged to be closer to a solution are extended first. Efficient selection of the current best candidate for extension is typically implemented using a priority queue.

Examiner's comments

Most candidates (82.14%) answered question 3 and made a good attempt (average 16.83/25; standard deviation 3.37). In part a answers tended to neglect discussion of rule inference in general and instead concentrated on chaining mechanisms. Part b was answered well except that the merits of different search techniques was not discussed adequately.

Question 4

Development of any software system requires careful assessment of risk in order to ensure that a quality product is produced. Consider the risks that should be anticipated for the construction of an interactive rule-based advisory system by addressing the following tasks:

a) Identify all the main stakeholders involved in the project, and explain their respective roles and responsibilities in the construction of an intelligent knowledge based system.

(10 marks)

b) Explain five significant risks to the success of a KBS development project, and for each describe suitable measures that should be taken to mitigate against the risk.

(15 marks)

Model Answer Pointers

Question 4.a: 10 marks. General distribution of marks according to salient features.
There are generally three individuals having an interaction with expert systems. Primary among these is the end-user; the individual who uses the system for its problem solving assistance. In the building and maintenance of the system there are two other roles: the problem domain expert who builds the knowledge base, and a knowledge engineer who assists the experts in determining the representation of their knowledge and who defines the inference technique required to obtain useful problem solving activity. Additionally, other stakeholders are the project manager, who takes operational control of the project, and the board of directors and clients, who are responsible for commissioning and overseeing the project.

Question 4.b: 15 marks. 3 marks for each risk identified and discussed with salient features.
Some possibilities include: Technology risks: new techniques are needed that prove too costly or complex to employ. Solution, ensure that such technology is fully tested and available, and that sufficiently skilled staff are available on the project. Business risks: the business environment could alter and render the KBS inadequate or obsolete if it takes extensive time to develop. Solution, use agile approaches to development in which prototypes are produced quickly and the client gains benefit from the project early on. Personnel risk: availability of necessary experts and developers may be unreliable over the course of time, particularly where expertise is a scarce commodity. Solution, recruit multiple experts and ensure knowledge elicitation is completed early in the project. Other risks to discuss could include financial and resource-based.

Examiner's comments

About a half of candidates (46.43%) answered question 4 and made a good attempt (average 16.42/25; standard deviation 3.64). Identification of stakeholders in part a caused little problem, though, discussion of their roles and responsibilities was often not discussed in sufficient detail. Part b was more problematic since candidates tended to list potential but *preventable* problems and their solutions (e.g. incomplete requirements), rather than discussing possible and *unavoidable* risks (e.g. changes in technical standards, or unavailability of key personnel and technologies.)

Question 5

Artificial Intelligence (AI) arguably has struggled to make an impact on real world solutions. Select any AI technology and discuss how it has successfully made the transition from laboratory to workplace or home. Focus on the real problems that the technology addresses and the benefits perceived to have been realised. **(25 marks)**

Model Answer Pointers

Question 5: 25 marks. General distribution of marks according to salient features.
Any AI technology is acceptable provided it is identified clearly and its application is illustrated with examples. For instance, fuzzy logic is often used in consumer products such as cameras (autofocus), or washing machine controllers. Data mining is now often employed in business, intelligent search engines on the web (Google), games, consumer robots with vision systems, handwriting reading in PDA, speech synthesis in word processors, intelligent tutoring, knowledge management, etc.

Examiner's comments

About one quarter of candidates (28.57%) answered question 5 and made a reasonable attempt (average 15.38/25; standard deviation 4.27). Different technologies were addressed. Common problems were (a) answers covered several technologies briefly rather than one technology in depth; (b) answers described the technology in depth but did not give attention to its application; (c) answers focused on the technology and/or its application, but did not discuss the benefits derived.