



APPARATUS FOR TRAINING OR ANALYSIS

5 This invention relates to an apparatus for developing and testing the knowledge and skills of an individual, particularly through the generation of stimuli customised to the requirements of the individual and the monitoring of the individual's responses to further customise the generation of stimuli to the individual's requirements.

10 In many applications the current state of an individuals' knowledge, ability, or reflex must be determined before the appropriate therapy, or instruction can be determined and then administered. Examples of such applications include: teaching of factual information; development of language skills; improving an individuals perception relating to colour, auditory pitch and motion; and psychological testing.

15 Many techniques exist for testing and measuring knowledge, reflex, and ability; however, in many cases they are manually implemented which requires trained supervisors and are time consuming. In all cases, methods for measurement are unique to the type of knowledge, reflex, or skill that is being tested, so no single method is universally applicable to a wide range of testing. This requires that diagnostic technicians be specifically trained for
20 each different test method.

The breadth of information that is tested for in the state of the art is limited again due to limited resources and time for each test. For example, colour blindness is not simply the inability to recognise a colour, but is also correlated to contrast, context, hue, brightness, etc., but the battery of tests
25 necessary to conduct all these measurements are not cost effective enough to justify their use in most cases. Some measurements such as the time used to determine each colour, which is useful to determine the confidence, ease, and



clarity with which each colour is perceived, is so labour intensive, and difficult to carry out accurately that they are never used clinically.

According to a first aspect of the invention there is provided an apparatus for training or analysis which is arranged to store a plurality of data elements
5 which are associable with further stored data elements to form outputs to a user and predetermined values against which a user's inputs can be validated, respectively. The outputs and inputs may be in the form of questions and answers.

Expressed differently, the normal way in which knowledge is acquired is
10 through the formation of associations, which may range from the very simple to the highly complex. An example of an association is how a written word is spoken. Any learning process is therefore dependent upon the subject matter, or content of the association. In the above example, this would be the written and the spoken word. Therefore, methods of teaching one type of subject
15 matter are generally not easily adapted to effectively teach another type of subject matter. However, in the present invention, by disassociating the mechanism by which associations are generated from the subject matter of the association, a very wide range of applicability is achieved. Thus, the present invention is applicable not just to the learning of factual based knowledge but
20 also the development of personal skills and even, for example, the implementation of psychological tests.

According to a second aspect of the invention there is provided an apparatus for training or analysis which is arranged to output stimuli to a user through the medium of a video arcade game, through the playing of which a user is
25 able to respond to the output stimuli.

The use of an interactive arcade game style interface has several valuable effects on the learning process of users. Firstly, it provides a stimulating learning interface which serves to motivate a user. However, more importantly by causing a reflexive response to an association, whilst at the

same time providing a distraction to the user, this results in the level to which the association is learnt, or acquired, being greatly increased by making the user's ability to manipulate the subject matter in question more fluent. This means that an individual can recall a given association and response with very little thought or effort. Experiments have shown that information or skills learnt to such a degree are most readily available when needed and less likely to be forgotten.

According to a third aspect of the invention there is provided an apparatus for training or analysis which is arranged to output stimuli to a user and to receive corresponding inputs from the user which are validated, the validation results being used to alter the frequency with which the subsequent outputs are made. Advantageously, the present invention increases both the speed at which an individual may learn or develop skills, and also the degree to which this learning is achieved. This is accomplished by instilling a sense of urgency in the individual, both to make an association and also to respond to it. This has the effect of increasing the level of concentration in the individual as well as increasing their level of motivation, through creating a sense of competition. The present invention achieves this through the use of encouraging an individual to respond in a given time.

According to a fourth aspect of the invention there is provided an apparatus for training or analysis in which the apparatus is arranged to output questions to a user, the questions being derived from a set of question data, and receive corresponding answers from the user, the validity of which is monitored and used by the apparatus in order to keep the proportion of data in the set which is known by the user at a specified level, by varying the size and/or content of the data set.

This feature allows a user to gradually progress to new data or information as he or she becomes more proficient with or learns the data with which he or she is currently dealing. This has the advantage of ensuring that progress is

made at a rate which is suitable for the user and avoids the shock of being introduced to a totally new set of information once the previous set has been fully mastered. Additionally, it also ensures that the user will not be subjected to too great a quantity of unfamiliar information at once, which would have the effect of hindering the learning or familiarisation process of the user with the information concerned.

According to a fifth aspect of the invention, there is provided an apparatus for training or analysis which is arranged to output stimuli to a user and receive corresponding inputs from the user; the apparatus being further arranged to monitor the types of stimulus to which the user inputs correct responses and the type of stimulus to which the user inputs incorrect responses; and to vary the future generation of these types of stimuli in dependence upon the success of the user in responding to them.

Advantageously, this ensures that the user of the apparatus may be made to concentrate upon stimuli, or questions for example, which they find difficult. By the same token this also ensures that the user does not waste time concentrating on subject matter which they are familiar with and thus respond correctly to.

According to a sixth aspect of the invention there is provided an apparatus for training or analysis in which the apparatus is arranged to generate stimuli for output to a user and to receive corresponding inputs from the user, the validity of which is determined. By periodically outputting stimuli which are either identical or similar to previously output stimuli, the rate at which the user is becoming familiar with the information of which the stimuli comprise may be determined.

This aspect of the invention provides several advantages, both for the user and for those involved with the implementation of training or analysis. It enables the user to obtain feedback on their rate of progress in a learning situation and

may be used by the user in order to maximise their progress in further learning situations. It also allows teachers or analysts to obtain feedback as to the effectiveness of particular teaching or analytical methods.

According to a seventh aspect of the invention there is provided an apparatus
5 for training or analysis which is arranged to output stimuli to a user and to receive corresponding response from the user. The generation of the stimuli is achieved by the selection of a first data item and an association, thereby defining a correct response; the stimulus and the response being further determined by the selection of a format and a medium in which they must be
10 made.

This process for generating output stimuli allows a great variation of output stimuli to be generated and hence the suitable responses to the stimuli may also be greatly varied. In a learning environment, this allows a student to learn to associate a particular item of knowledge in a great many ways, which
15 has the effect of ensuring that it may be thoroughly learnt. This process is further augmented by the fact that the student may be asked to respond through greatly varying media and format combinations.

According to a eighth aspect of the invention there is provided an apparatus
20 for training or analysis which is arranged to output a series of questions to a user and to validate the user's responses. By measuring the accuracy of the responses in a particular field of questions, the apparatus is arranged to measure the degree to which the user knows the information of which they comprise. Thus, when the user makes an error in responding to a question, the apparatus is arranged to implement a remedial course of questions which may
25 be adapted in accordance with the user's knowledge of that subject matter.

In the event that a user has a good knowledge of the subject matter under consideration, a momentary lapse in memory may result in an incorrect response. This may be rectified by briefly reminding the user of the correct

response. However, in the event that the user has only a limited knowledge of the subject matter under consideration, an incorrect response may be more suitably responded to by the provision of a series of more simple questions on the same subject, which lead the user gradually back to the level of difficulty of the question which he or she got wrong.

According to a ninth aspect of the invention there is provided an apparatus for training or analysis in which the apparatus is arranged to output a series of stimuli to a user, where each stimulus comprises two or more outputs which may be distinguished between by the user; the apparatus being arranged to make the task of distinguishing between the outputs of each stimulus more difficult with each successful stimulus.

In this aspect of the invention, a user may be lead gradually through a series of tests in order to find out his or her ability in a particular area; for example sensory perception. Once the user's ability has been determined, the apparatus may be used to generate further stimuli at a level of discrimination difficulty that runs from slightly below to slightly above the user's maximum ability; thus allowing the user the opportunity to improve their skills in this area.

Other inventive aspects and features of the invention will be apparent from the following description, and claims, as will advantages of these aspects of the invention.

The invention will now be illustrated, by means of example only, with reference to the following drawings in which;

Figure 1. **Basic system hardware.** Demonstrates the hardware configuration of 1 possible embodiment of the invention. Also details the connection and interaction of the different parts.

Figure 2. **Basic technology diagram.** Demonstrates in broad terms the general interaction of the user and the invention.

Figure 3. **Object diagram.** Demonstrates in broad terms the overall structure and mechanism of the invention.

Figure 4. **Main menu screen.** Demonstrates the choices and options available to the user while in main menu screen.

5 Figure 5. **Content data structure I.** Demonstrates the general data structure for language content.

Figure 6. **Content data structure II.** Demonstrates the general data structure for knowledge content.

Figure 7. **Content data structure III.** Demonstrates the general data structure for scenario content.

Figure 8. **Content data structure IV.** Demonstrates the general data structure for psychometric test content.

Figure 9. **Quiz matrix structure.** Demonstrates the organisation and structure of the quiz types.

15 Figure 10. **Quiz type example I.** Demonstrates one possible question type generated by the invention.

Figure 11. **Quiz type example II.** Demonstrates one possible question type generated by the invention.

Figure 12. Quiz type example III. Demonstrates one possible question type generated by the invention.

Figure 13. **Module diagram I.** Demonstrates the structure of the quiz module.

Figure 14. **Module diagram II.** Demonstrates the structure of the notebook module.

25 **Figure 15. Module diagram III.** Demonstrates the structure of the content
module.

Figure 16. **Module diagram IV.** Demonstrates the structure of the course module.

Figure 17. **Module diagram V.** Demonstrates the structure of the ICC module.

Figure 18. **Module diagram VI.** Demonstrates the structure of the ICC configuration module.

5 Figure 19. **Module diagram VII.** Demonstrates the structure of the quiz configuration module.

Figure 20. **Notebook type examples I.** Demonstrates the notebook List view.

10 Figure 21. **Notebook type examples II.** Demonstrates the notebook Association view.

Figure 22. **Content selection.** Demonstrates the means by which the system selects and accesses content.

Figure 23. **Question type selection I.** Demonstrates the means by which the system selects the question types from a possible 2,500.

15 Figure 24. **Question type selection II.** Demonstrates the means by which the system selects the question types from a possible 2,500.

Figure 25. **Question type selection III.** Demonstrates the means by which the system selects the question types from a possible 2,500.

20 Figure 26. **Question matrix rules.** Demonstrates the rules and structure to determine valid question types.

Figure 27. **Question type display I.** Demonstrates the means by which the system displays a question once the question types have been chosen.

25 Figure 28. **Question type display II.** Demonstrates the means by which the system displays a question once the question types have been chosen.

Figure 29. **Quiz question pattern.** Demonstrates the means for generating the different possible question patterns.

- Figure 30. **Quiz flow.** Demonstrates the real time flow of events during a quiz.
- Figure 31. **Answer response.** Demonstrates the means by which the system analyses the users response.
- 5 Figure 32. **CDR calculation.** Demonstrates the different methods of calculating the CDR as determined by the user response.
- Figure 33. **QDR calculation.** Demonstrates the different methods of calculating the QDR as determined by the user response.
- Figure 34. **Timer calculation.** Demonstrates the different methods of calculating the timer value as determined by the user response.
- 10 Figure 35. **Error response I.** Demonstrates the mechanism for the system to respond to user errors.
- Figure 36. **Error response II.** Demonstrates the mechanism for the system to respond to the users errors.
- 15 Figure 37. **Error response III.** Demonstrates the mechanism for the system to respond to the users errors.
- Figure 38. **ICC-depth & width** Demonstrates the algorithms and the meaning of depth & width.
- Figure 39. **Content flow.** Demonstrates the measurement, meaning, and application of content flow.
- 20 Figure 40. **CDR / QDR / familiarity / depth / time.** Demonstrates the relationship between CDR, QDR, familiarity, depth, and time.
- Figure 41. **ICC mechanism I.** Demonstrates how the entire system is automated by the values calculated by ICC.
- 25 Figure 42. **ICC mechanism II.** Demonstrates how the entire system is automated by the values calculated by ICC.
- Figure 43. **Intelligence type A I.** Demonstrates one type of intelligence drill.
- Figure 44. **Intelligence type A II.** Demonstrates one type of intelligence drill.

- Figure 45. **Intelligence type A III.** Demonstrates one type of intelligence drill.
- Figure 46. **Intelligence type A IV.** Demonstrates one type of intelligence drill.
- 5 Figure 47. **Intelligence type A V.** Demonstrates one type of intelligence drill.
- Figure 48. **Intelligence type A VI.** Demonstrates one type of intelligence drill.
- Figure 49. **Intelligence type B I.** Demonstrates one type of intelligence drill.
- 10 Figure 50. **Intelligence type B II.** Demonstrates one type of intelligence drill.
- Figure 51. **Intelligence type B III.** Demonstrates one type of intelligence drill.
- Figure 52. **Intelligence type B IV.** Demonstrates one type of intelligence drill.
- 15 Figure 53. **Intelligence type B V.** Demonstrates one type of intelligence drill.
- Figure 54. **Intelligence type B VI.** Demonstrates one type of intelligence drill.
- 20 Figure 55. **User profile for colour.** Demonstrates the structure of computer generated content.
- Figure 56. **SDI (*Stimulus Drive Interface*) type A.** Demonstrates one type of SDI.
- Figure 57. **SDI (*Stimulus Drive Interface*) type B.** Demonstrates one type of SDI.
- 25 Figure 58. **SDI (*Stimulus Drive Interface*) type C.** Demonstrates one type of SDI.
- Figure 59. **SDI (*Stimulus Drive Interface*) type D.** Demonstrates one type of SDI.

Figure 60. **User profile for pattern.** Demonstrates the structure of computer generated content.

Figure 61. **Quiz type example IV.** Demonstrates two possible question type generated by the invention.

5 Figure 62. **Customised Diagnostic Profile.** Demonstrates the content and structure of the user profile.

Figure 63. **Computer generated data structure.** Shows the structure of computer generated content.

First Embodiment

10 Referring to Figure 1, an embodiment of the invention comprises 3 general categories of components grouped by function: output apparatus 1.10, input apparatus 1.12, and measurement/analysis and processing apparatus 1.11.

Output apparatus 1.10 includes: a display unit 1.1 (typically a colour CRT VDU capable of average resolution displays); speakers 1.2 capable of
15 reproducing sound within the audible hearing range; a printer 1.4, capable of at least black and white text output. Other output devices could also be used in this embodiment, either in conjunction or instead of those already mentioned. Their primary purpose is to produce a perceptible stimulus, to which the user may respond through the input apparatus 1.12; therefore such
20 devices may include: flashing lights; tactile stimulators; scent stimulators, and so on.

Input apparatus 1.12 includes a microphone 1.7, capable of detecting the human voice; a keyboard 1.8 such as that of a standard personal computer, or a key board typically used for producing music; a mouse 1.9 which enables
25 the user to make selections displayed on the display unit 1.1. The primary purpose of the input device is to allow the user to respond to the stimulus generated by the output devices 1.10, therefore any other suitable device could also be used in this embodiment, either in conjunction or instead of those

already mentioned. Input apparatus 1.12 has further functions that will be described below in further detail.

Both the input and the output apparatus also include removable storage 1.3, such as standard 3.5 inch floppy disks, for the purpose of storing instruction modules, course content, user profiles, and other products of the system.

5 The measurement/analysis and processing apparatus 1.11, includes; a central processing unit (CPU) 1.6 together with its associated program memory 1.5, and is conveniently an IBM compatible personal computer. The primary purpose of the measurement/analysis and processing apparatus 1.11, is to
10 manage the user input through input apparatus 1.12, make the required calculations specified in detail further on, store required information in memory 1.5, and manage the system output through the output apparatus 1.10.

The general interaction between the user and the system will now be described with reference to figure 2. Figure 2 is divided into 2 distinct parts,
15 2.14 and 2.15. 2.14 includes all parts to the left of the dotted line and illustrates in broad terms the interactions of the constituent parts of the system of the present embodiment, as implemented through the measurement/analysing and processing apparatus 1.11. 2.15 includes all parts to the right of the dotted line, and illustrates how users may interact with the
20 system.

To measure the user's response to stimulus of any sort, the system must first generate a stimulus. A stimulus is often generated within the context of a question by the stimulus generator 2.1. The stimulus generator 2.1 is comprised of content 2.2, context 2.3, and the Intelligent Course
25 Customisation (ICC) 2.4.

Content 2.2 is the information of which the stimulus comprises and is stored in memory 1.5. For example, if the content is language, then the stimulus

could be in the form of synthesised or textual words or phrases. Various content types and their data structures will be described below in detail.

Context 2.3 determines the structure of the content, i.e. it is the medium in which the stimulus is conveyed to the user. One example of a language based context is a textual word and its text definition. Various context types and their structure will be described in detail below.

ICC 2.4 determines which content to use with which context type, and sets the order and frequency with which they are used. ICC 2.4 manages this process with data obtained from previous user interaction.

10 The stimulus generator 2.1 produces a stimulus 2.8 output through any of the devices of output apparatus 1.10. At this point the user 2.11 may respond to the stimulus in a predetermined manner for each context type. The user 2.11 or users 2.10, e.g. a class, may use any predetermined input device of input apparatus 1.12 to respond to the stimulus.

15 The user's response is passed to the QED (*Quantified Educational Data*) 2.6 where a quantity of data is captured and then stored in the CDP (*Customised Diagnostic Profile*) 2.7. Once the data has been captured and stored in the CDP 2.7, the ICC 2.4 may access them in order to prepare the content 2.2 and the context 2.3 for a further stimulus 2.8.

20 Periodically, the data from the CDP 2.7 is transferred to the EDSS (*Educational Decision Support System*) 2.9. The EDSS 2.9 performs calculations and analysis on this data which enables the ICC 2.4 to work more effectively in customising future stimuli to the requirements of the user or users, as well as generating reports, for use by the TPM (*Training Program Manager*) 2.13 and the instructor 2.12, on the progress of the user or users.

25 The overall structure of the system of the present embodiment, as implemented through the measurement/analysis and processing apparatus 1.11, will now be described with reference to figure 3. A system object 3.1

manages the interaction between the system and one or more users, the content which they are using and the necessary interfaces, which may be required between them. The system object 3.1 also manages the import and export of modules which will be described in detail further on.

- 5 The network object 3.2 manages the transfer of data from the CDP (*Customised Diagnostic Profile*) 2.7 to the EDSS 2.9. The network object 3.2 also manages the transmission of modules which will be described in detail further on.

- 10 The interface object 3.3 manages all the inputs and outputs of the system. It determines how the content, within a specified context, will be generated for the user. It also determines how the user may respond to the stimulus once it has been generated. The interface object 3.3 requires that the content and context be determined first before it acts. The content is determined and prepared by the content object 3.9. The context is generated by either the
- 15 Quiz object 3.8, or the notebook object 3.4. These objects co-ordinate a range of settings and configurations that define the parameters by which the interface object 3.3 operates.

- The system is configurable in a variety of ways which may be done entirely by the ICC (*Intelligent Course Customisation*) 2.4 or partly by the user
- 20 through the use of the wizard and configuration object 3.6. These configurations will determine how the interface object 3.3 operates. Finally, the information palette object 3.5, enables the user to configure the information about their performance which is displayed.

- Figure 4 depicts the initial screen , as displayed on screen 1.1, which displays
- 25 three option buttons. The first button 4.1 is the course button. The user may click on it using the mouse to start a course. The course may be set by the user, or set by the ICC. To change and configure any of the settings, the user may right click on the course button. A configuration wizard appears that

leads the user through all the possible steps. The content of the configurations will be described further on. The quiz button 4.2, enables the user to manually select and start a quiz. The notebook button 4.3, enables the user to select and start a notebook module. The details of each the course, quiz, and
5 notebook options will be described further on.

Notebook display

The notebook displays content for the user to view. The design improves content assimilation by actively leading the user through it at a steady pace. Two examples of the notebook display are shown in figures 20 and 21.

10 Figure 20 shows the list view 20.1 which displays information in a list of a variable number of rows with 2 columns. The left column shows or plays any combination of the association element and the media type 20.4. The right column shows any other combination of associative element and media type 20.5. These may be selected manually by the user, or generated by the ICC
15 (*Intelligent Course Customisation*) based on the user profile or CDP. For example, the list view 20.1 may display a specified session's error content, and only the specific association elements that the user erred on. With the automated notebook feature, the content is displayed element by element at a suitable rate. The mechanism for the automated notebook uses the
20 information in the notebook module.

The notebook module contains all the information necessary to automate the notebook. The association element, media type, timer value, and sequence, are the values that are used directly by the display objects. If the display object is a text object, then simply loading the above values into the object
25 produces the desired effect. A timer object accepts the time values and turns the objects on or off as required. The result is information that is displayed (or stimulus which is generated) at a rate suited to the users ability to follow.

The association view 21.1 displays only one content record at a time, but may show 2 media types for each association element. For example, the left side 21.3 may display native text word on the top, and below it target text word. The right side 21.2 may display target text synonym, and target audio synonym. With the automated notebook, the display may be manipulated to show the content in any timed sequence.

Content Data Structure

A detailed description of the data structure for several content types will now follow; these include language, knowledge, scenario and pre-defined output content types. Language content may contain the following types of content: words in any language, alphabets or characters of any language, idioms and expressions. Knowledge content includes any fact based information such as history, geography, medicine, the sciences, philosophy, engineering, linguistics, statistics, agriculture, architecture, law, corporate training methodologies etc. Scenario content may contain any information which conveys a process or scenario, such as video images, animations and flow charts. Predefined output content types will be explained in greater detail below, but may include any of the aforementioned information types. The data structure in each of figures 5 to 8 shows the data of only one record per table. In all cases when the system is operating, the data structures will be populated by a plurality of records; the structures of which are identical to those shown in the figures.

Figure 5 illustrates the data structure of a language based content record. This comprises content information 5.2 and non-content information 5.8; as previously mentioned the content information consists of the substance or information of which the stimulus, or in specific cases, questions, comprise. The non-content information comprises information which allows each record to be effectively used.

When learning a language it is common to start by making associations with something with which the student is already familiar, i.e. the student's native language. In this case the native language is English and the target language is Japanese. Therefore Fig. 5 also shows the corresponding Japanese record

5 5.6. Within the native language 5.1 record there are 3 possible media types; text 5.3, audio 5.4, and pronunciation 5.5. However, the available media types vary according to the actual content type under consideration.

Each native record language 5.1 contains associative elements 5.2. These, in combination with a specific media type yield content. For example, word

10 5.2a and text 5.3 contain a text depiction of a word, such as the word "dog". The same associative element 5.2a in combination with the audio 5.4 media type contains an audio representation of the word "dog", such as a recording of someone saying the word. The same associative element 5.2.a, in combination with pronunciation 5.5 yields a phonetic depiction of the word

15 using phonetic notation.

In this case, the other possible associative elements 5.2 include synonym, antonym, definition, and sample sentence. The actual number of associative elements may be more or less depending on the content type together with the perceived needs of the users. The system is arranged to allow for a variable

20 number of associative elements for different content.

The target language 5.6 record in this case is Japanese and contains the same media types, non-content and content information. The associative elements 5.7 are the same as those in the native language record 5.1, except that are in Japanese as opposed to English.

25 Native language content is stored in one table which contains sufficient records to meet the purposes of the system, which in this case is the instruction of English students in learning Japanese. Each record contains the combination of associative elements and media types described in figure 5.

Target language content is also stored in its own table which contains the same number of records as the native language table. Each record in both tables has a unique identification (ID) number which links it to the equivalent record in the other table. For example, if the word "dog" is number 102 in the native language table, then the word "inu" which means "dog" in Japanese, may be number 102 in the target language table.

It is not significant whether the system can or cannot store multiple media types together, because in the case that it does not, directory addresses can be stored in place of the media type or types which cannot be stored. This will enable the system to find the location of a particular audio file, or graphic file, for example.

Native language content record 5.1 contains 5 associative elements, and 3 media types, giving rise to 15 separate pieces of information associated with each record, or 15 content fields which do not include non-content information 5.8, which in this example includes information such as ID number, lesson number, lesson name, and part of speech. Target language content 5.6 also contains 5 associative elements, and 3 media types which are directly equivalent to the same association element and media type in the native content. These also total 15 content fields.

Totalling the number of content fields in native and target content produces 30 content fields. The data contained in the 30 separate fields may be combined in any pair combination to form an association. 900 such associations are possible with 30 fields. An example of an association is text word, and audio synonym. Each associative element may be assigned a degree of difficulty, thereby producing a value for each pair combination as well. Some associations are so easy as to be meaningless in certain contexts. For example when learning a foreign language, associating a word in your native language with its definition in your native language, is unnecessary. However, in a

different context, where one is trying to expand ones vocabulary in their own native language, then the association is pertinent. An association which pairs an antonym in the target language with a sample sentence, using the word in question also in the target language, is considered a difficult association.

5 Figure 6 demonstrates the data structure for a knowledge based content record, in this case, human anatomy 6.1. Similar to language content, knowledge based content is comprised of associative elements 6.2, and media types, although there is only one table. There is no equivalent to the target language content tables as in the language example. Content ID, lesson
10 number and lesson name make-up the non-content information 6.5 contained in the record. The associative elements 6.2 in this record are: anatomical part; function; description; location; component parts; and sample sentence. Again, 3 different media types are available (text 6.3, audio 6.4, and graphic 6.4a) although the graphic media type 6.4a is only available for the anatomical part
15 in this instance. Therefore, in this instance there are 6 associative elements and 3 media types. Therefore, the total number of content fields in the human anatomy content 6.1 equals 13.

The data contained in the 13 fields may be combined into 169 different pairs, or associations. If each associative element and media type is given a
20 difficulty weighting, then each pair will also contain a difficulty value. Knowledge based content is different from language based content; however, the basic units, associative elements and media types, can be manipulated in the same manner. The system may manipulate these basic units in identical ways despite the difference in content types, as will be described in full detail
25 below.

Figure 7 illustrates the structure of a scenario based content record. Scenario 7.1 refers to content which represents a process, for example, a banking procedure which requires several steps. The association types 7.2 include

procedure name, procedure purpose, procedure result, and procedure steps-a to c. The available media types in this case include animation, which graphically demonstrates each step in the process (a to c), in addition to the audio and text media types previously discussed. If the process concerned
5 involves computers, for example, the animation would display first the method for starting the particular process, then the sequence of events that follow in order to complete the task. If the process has multiple steps then it may be broken into several separate animations as in 7.2 which includes 3 steps, step-a, step-b, and step-c. The number of steps depends on the nature of
10 the content, and could be greater or fewer than the three shown. If the process branches into different paths, then each path may also be represented as a separate step.

In this case, animation is only available for those associative elements which are procedure steps. The total number of content fields for figure 7 is
15 therefore 15. With 15 content fields the possible number of pair combinations equals 225. As previously discussed each associative element and media type may be assigned a difficulty weighting, which may then be used to calculate the difficulty of each pair, or association.

The content itself is different from knowledge based and language based
20 content, but again, the structure is similar in that it may be broken down into associative elements and media types, which may be manipulated similarly, regardless of the content type.

Figure 8 demonstrates the data structure for records which result in output stimuli are pre-defined, and are not dynamically generated; i.e. where set
25 stimuli are taken in isolation from a stored list without first being paired or combined with further associative elements to create an association. This type of stimulus is required, for example, in the implementation of psychometric tests.

Psychometric test 8.1 contains only one possible association, or combined pair per content record. In other words, each record 8.1 comprises a single pre-defined question 8.2 to which the user may respond by selecting any one of pre-defined answers 1 to 5 8.3. Therefore, a psychometric test will utilise a
5 separate record for each question. However, since each answer is not simply a matter of being correct or incorrect, and usually has its own significance in terms of identifying tendencies, answer codes 8.3 are included for each possible answer. Answer codes 8.3 show which answers in a group of questions are related in meaning, thereby enabling the system to calculate the
10 relative weight of the tendencies which are being measured. The methods of calculation and analyses will be explained in more detail further on.

The basic structure for the type of content in psychometric test 8.1 is nevertheless the same as the previous content types listed in figures 5, 6, and 7; albeit, uni-dimensional. The media type need consist of only one media
15 type, although more may be used if desired. This is in most cases text. The association types consist only of an output and a corresponding input, (a question and an answer), these must be paired in a specified order, namely, "question", then "answer". The total number of paired combinations is therefore 1.

20 Skill based content is a further content type, which is different from the content types described above in that the content itself is mostly generated by the system. The data structure is more complex than the structures so far described in that it involves more tables or matrices. Rather than just the associative elements, and media types being required several other tables or
25 matrices which may vary from one skill based content to another. A detailed explanation of the skill based data structure will follow further on.

Quiz Matrix Structure

A quiz is defined as a series of output stimuli in the form of one or more question types applied to one or more content records. A question type refers to the particular combination of association types, media types, and, as will be detailed further on, stimulus formats. In the description of the content data structure, the number of content fields for each record was used to determine the possible number of combined pairs, or association types. One combined pair, or association type includes, for example, audio word and text synonym from the native language content type 5.1. All association types may be formatted as a question. For example, the audio word could be played as the question, and the text synonym, along with 3 other text synonyms from other content records, may be displayed as multiple choice answer selections from which the user must choose the text synonym which corresponds to the audio word in the question. For the native language content type 5.1, the total number of paired combinations, or association types was 900. This means that within the multiple choice question format, 900 possible question types are possible by changing the media and the association types.

The content utilised by the system is organised by type using matrices. Figure 9 illustrates in more detail, the matrices by which the content is organised. It also illustrates further matrices relating to the question type format of the content.

The first matrix is the stimulus format matrix 9.1. Each stimulus format represents a format in which a stimulus or question may be delivered, together with the format in which the user's response or answer must be delivered. For example, one type of stimulus format is multiple choice, where the response must be given by selecting the correct answer or answers from the choices given.

The system of the present embodiment operates with 57 different stimulus formats, as represented by 9.4; however, this number could be greater or

smaller depending on the particular requirements of the content. An example of a stimulus format is dictation, where an audio association is played after which the user must type what was played. A more detailed illustration of some of the possible stimulus formats can be found in figures 10 through 12.

5 The stimulus format matrix 9.1 includes QDR (*Question Difficulty Rating*) 9.2 which is a relative weight which determines how easy or difficult a stimulus format may be. For example, a multiple choice stimulus format would be easier than a "rote memory" format where the user, upon hearing the target audio definition, is then requested to type in the corresponding antonym with
10 no choices to select from. Each stimulus format also has a corresponding ratio 9.3 associated with it. The ratio determines the relative frequency with which a particular stimulus format is used within a course. Courses will be described in more detail further on.

The media matrix 9.5 contains information on the type of media used to
15 express a stimulus; for example, text, audio or pronunciation, as was described in the explanation of the language data structure. The media matrix media types 9.8 are determined by the content type, and may vary accordingly to suit the requirements of the content. The QDR (*Question Difficulty Rating*) 9.6 and the media matrix ratio 9.7 serve the same functions as previously
20 described with reference to the stimulus format matrix.

The association matrix 9.9 contains information concerning the types of association which can be made with the content type under consideration. For example with language based content type suitable associations include antonym and synonym, whereas for a knowledge based record such as human
25 anatomy suitable associations include location and function. Again, the QDR (*Question Difficulty Rating*) 9.10 and the association ratio 9.11 serve the same functions as previously described with reference to the stimulus format matrix.

Matrix n 9.13 represents the possibility for other matrix types to be added as may be required by the content under consideration. An example of content that uses added matrix types is "skill" based content. The matrix content 9.16 is necessarily also dependent on the content requirements. Again, the QDR
5 (*Question Difficulty Rating*) 9.10 and the association ratio 9.11 serve the same functions as previously described with reference to the skill matrix.

In most content types, only 3 matrices are necessary to represent all the possible question type choices. An example of this will now be described.

One value each from the stimulus format matrix 9.1, the media matrix 9.5,
10 and the association matrix 9.9, when combined produce the parameters with which the system may build and display a stimulus, which in this case is a question. In the language content example, over 57,000 possible question types may be generated. For example, if the stimulus format is multiple choice (which may be still type 3 in the stimulus format matrix), the media
15 matrix is text for the question and audio for the answer (which may be media types 1 and 2 in the media matrix), and the association matrix is word and definition for question and answer respectively (which may be association types 1 and 2 in the association matrix), then the combined value representing the selection is (3,1,2,1,2). Each combination of values from the various
20 matrices produces one unique stimulus, or question type. All of these possible types are available to any, and all of the content records. Below four different types of question are exemplified.

Examples of Question Types

Figure 10 demonstrates an example of a multiple match stimulus format
25 question from the stimulus format matrix. In this example, there are 5 rows 10.4, each row 10.5 represents any associative element, and any media type such as audio 10.1, graphic 10.2, or text 10.3. The right side 10.6 represents the same associative elements and media types as on the left side 10.5, but in

a random order. The user must match the left hand side 10.5 to the right hand side 10.6 by connecting similar rows using a line 10.7.

Figure 61 represents a true or false type of question. The question part 61.1 is comprised of two parts, the main part 61.2 and the sub part 61.3. The main part may be comprised of any combination of any associative element and media type, such as target text word. The sub part 61.3 may also be comprised of any combination of associative element and media type, such as native audio synonym. If the main part 61.2 and the sub part 61.3 are in fact a correct match, meaning that the synonym is the correct synonym of the word, then the answer 61.5 would be "True" 61.6. If the match were incorrect, then the answer would be "False" 61.7.

Figure 11 graphically depicts a multiple choice question type. The question part 11.1 is comprised of any combination of association element and media type as depicted by text 11.2, audio 11.3, and graphic 11.4. The answer part 11.6 is comprised of 4 choices 11.10, each of which contains the same combination of association element and media type 11.7, 11.8 and 11.9, but which is different from the question part. One of the 4 choices comes from the same content record as the question content and is therefore correct. The other 3 choices are culled from any 3 other content records, and are therefore incorrect. The user must select one of the possible answer choices.

Figure 12 shows a "rote memory" type question. The question part 12.1 may be any combination of association element, and media type 12.2, 12.3 and 12.4. The answer type displays only a "type in" input box. The answer label 12.6 specifies the associative element which the user must type in. For example, if the question is target language word text, and the answer is target synonym, then the user must type in the synonym which corresponds to the question 12.1.

Module Structure

Modules are a convenient means of organising large amounts of related data and functions. Further, modules allow people responsible for the preparation of a course, such as a course co-ordinator, a simplified method customising a course to the exact needs of the user or users. Each aspect of the system is
5 organised into a separate module.

Most modules comprises certain information which is common to the majority of modules. This includes details of record identification number, author, module identification, module name, and module description. These are shown in Figure 13, for example, by 13.1, 13.2, 13.2, 13.4 and 13.5,
10 respectively. This information is used in the management and implementation of each module. For example, when ICC wants to combines specific content modules with specific quiz modules, the individual modules are identified by their ID numbers. The remaining details provide the user with useful information for identifying and using the modules.

15 The remaining data in each of the fields of each module is used to carry out the function of the respective module. For example, in the quiz module it is used to generate a quiz. The modules of the present embodiment will be described below:

The Quiz Module

20 The quiz module provides a means of saving pre-set configurations of the quiz selector (i.e. the way in which questions for a quiz are selected). This allows both a course co-ordinator and a user to easily configure the system to match the requirements of the user. As described above the quiz module shown in figure 13 contains information for the management and implementation of a
25 quiz including: identification number; author; module identification; module name; and module description.

The remaining fields of this module include the following: the number of questions 13.6, which is used to determine how many questions the system

must generate for a specific quiz. In this case, the number of questions 13.6 is 5, however, 5 question types does not necessarily mean a total of 5 questions. The total number of questions also depends on the number of content records, and the QQP (*Quiz Question Pattern*) 13.7. The number of content records is
5 determined by the content module, which will be described in detail further on. If the content module selects words, from lesson 1 for example, and there are 25 words in a lesson, then there would be at least 25 questions in the quiz. This is however is determined by the QQP (*Quiz Question Pattern*) 13.7. There are 3 types of QQP 13.7, type one steps through each question type one
10 at a time, while also stepping through each word in the lesson. Therefore, the first content record would be questioned with the first question type, the second content record with the second question type. After the 5th and last question type, the 6th content word would be questioned with the 1st question type, and so on until all the content words had been questioned. The total
15 questions in this case equals 25.

The 2nd QQP type starts with the first content word being questioned with the first question type, and all the content words being questioned with the same question type until all 25 content words have been questioned. The next question type then starts with the 1st content word again, and continues until
20 all the content words, and all the question types have been used. The total number of questions in this case is 125.

The 3rd QQP is similar to the 2nd, but instead of first going through all the content and then switching question types, each content record is questioned by each question type before changing to the next content record.

25 The media values 13.9 contain a ratio for each of the possible combinations of media types specified for the stimulus or question and the response or answer; in this case, 16. The association values 13.10 contain, as above, a ratio for each of the possible combinations of association types specified for the

stimulus or question and the response or answer; in this case, 49. The stimulus format values 13.8 in the quiz module contain a ratio for each of the possible stimulus formats. In the present embodiment of the invention there exist 57 stimulus formats. Therefore, 57 ratio values between 0 and 1 are stored in this field.

The system will use these values to select 5 question types. The ratios themselves are determined by user performance. The mechanism for selecting question types and generating the ratio values will be described in more detail below.

10 **The Notebook Module**

The notebook module is next described with reference to figure 14. The notebook module is the means by which the visual output format of the system delivered through monitor 1.1 are pre-specific. The common module information is included in fields 14.1 to 14.5.

15 The specific content of the module starts with view type 14.6.

There are several possible view types; two of which are illustrated in figures 20 and 21. The value in the view type field 14.6 determines which view will be used for the selected content. One example of a view is list view. List view may display a native text word on the left side, and the target text word on the right side as shown in Fig. 20. A list of 10 rows may be displayed at one time. Notebook automation allows for each row, to be displayed one at a time in a timed sequence. With each row, the left side may appear first, followed, after a set time, by the right side. Each event in the sequence for each content record, or in this case, each row, is referred to as a step. Step 1 14.7 in the list view example contains the numbers which refer to the association element, the media type, the amount of time before moving to the next step, and so on. Step 2 14.8 contains the information necessary to display the right hand side of the row. In this example of list view there are

only 2 steps, so only 2 step fields will contain data. Some notebook views may contain more steps, hence the extra fields to accommodate them.

Content Module

Figure 15 shows the structure of the content module. The content with which a user wishes to work may be pre-specified either by the user or the course coordinator prior to the commencing of a course. The pre-specified configuration of the content may be saved as a module for subsequent use. The common module information is included in fields 15.1 to 15.5. The specific content of the module starts with SQL (*Structured Query Language*) 15.6. SQL is a standard language for manipulating databases. The SQL statement which is stored in the SQL 15.6 field queries the database and produces a group of content records which can then be used in combination with other modules to produce a specified output. The power of the SQL 15.6 is that it enables the system easily to generate complex content record groups.

An example of the SQL 15.6 queries can be referenced from any of the many published books on SQL. Within the scope of the invention, SQL 15.6 is used to select content records in the following way. Selection criteria 1, 15.7 may contain any of the possible criteria used to select content. Examples include: content within a specified lesson range; content that was erred during a specified date or time range; content that was erred using a specified association type; media type or stimulus format; content which has been deemed important or unimportant by the user; content which has been erred most often; content where the synonym contains the letters "ght", etc. The flexibility of SQL enables the content to be queried in many ways. As long as the data exist to enable a query, then SQL can generate the appropriate content records. The data referred to consist of , for example the date of each error. As long as the date is saved along with the content ID number, then SQL may access those errors made in a specified period.

The mechanism for storing, and modifying data for the purpose of making SQL queries will be described in detail below.

Selection criteria 2, 15.8 is identical to selection criteria 1, 15.7, except that it determines that more than 1 selection criteria must be met in order before the content is selected. For example, content from lesson 1 that has been
5 incorrectly answered in the audio media type. The 2 criteria may also be joined using Boolean logic. For example, lesson 1 content and/or audio media type errors.

Selection criteria 3, 15.9 and so on enable the SQL to select content records
10 that are precisely representative of the users performance. This detailed degree of selection criteria enables the system to customise the content precisely to the users response. The mechanism detailing these features will be described in depth further on.

SQL may also be applied to question type selection, however there are a
15 number of other ways that in certain circumstances prove more effective. These will also be detailed further on.

Course Module

The course module is shown in figure 16. The course module groups quiz
modules and content modules, as well as notebook modules and content
20 modules together. Any number of modules may be grouped together in the course module. The grouped modules may then be run automatically in a specified sequence, and in accordance with the settings of each individual module.

The common module information is included in fields 16.1 to 16.5. Each
25 notebook or quiz module must be associated with a content module. The fields ranging from 16.6 to 16.15 simply list the modules which are contained in the course module. They also specify which content module is associated

to which quiz / notebook module. The number of possible module sets in this case is 20, however, this could be more or less.

The user may select the modules which comprise a course module from a group of available content, quiz, and notebook modules. He may also create
5 any of the content, quiz and notebook modules through a user interface that allows him to enter values for any of the data fields described in figures 13, 14, and 15. For example, in the content module, the user may enter "1" for the lesson number, when that content module is then used, it will cause the system to produce lesson 1 for the quiz or notebook. The mechanism that
10 runs the course first checks the course module to determine the first set of modules to be used. If they happen to be the content module and the quiz module, then by checking the contents of the content module, in this case lesson 1, and the content of the quiz module, it uses those values to generate a quiz, or a series or stimuli for the user. The actual method of stimulus
15 generation is described later with reference to figures 23, 24, and 25.

ICC Module

The ICC (*Intelligent Course Customisation*) module is detailed in figure 17. The ICC module is similar to the course modules. In content and data structure the ICC and course modules are identical. The ICC module may be
20 described as a special instance of the course module. The mechanism that uses and controls the ICC module will be described in detail with reference to figures 43 and 44.

ICC Configuration Module

The ICC (*Intelligent Course Customisation*) configuration module is shown in
25 figure 18. The ICC configuration module contains information that determines how the ICC mechanism works. The purpose of ICC is to customise a course to the specific needs of each user, and the ICC configuration module determines the type of configuration ICC will perform.

The information contained in the ICC configuration module has 4 categories, default 18.22, user set 18.23, actual 18.24, and forecast 18.25. Default 18.22 contains pre-set configurations based on average values that have provided optimised learning speed for previous users. User set 18.23 contains the settings that the user configures themselves. Actual 18.24 contains the values that are actually in place while the system is running, and the forecast 18.25 contains values based on the user performance and calculated with the ICC algorithms to determine an optimised setting for a specified user.

Hours per session 18.1, is the amount of time spent using the system per session. One session is defined as the point where a user starts the system, to when they shut it off. This value will be stored in 4 different locations (records in this case) according to the 4 categories, default, user set, actual, and forecast. The remaining information will also be categorised similarly. Sessions per week 18.2 is the number of sessions per week. Number of weeks 18.3 is the length of the course. End date 18.4 is the specific date that the course ends.

Depth 18.26 is the average number of times each associative element is used in the generation of a quiz. The higher the value of the depth, the "deeper" the content will be learned. Width 18.5 is the number of associative elements used for each content record. For example, if the user only wants to do a course with the word and definition elements of each record, then the width would be 2. Content range 18.6 determines which content will be included in the course. Familiarity 18.7 is the degree to which the user is familiar with the content he has not viewed yet. Content bandwidth 18.8 is the number of content records which may be quizzed in one session that are not new content, and not content which has already been mastered. In other words, content bandwidth refers to the amount of content which is in the process of being learned.

Efficiency - total 18.9 is the ratio of time spent making associations versus total time. Efficiency - quiz 18.10 is the ratio of time spent making associations in the quiz versus total quiz time. Efficiency - notebook 18.11 is the ratio of time spent making associations in the notebook versus total notebook time. Time per association - total 18.12 is the average time to make 1 association. Time per association - quiz 18.13 is the average time to make 1 association in the quiz. Time per association - notebook 18.14 is the average time to make 1 association in the notebook.

The CDR (*Content Difficulty Rating*) is a value assigned to each content record which moves up or down according to user performance. The CDR is explained below with reference to figure 22. CDR- $\alpha\Delta$ 18.15 is the amount that the CDR is incremented when the user answers correctly. CDR- $\beta\Delta$ 18.16 is the amount that the CDR is decreased when the user answers incorrectly. CDR- $\pi\Delta$ 18.17 is the amount that the CDR is increased when the user answers correctly, and has never erred on the specified content record. CDR- $\lambda\Delta$ 18.18 is the average amount of change in the CDR per question for all content records.

Level 1, 18.19 determines whether or not certain features in the ICC mechanism are used or not. Level 2, 18.20 and level 3, 18.21 are similar, but control different features. The features that are controlled by the level settings will be described in detail further on.

Quiz Configuration Module

The structure of the quiz configuration module is shown in figure 19. The quiz configuration module stores the settings which determine how certain aspects of each question will perform during a quiz. The values stored in the quiz configuration module are used during the quiz flow which is described in detail in figure 31.

QDV (1~n) 19.1 (*Quiz difficulty Variance*) is the degree that one question type can change in difficulty according to the configuration settings. For example, in a multiple choice type question the degree of difficulty can be altered by increasing or decreasing the number of options from which the correct answer has to be selected. The QDV is a value that is attributed to each of the stimulus format values in the stimulus format matrix. The QDV sets the difficulty range within each question category. For example, multiple choice type questions may contain from 2 to 6 answer choices. The greater the value of the QDV, the more difficult the question. This is a form of fine tuning, as there is already a wide range of difficulty variance among the available question types. Other question type categories use the value of the QDV in different ways.

Text display 19.2 determines whether the text of the correct answer will be displayed when the user makes an error. Text display time 19.3 determines the length of time to display the text if text display is on. Audio 19.4 determines whether to play the audio of the correct answer if the user answers incorrectly. Error type in 19.5 determines whether the user must type in the correct answer if he answers incorrectly. Error type in maximum 19.6 determines the maximum number of times the user must type in the correct answer if they continue to type incorrectly. Correct type in 19.7 determines whether the user must type in the correct answer after they answer the question correctly. This is usually only used for preliminary learning when the user is very unfamiliar with the content, as a form of reinforcement. Redo question 19.8 determines whether the user must redo the question if they answer incorrectly. Redo question maximum 19.9 determines the maximum number of times the user may redo the question if they continue to answer incorrectly. Error beep 19.10 determines whether the system plays an audio sound when the user makes an error. Correct beep 19.11 determines whether the system makes an audio sound when the user answers correctly.

Error choice content 19.12 determines the type of content that is used to generate the incorrect answer choices for a question. Some examples of the error choice content are content that the user has not used in a previous quiz, content that the user has already mastered, or content that is the same as the content being quizzed on (except that the content record for a particular question may not show as an error selection for that same question, it could however, show in subsequent questions from the same quiz). The error selection content will be described in further detail further on.

Content selection

Figure 22 shows the mechanism for selecting content that the system will then use to generate outputs, perceived by the user, such as stimuli, questions or outputs to the monitor screen. Start content selection at step 22.1 sets the course of events that will select content. Load meta values for content selection at step 22.2 determines which meta values to use from the content meta database 22.6, which in turn is determined by the content type that is loaded at the time the content selection procedure starts. Because the system abstracts the structure of the content from the content itself, the associative elements of one content type may not exist in another content type. In other words, native word text may be field number 1 in one language content; however, field number 1 may be equivalent to "anatomical part" in medical based content. The content meta database 22.6 is similar to a mapping table to ensure content independence so that the invention may be applied to any content that may be categorised into associations.

Next the values of the content module 22.7 are loaded at step 22.3 and prepared to be placed into the appropriate SQL code. For example, if the search criteria is lesson 1, then that value (1) is located in a position dedicated for the lesson criteria. The value (1) is then added to a predefined SQL query. With the criteria value (1) the SQL will load the appropriate content records at

step 22.5 from the content database 22.8. If the value is nil, then lesson was not selected as a search criteria and the lesson SQL is not executed.

In the case of multiple search criteria, then each separate value is loaded into the appropriate SQL query and combined before being executed.

- 5 Once the query result has been loaded, the content fields are in a predetermined order which may be accessed with the mapping data, or content meta data. The query result produces an array of content records which may be sorted, and then accessed for display in a question, or notebook view.

- 10 Therefore, regardless of which associative element is required, the content selection procedure results in all the associative elements for the selected content records being loaded. However, it makes no difference whether the specific associative elements are first loaded into the SQL query so that only the necessary associative elements are prepared for use.

- 15 Certain quiz types require that additional content is loaded apart from the content specified by the content module. This content is, for example, for the purpose of producing the erroneous answer selections for multiple choice type questions. The basic procedure to obtain this content is the same as the content selection described in figure 22. However, the content module criteria is usually different, although this is not necessarily the case. These criteria are
20 as variable as the normal content criteria described in figure 15. However, in all cases of content selection, the SQL allows for highly configurable and customisable selection criteria to be easily implemented. This flexibility allows effective content customisation for each individual user.

- 25 In order for the SQL to generate meaningful content, data must be stored in a user profile, or CDP (*Customised Diagnostic Profile*). This process is usually called tracking. For example, if a user makes an error, then that information must be stored. It may include other information such as the date, etc. This enables specific content records to be queried by date of error in the future.

The extent to which data is tracked, and the simultaneous manipulation and tracking of data enables the results of the tracking process to be immediately made available for the benefit of the user.

Figure 62 shows some of the information that is tracked, and how it is organised. CDP 1, 62.1 tracks the user performance with each of the available question types as identified by 62.2 and 62.3. Total question time 62.4 for each question type is updated after every question. Set time 62.5 is a default timer setting which determines the amount of time allotted to each question type. User time 62.6 is based on the set time 62.5; however it is adjusted according to user performance. The time allotted to each question is thereby dynamically determined by user performance. This will be explained in detail further on. Times used 62.7 is the sum of the number of times a question type has been used. Number correct 62.8 stores the sum of the number of times the user correctly answers each question type.

Content information 62.9 is a database table that tracks user performance relating to each content record. Content ID 62.10 refers to a specific content record in the content database. Times used 62.11 stores the number of times each record is displayed as a quiz and notebook. Total question time 62.12 is the sum of the time each content record is displayed in a question. Total notebook time 62.13 is the sum of the time each content record is displayed in any of the notebook views.

First used date 62.14 is the date that each content record first appears in a quiz. First error date 62.15 is the date that the user first answers a question incorrectly with each content record. Last error date 62.16 stores the date of the last error for each content record. Number correct 62.17 is the sum of correct answers for each content record.

CDR (*Content Difficulty Rating*) 62.18 is a value which is dynamically adjusted according to user performance. It starts at a specified number, and

then may be incremented, or lowered. A low CDR implies the users inability with that specific content record, a high CDR implies familiarity with the content record. Ω -CDR 62.19 is the highest level that a content record may achieve. At the Ω -CDR level, content is assumed to have been learned to the degree (depth) specified by ICC (*Intelligent Course Customisation*) or by the user.

Tag (1~n) 62.20~.21 are flags that may be set by the user to make selected content records unavailable to a course. Conversely, the user may choose only the tagged content for a specified course. The number of tag types enables the user to categorise the content in numerous ways.

Degree of loss 62.22 is the number of times that a content record drops from the Ω -CDR plus the number of questions necessary for it to return to the Ω -CDR.

Error content profile 62.23 tracks information regarding each error. Error number 62.24 tracks the sequential order of the each error. Content ID 62.25 identifies the specific content record that the user answered incorrectly. Question type 62.26 identifies the question type which also includes the values for the stimulus format matrix, media matrix, and association matrix. Quiz ID 62.27 identifies the quiz module associated with the error. Error date 62.28 stores the date the error was made.

As matrices are added or subtracted according to the requirements of the content, the information that is tracked is adjusted to correspond to the content type.

Question Selection

In preparation for a quiz, the system must select the specific question types that will be used. The question selection mechanism will be described referring to figures 23, 24, and 25. The mechanism is based on ratios to

determine a specific value for each of the 3 matrices, stimulus format, media, and association which, when combined, point to a unique question type.

Figure 23 starts at step 23.1 by initialising the values of an [I] index to count the number of selected questions, and a [j] index to specify 1 of the matrices.

5 [j] is then set at step 23.2 to 1 which, in this case is the stimulus format matrix. The stimulus format matrix ratio values for each possible stimulus format is then loaded at step 23.6 as α -question_ratio[j] from the question module 23.5. α -question_ratio[j] contains the ratio for each stimulus format, so they must be parsed into an array as α -question[j][k] at step 23.7 so that
10 each ratio value may be dealt with separately. "k" is an index to identify the ratio for each specific stimulus format in the stimulus format matrix. If j, the matrix index equals the number of matrices used in the specified content type at step 23.8, then the ratio values of each matrix, in this case stimulus format, media, and association, will have been loaded and parsed into the α -
15 question[j][k] array. If this is not so, then the above steps from 23.3 are repeated until all the data has been loaded.

The matrix index is reset to "1" at step 23.9. The ratio_sum is set to "0", as is the "k" index at step 23.10 which will move through each of the ratio values for each of the matrices. A random number is generated "rnd" which ranges
20 from 0 to the sum of the ratios in the matrix specified by [j] 23.11. "k" is incremented by "1" at step 23.12. The value of ratio_sum is then increased by the value of the ratio in the matrix [j] which k is now pointing to at step 23.13. At this point, if the randomly generated number "rnd" is less than, or equal to the sum of ratio values up to the k matrix type value, then the present value of
25 k is selected at step 23.14. If the random number "rnd" is greater than the ratio sum at step 23.14, then k is incremented by 1, and the above steps from 23.12 are repeated until a k value is selected.

So far, the value of the stimulus format matrix has been determined. If k equalled 3, and 3 was the stimulus format for multiple match, then multiple match would be the stimulus format of the first question type.

5 Next, the values of the remaining matrix types must be selected. This process continues in figure 24. 24.1 shows the transition from figure 23 to 24. Choice[j] equals k thereby placing the value of the selected stimulus format as choice [j] at step 24.2. If all the values of the matrices have not yet been selected at step 24.4, then the value of j is checked to determine if it is on the 1st matrix type. When j equals the 1st matrix type, then the rules for the
10 stimulus format k are loaded at step 24.6 from the rule module at step 24.5 as α -rules[l]. After having loaded the rules once, there is no need to repeat the same step again for the same question type, so this step will be skipped when the matrix type is media or association. The rule module and its purpose is explained in detail in figure 26.

15 Figure 26 shows the structure of the rules module. A set of rules exist for each stimulus format matrix type. For example, multiple match has a unique set of rules, as does dictation, and any other category of stimulus format. This is demonstrated by 26.1. The rules consist of either a 0 or a 1 for each possible media type category. If, for example, text question and text answer is
20 set to 1, then it is a possible media type combination. On the other hand, if the rule value is 0, then that particular media type combination is not possible. The same is true for the rules determining the association type possibilities. The necessity of the rules is evident when the total possible combinations of the matrices are calculated. Over 997,000 possibilities exist for some content
25 types. However, this far exceeds the possible number of question types. Most of the combinations yield meaningless matrix combinations that cannot generate questions. Therefore, the rules check to allow only the possible matrix combinations. This is shown by media matrix values 26.2 and 26.3.

After having loaded the appropriate rules, the l index, which points to a specific value in α -rules[l] is set to 0, as is the k index which points to a specific matrix value at step 24.7. The j index is then incremented by one to change to the next matrix type at step 24.9. In this case, this will be the media
5 matrix.

Moving to figure 25, "l" and "k" are incremented by (1) at step 25.2. When the value of the selected rule determined by the l index equals (1) at step 25.3, then the specified matrix type value is permitted. In other words, if media type matrix text question and text answer are presently selected, and the
10 rules determine that that combination is valid within the context of the stimulus format, then text and text is allowed. However, in the case that text and text is not allowed, for example when the stimulus format is dictation, then since the question type must be audio for the question type to be considered dictation, then text and text will be given a ratio value of 0 and it
15 will not be selected. Once each rule is applied to the each possible matrix type, then the system will then determine if there are any values of the matrix type which remain to be selected. If text and text had been the only media type allowed by the quiz module, and within dictation it could not be used, then no questions could be generated, in which case the system starts the
20 procedure again from the beginning. This sequence is detailed in steps 25.4 and 25.5, 25.6 and 25.7. The purpose is to eliminate the chance of impossible question types being selected.

The ratios of the remaining matrix types that have not been ruled out are summed up at step 25.9. Subsequently, the sequence returns to figure 23, and,
25 the procedure to select one of the remaining matrix types is repeated starting from 23.15.

Once the values of each matrix type have been selected at step 24.3, in the above stated fashion, then the values necessary to generate one question type

have been set. The values of each matrix type are then concatenated into Ω -question[i] at step 24.13. One question type has been determined, and the procedure may repeat until the desired number of questions, as determined by the quiz module settings, have been selected.

- 5 The benefit of selecting the question types by ratio is that it provides a high degree of control over the question types over a period of time. At any given time, the result is unknown, but over time the selections will be equivalent to the set ratios. In other words, users, instructors, or content providers may determine 40% audio and 60% text, as well as 100% target word and target
10 definition, may be an optimal setting for some users. The absolute degree of control, and the means of monitoring the results or effects, produces a feed back loop which may be used to continually optimise the settings according to each user. Apart from optimising the settings, specific ratio values may be set to obtained a desired result. The modular structure of the ratio values also
15 provide a means for their distribution and implementation to a class, or any group over a network, or by mail.

Once the question types have been determined for a set quiz, then the values of the various matrices must be used to actually generate the question type they point to. This procedure is detailed in figure 27.

- 20 One specific question type value is first combined with the content from one content record. The first question might display the first content record at step 27.2. From the matrix values, the association elements may be discerned and separated into question type associative element, and answer type associative element at step 27.3. The system then proceeds to set the question part at
25 step 27.4. The value of the stimulus format matrix is used to search a database of preset question structures 27.6. As necessary, the values from the database, in combination with the content record values, set for example the target text word into the text object in the question part position. The same is

repeated for all the media types. For example, audio is then loaded into the audio object, and graphics into the graphic object. Some question types may have more than one text display such as multiple match. The number of rows is therefore set, and the appropriate text content is displayed following the sequence from 27.4 to 27.11.

The answer display is similar to that of the question display. The exception is that for some answer types, content must be retrieved to display the incorrect answer choices. Having done that, the settings such as the QDV (*Question difficulty Variance*) must also be set. If the question type is multiple choice, and the QDV is 3, then 3 answer selections will display. 2 will be incorrect, and 1 will be correct.

Continuing with reference to figure 28, extended question and answer types must be handled. An extended question type may include a sample sentence where a word must be removed and replaced either with a blank, or another word. This procedure only occurs on some question types, and may include other types of modification as well. Letters of a word may be scrambled for example. The same procedure is also used to handle any extended answer types that are required by the specific question types. Once all the interface elements have been prepared, then they are made visible, or, in the case of an output mechanism that is not a display, then the actions which are relevant to the specified output device are activated. For example, audio is played through a speaker, the surface of a Braille machine for blind people configures the letters of a word, depending on the content type, the keys of a music generating electric keyboard may be depressed.

A question has now been selected and displayed; however, courses, and quizzes are composed of a many of questions. The pattern with which they are displayed becomes an issue, and a tool for optimising learning speed. The 3 QQP (*Quiz Question Patterns*) have already been described, now, with

reference to figure 29, the mechanism for applying the different patterns will become apparent.

The QQP may be set or determined by the QQP value in the quiz module, or the settings may be overridden by the ICC (*Intelligent Course Customisation*).

5 Using the settings in the quiz module, the course of events is defined by the QQP value. In other words, when QQP is 1, then the content record is not changed until all the question types have been performed. The next content record is put in place, and again the question types are performed. With a QQP value of 2 the opposite occurs, one question type is kept in place while
10 all the content records selected by the content module pass through it. Finally, a QQP value of 3 yields a pattern where each content record, and question type rotate in synchrony.

Once a quiz starts, the flow of events is repeated for each question. The sequence of events is shown in figure 30. The mechanism described in
15 figures 27 and 28, displays a question at step 30.2, and starts the question timer at step 30.4 with the "set time" value loaded from the question type. While the question is waiting for the user to answer, there are several possible event scenarios. The user may temporarily pause the question at step 30.5 from which he may either continue on with the quiz or stop it altogether.

20 The user also has the option of using hint at step 30.7 which provides assistance in answering the question. The hint option provides only a clue for the purpose of jarring the user's memory. The hint works by showing the user an association element of the correct answer that is not used in the quiz. For example, if the question is multiple choice, and the content is native text word
25 and target audio synonym, then the hint may show the target text definition. This is not the answer; however, it helps the user recall the correct answer.

The user also has the option to answer the question at step 30.9. A question may be answered in many ways depending on the content type. A typical

multiple choice question type based on language content, may be answered with a mouse click on the correct answer button, by pressing a key on the key board, or by saying the answer into a microphone. In each case, the input device is a standard off the shelf product, as well as the mechanism which renders the answer to a form interpretable by the system.

Once a question has been answered, the response is evaluated at step 30.10 by the system to determine if the response is correct or incorrect (Psychometric tests are an exception that do not evaluate answers as correct or incorrect. Instead the responses are simply stored in a database and processed according to rules provided by the test developers.) The answer response mechanism is described in detail in figure 31.

The response time for the question is recorded at step 30.12, along with the evaluation results and other information that will be explained further on. Next, the system may respond to the users answer in one of several ways. An incorrect answer may prompt the system to display the correct answer at step 30.13, and a correct answer may be followed by the next question. When all the questions have been completed, then the quiz may end.

Finally, information regarding the question and the user's response to it are adjusted according to algorithms at step 30.14 that will be explained in detail further on.

If the user does not answer the question within the allotted time at step 30.11, then the system records the time, and determines the appropriate system response at step 30.13, before moving to the next question.

Alternatively, the user may select to stop the quiz at step 30.17. In which case, the quiz ends without going through all the questions.

The system distinguishes between correct and incorrect answers using the mechanism outlined in figure 31. As the question is loaded and prepared for execution, the associative element that will be the answer, is loaded from the

question content record, into a variable: α -correct_answer[i]. The user's selection is stored in user_response and the two values are compared at step 31.4.

5 If the values match, then the response is correct and a correct flag is triggered 31.12. The correct answer counter for the quiz is augmented by 1 (at step 31.11). If the response is incorrect, then an error flag is triggered at step 31.9, and the error counter is augmented by 1 (at step 31.8).

10 The wide variety of questions require exception handling because of the difference in formats. For example, some questions have more than one possible correct answer. In which case, after the α -correct_answer[i] is compared with the user response, and no match is found, then β -correct_answer[i] is compared with the user response 31.5. β -correct_answer[i] corresponds to the alternate correct answer. If more than 2 correct answers are possible for a particular question, then the matching routine may be increased as necessary.

20 In some cases, one question may require more than one answer. For example, a question may require the user to type in different verb tenses for a selected word. Each of the user's inputted responses must be checked separately, hence the i index is augmented at step 31.3, in order to process and check each answer separately. If only one of the responses is incorrect, then the entire question is deemed incorrect at step 31.7.

25 For the multiple match question type format, many content records are tested simultaneously. However, all the responses are checked at the same time, and some may be correct while others are incorrect. The i index allows each record to be processed separately and handled as a separate question (31.6, 31.10, 31.13).

Once the user response has been evaluated, then the course of events branches accordingly. Correct and incorrect responses are processed differently

according to the value of the flag that is triggered in the response evaluation sequence.

CDR (*Content Difficulty Rating*)

A number of calculations are performed on the user data according to their responses. One such calculation involves the CDR (*Content Difficulty Rating*) as depicted in figure 32. The CDR value is a detailed measure of the user's performance with each content record. The CDR enables a high degree of differentiation of levels so a distinct course may be used for each separate level, as well as differing usage frequencies. Content that has never been viewed before starts with a CDR value of null. As the user responds correctly, or incorrectly, the CDR is adjusted upward or downwards accordingly. When a content record reaches the Ω -CDR it is deemed to have been learned to a degree (depth) previously specified by the user or ICC.

The rate at which a content record advances to the Ω -CDR is dependent on 2 factors: the ratio of correct responses by the user, and the CDR- $\alpha\Delta$, CDR- $\beta\Delta$, and the CDR- $\pi\Delta$ values. The CDR- $\alpha\Delta$ is the amount the CDR is incremented when the response is correct. If CDR- $\alpha\Delta$ equals 1.5, then it will take at least 10 correct answers to get to CDR-15. Alternatively, if the CDR- $\alpha\Delta$ is 7.5, then it will only take 2 correct responses reach CDR-15. The rate at which the CDR drops when the response is incorrect also determines the overall speed with which the CDR may reach the Ω -CDR. This is determined by the value of CDR- $\beta\Delta$ which functions as CDR- $\alpha\Delta$ but in reverse. Furthermore, it may be assumed by the system in the following manner, that if the user has made no errors with a specified content record, then he is deemed to have been previously familiar with that content record, and therefore need not spend much time with it. In the case that no errors have been made, then the CDR is incremented by an amount CDR- $\pi\Delta$. This value is usually higher than CDR- $\alpha\Delta$.

The user may set the values of $CDR-\beta\Delta$, $CDR-\alpha\Delta$, and $CDR-\pi\Delta$; however, it is usually more effectively done by the ICC. The mechanism for establishing these and other values via the ICC will be described in detail further on.

Using information previously stored in the user profile 32.2, the system
5 checks if the user has ever incorrectly answered the specified content record before at step 32.3. If yes, and the response is correct, then the CDR is incremented by $CDR-\alpha\Delta$ at step 32.11. The boundary limits are checked at step 32.12 to make sure the CDR value is within an allowable range. If on the other hand the response is incorrect, then the CDR is decreased by $CDR-\beta\Delta$ at
10 step 32.14, and the boundary limits are checked at step 32.15.

When no errors have been made previously, the CDR is incremented by $CDR-\pi\Delta$ at step 32.8. For a first time error, the content record is marked with the error date at step 32.5.

QDR (*Question Difficulty Rating*)

15 The QDR is the value given to the matrices which combine to generate a question type. The QDR value is a ratio from which the question selection procedure generates a specific question type. Generally, ICC (*Intelligent Course Customisation*) generates a question type more frequently as the user answers it incorrectly, and less frequently when the response is correct.
20 However, the adjustment is not limited to one specific question type. The values that are adjusted are the ratios for each matrix that comprise a question. In other words, if the QDR of a question is adjusted upwards, then the component matrix values are adjusted upwards. For example, a multiple match native text word, target text synonym question type would have the
25 component association pair, media type pair, and stimulus format values increased. This will affect not only the one specified question type, but all word, synonym association pairs, native text, target text media pairs, and the multiple choice stimulus format. Each would be selected proportionally more

frequently. When a question is answered correctly, then inversely, each component value is decreased.

The resultant QDR ratio values are customised to the user according to the basic elements that comprise a question. Over time, if the user is weak with target audio questions, but relatively strong with the word associative element, then, although they may be combined in one question type, their ratio values may change differently reflecting the users relative ability with each component.

The use and application of such accurate and detailed customisation will be elaborated in detail further on.

The actual QDR mechanism uses the matrix position values of the currently loaded question type and, according to the response evaluation flag at step 33.5, increases or decreases the ratio values at step 33.6 and 33.9. A boundary limit may be set to keep the values within a reasonable range at step 33.7 and 33.8. When this procedure has been completed for each of the matrix types at step 33.12, then the resulting changes in the ratios update the ICC Quiz module 33.14. The value by which the ratios are increased and decreased, QDR- $\alpha\Delta$ and QDR- $b\Delta$ respectively may be adjusted by the user, content provider, or calculated by ICC. These values determine the speed by which the QDR responds to the users performance. The use of the ICC Quiz module will be described with the ICC explanation.

Timer Calculation

The timer, similar to the CDR and QDR, may be adjusted according to user performance. As a user correctly or incorrectly answers a question type, the timer for that specific question type is increased or decreased accordingly by a value timer- $\alpha\Delta$ at step 34.7 or timer- $\beta\Delta$ at step 34.3. A boundary limit check 34.4, 34.8, keeps the timer within a reasonable range. The resulting timer value is stored in the question type data table referred to in figure 22a.

The dynamic customisation of the timer for each question type keeps the user moving at a pace that precisely matches their ability.

Error Response

5 The system may respond to the user's answers, be they correct or incorrect, in numerous ways. The possible responses may be configured by the user, or set dynamically by the ICC. Some of the possible response types when the user answers incorrectly include displaying the correct answer before moving to the next question. The amount of time the text is displayed may also be set. Normally, content the user is very unfamiliar with would be displayed longer, 10 and familiar content would be displayed no longer than absolutely necessary to jar the memory. Time is an essential element, as will be explained in detail further on, so these settings are important. Too fast and the user doesn't have time to recognise or assimilate the correct answer, too slow and precious time is wasted that could be better spent elsewhere.

15 Apart from displaying text, it is also possible to display audio, or any other media type. The specifics of the content type, and the output device will determine the exact range of possible error responses.

Further error response types include typing the correct answer in, or even redoing the question until getting it correct. Any of these may be combined 20 with one or all of the response types. The strategic application of the error response will be described further on, for now it's mechanism will be elaborated as shown in figures 35, 36, and 37.

The values stored in the quiz configuration module 19 determines which set of response types will be executed. Branching by the flag value at step 35.4 set 25 by the answer evaluation, allows the error responses to be handled differently from the correct ones. If the configuration is set to play an error beep, than it will play at step 36.3, If the text display is on, it will display at step 36.7.

Similarly for the audio at step 36.10. A timer determines the length of the display at step 36.12.

The type in feature is checked at step 35.14, and if it is on, will continue until the correct answer is typed in at step 37.5, or a limit is reached at step 37.2.

5 Similarly, the redo question feature is checked at step 37.8, and will continue until the correct answer is selected at step 37.11, or a preset limit is reached at step 37.9. When either the type in, or the redo question is incorrect, the correct answer is redisplayed before the user attempts another answer at step 37.14, and at step 36.2.

10 **ICC (*Intelligent Course Customisation*)**

The main purpose of the ICC is to generate output stimuli, often in the form of questions, of an appropriate difficulty for each particular content record. In other words, if a user is unfamiliar with a content record, then ICC is arranged to produce a stimulus or question which is of a difficulty just within or slightly beyond the ability of the user for that particular content. This is achieved by increasing or decreasing the difficulty of responding correctly to a particular stimulus, or question, as appropriate by adjusting the difficulty of the association which the user is required to make.

Content width is defined as the number of associative elements of a content record which may be applied in a course. For example, if a user selects to use only 3 associative elements, say word, definition, and synonym, then the width is 3. The possible number of associations is 9 according to the formula:

$$\text{associations} = (\text{associative elements})^2$$

Referring to figure 38, two associative elements 38.5 comprise 1 association 38.6. As the number of associative elements increases from 2 shown in fig. 38a to 6 shown in fig 38e, the possible associations also increases.

Depth is the number of times each associative element in a content record (as determined by the width) is used to make an association. For example, if the depth is 4 for a width of 2 as shown in fig. 38a, then 4 associations must be made with each content record to achieve the specified depth. On the other hand, if the width is 5, as shown in fig. 38d then 10 associations with each content record will be necessary to achieve the same depth. The number of necessary associations is calculated according to the formula:

$$\text{associations} = (\text{depth} * \text{width}) / 2$$

Inversely, the depth may also be calculated after a course by altering the formula:

$$\text{depth} = (\text{associations} * 2) / \text{width}$$

The significance of these relationships can be seen from the potential for forecasting the necessary number of associations necessary to learn a content record to a specified degree; or conversely, the ability to set the degree to which a content record is learned.

Forecasting the necessary number of associations necessary to learn a content record to a specified degree is impossible on a content record by record basis; however, using averages from a group of content records, can provide statistical accuracy. In other words, it can not be known how many associations will be necessary to learn a specific content record, but, the average number for all records can be known.

At this point, depth is an arbitrary number which simply determines the number of associations the system will generate for each content record; however, it must be connected to the users actual performance for it to be meaningful. Specifically, depth should correspond to the actual degree to which a content record is learned. This is accomplished using the "loss rate" and the "degree of loss".

The loss rate is the ratio of content records which, having reached the Ω -CDR, then drop down again into lower CDR levels. If the content has not been learned well, then a high loss rate is expected. The degree of loss is a measure of the number of associations necessary to return the CDR back to the Ω -CDR. The higher the degree of loss, the lower the degree of learning that specific content record is deemed to have reached upon arriving to the Ω -CDR. By determining an acceptable loss rate range, the depth may be calculated accordingly. Adjusting the depth up or down until a desired loss rate is achieved. The loss rate is calculated as follows:

$$\text{loss rate} = (\text{total loss records} / \text{total } \Omega\text{-CDR content})$$

The depth determines the average number of associations a content record will need to reach the Ω -CDR. Therefore, depth is inversely related to loss rate. As the depth increases, the loss rate decreases, and vice versa. By setting the acceptable loss rate range, the user may indirectly set the depth. The user may determine that they only have time to master the content to a loss rate of 15%, or they may determine that the time required to achieve a 0% loss rate is better spent learning more content to a 15% loss rate level.

The ICC mechanism for adjusting the depth is similar to the QDR, CDR, and timer mechanisms. It includes depth- $\alpha\Delta$ and depth- $\beta\Delta$. Depth- $\alpha\Delta$ is the value by which the depth is incremented if the loss rate is above the set range, and depth- $\beta\Delta$ is the value by which it is decreased when the loss rate is below the set range.

Once the depth has been determined, the mechanism for implementing its value is shown in figure 32. The CDR- $\pi\Delta$, CDR- $\alpha\Delta$, and CDR- $\beta\Delta$ values determine the actual number of associations necessary to reach the Ω -CDR. The users performance will also effect the CDR flow; however, when the average user score is included into the calculation, the precise average depth

may be set. The formula to determine the average amount that the CDR changes after each question, CDR- λ , is as follows:

$$\text{CDR-}\lambda = \{(\text{total errors} * \text{CDR-}\beta\Delta) + ([\text{total correct} - \text{total 1}^{\text{st}} \text{ error}] \text{CDR-}\alpha\Delta) + ([\text{total no error} + \text{total 1}^{\text{st}} \text{ error}] \text{CDR-}\pi\Delta)\} / \text{total number of questions}$$

- 5 In the above formula, the CDR- λ is the rate at which the CDR changes, or flows, per question for both Ω -CDR and μ -CDR (μ -CDR is the content that the user has started to study, but has not yet mastered). The depth and loss rate must be measured from the Ω -CDR content, as the μ -CDR, by definition, only provides data on part of the entire process. μ -CDR is content that is still
10 being learned, as opposed to content that has been learned. Therefore, the formula is further refined to:

$$\Omega\text{-CDR-}\lambda = (\text{total number of questions to reach } \Omega\text{-CDR} + \text{total loss}) / \text{number of } \Omega\text{-CDR content records}$$

- Figure 39 shows a typical content flow pattern over the duration of a course.
15 39.1 is the starting point. 39.2 is the number of content records axis, and 39.3, the CDR value axis. 39.4 shows all the content to equal CDR 0, which is how it is set at the beginning, before any questions have been answered. After a period of time, the user answers questions, and the CDR is adjusted accordingly. The content records the user has had the most difficulty with,
20 have the lowest CDR 39.6. The higher CDR of 39.8 show a better understanding of that content, and the content that the user masters is grouped as Ω -CDR 39.9. New content that the user has not yet started yet also shows a corresponding drop 39.7 as the course progresses.

- The course continues, and more content is spread through the different CDR
25 categories according to user performance. The remaining content that has not yet been started continues to drop 39.12, and more content reaches the goal of Ω -CDR 39.14.

Lastly, when the course has been finished 39.15, all the content records have been learned, and therefore all have a value of Ω -CDR 39.17.

Generally, there are 3 states to the flow of content. The first is new content (CDR-null). The 2nd is content that is being learned (CDR-[-n] to CDR-n,
5 these refer to all the CDR levels between CDR-null and Ω -CDR, collectively referred to as μ -CDR). The 3rd state is learned content (Ω -CDR). The amount of content in each stage is determined by different factors. The amount of CDR-null content, or new content, is determined at the beginning of a course when the user, or instructor selects the desired content. As the course
10 proceeds, the amount of new content diminishes until there is no more left and the course is finished. The amount of Ω -CDR, or learned content starts off as 0, and increases as content is learned, until eventually, all the content that was initially new content, becomes learned content. The rate at which content moves from the 1st state to the last state is called the content flow, or CDR- λ .

15 Finally, the 2nd state, which is a measure of the amount of content being learned at one time, is called the content bandwidth. Content bandwidth may differ from user to user, according to the length of the learning session, or the users aptitude. Content bandwidth is also a factor in the users overall learning speed, hence its setting is significant.

20 At the extreme, content bandwidth may be set to 1, or to the total number of content records. For example, at a bandwidth of 1, the user does a series of questions relating to the one content record, when new content is introduced. Then, once the user completes the introduction sequence, he continues to do exercises that enable him to view the content in a variety of question type
25 contexts as set by the content width. When the user answers enough questions correctly regarding the one content record, as determined by the CDR- $\alpha\Delta$, CDR- $\beta\Delta$, and CDR- $\pi\Delta$, then the CDR value will reach Ω -CDR.

Obviously, a content bandwidth of 1 is not an effective value when the object is long term retention. The reason being that when a user is faced with many questions all using the same content record, then they will quickly retain the information in short term memory, and, after the initial few questions, will usually answer the remaining ones correctly. However, answering a large number of questions correctly does not necessarily mean that the content has been learned, or will be remembered in the future.

On the other extreme, if the content bandwidth is set to the limit (the total number of content selected by the user) and the user has selected 1,500 content records, then they are likely to be overwhelmed.

On the other hand, if the content bandwidth is set to a number that the user is capable of handling in one session, then long term retention is augmented. In other words, if the user handles content in units which are suitable to their aptitude, then learning may be optimised.

The algorithms the system uses to determine the appropriate content bandwidth is as follows:

$$\text{content bandwidth} = (\text{session time}) / (\text{time per association}) / (\text{frequency per session})$$

Where session time is the amount of time the user spends in one session, time per association is the time it takes the user to answer a question, and frequency per session is the number of times a question will appear using one content record in one session. Content bandwidth is therefore a function of 3 factors. 1. The session time which is set before each session by the user, or determined as the average session time of all sessions. 2. The time per association, which is the average time for all questions previously answered by the user. 3. The frequency, which is the number of times a content record is viewed in one session. Clearly, the optimum content bandwidth 1, likely to vary between sessions. Therefore, the bandwidth is in the present

embodiment recalculated prior to each session by the ICC. However, depending on the application this calculation may be performed more or less frequently.

Once the content bandwidth has been determined, then a similar calculation
5 may be made to calculate the number of new words that are introduced at one time. If the content bandwidth is 250 content records, and the maximum is also 250, then no new content may be added. If the amount of new content to be added is 25 records, then only when the content band width drops below 225 will new content be added. The content bandwidth decreases as the user
10 answers questions correctly and the content records reach Ω -CDR. The number of new content records to be added is calculated using the following algorithm:

New content unit = (session time) / (time per new content association) /
(frequency per new content session)

15 The session time is set by the user. Time per new content association is the average time per question for new content (new content questions generally take longer than the average time for all question). Frequency per new content is the number of times each new content record will be questioned. The new content unit calculation therefore determines when new content is
20 added, and how much to add.

The Ω -CDR content is used only once there are enough records to fill an entire session. This allows the user to learn some content first before the review procedure starts. The following algorithm determines when the review procedure may start:

25 Ω -CDR content unit minimum = (session time) / (time per Ω -CDR association) / (frequency per Ω -CDR record)

A detailed explanation of when Ω -CDR content is used is shown in figures 41 and 42. The same figures also details the 3 states that comprise the flow of content, and specifies their relationships

Frequency is an essential feature of the preceding content bandwidth algorithms. It may be predefined by the user, instructor, or content provider;
5 however it may only be manipulated dynamically by the ICC (*Intelligent Course Customisation*). The frequency may be set for each value of the CDR. In this case, the frequency refers specifically to the number of different question types set for a quiz module. When the QQP is not 1, then the
10 number of times a content record is viewed in each quiz module is equal to the number of question types. This allows for frequency to vary according to user performance with each content record. Therefore, the change in frequency is determined by the $CDR-\alpha\Delta$, $CDR-\beta\Delta$, and $CDR-\pi\Delta$ values. The frequency values for each CDR level are determined according the following algorithm:
15 if ($CDR(n)$ score) > (score minimum), then increase frequency $CDR(n-1)$ by ($frequency-\alpha\Delta$)

Otherwise decrease frequency $CDR(n-1)$ by $frequency-\beta\Delta$.

The frequency of each CDR level is thereby increased or decreased according to the users score on the next CDR level content. If the score is too low, then
20 the system determines that the frequency is too low and should be increased. Alternatively, if the score is too high, then the frequency of the previous CDR level is determined to be too high, and is therefore decreased. Boundary limits determine the minimum and maximum frequencies for each CDR level.

Frequency on a larger scale is determined by the sequence and order of the 3
25 content flow states. These will be detailed further on with figures 41 and 42.

The general relationship that the system strives to maintain between CDR, QDR, familiarity, time per question, depth, and frequency is depicted in figure 40.

As any of the factors move up, their value increases as indicated at 40.1, and decrease as they move down as indicated at 40.7. As the CDR 40.2, increases the question difficulty QDR 40.3 also increases, as well as the users familiarity 40.4 with the content record. At the same time, the time allotted to each question 40.5, decreases as does the remaining number of questions for that content record (depth) 40.6 and the frequency that content record is viewed 40.8 also decreases.

More important than the actual relationship between the above mentioned factors is the fact that the system is able to manipulate each of the factors on a content record by record basis. The actual relationships between these factors may be altered to suit the particular needs of the content, or of the user by defining the values of the parameters by which they operate.

For example, the CDR may be set to increase when the user answers incorrectly by setting the $CDR-\beta\Delta$ to a positive rather than a negative number. The relationship between CDR and the remaining factors is thus reversed, showing that the actual relationships may be manipulated by setting the parameter values in different ways.

The mechanism by which the ICC (*Intelligent Course Customisation*) is executed to implement a course is shown in figures 41 and 42. The purpose of the routine is to generate a customised course based on the users performance. To this point, the means of collecting user information and transforming it into questions has been described. The ICC mechanism uses the results of the CDP (*Customised Diagnostic Profile*) to actually generate a course. The ICC determines what content to use, and how it should be used in a course. This includes when to add new content, when to work on learning content, when to review learned content, as well as the many other factors that have been explained previously.

Referring to figure 41, the system starts with CDR-null, or new content 41.2. When the system checks for new content it loads the μ -CDR content 41.5 to see if it is under the content bandwidth by an amount equal to the new content unit 41.9. If it is determined that new content may be added, then an amount
5 of new content equivalent to the new content unit is prepared 41.8, 41.10. Before setting the notebook and quiz configurations 41.14 and 41.16, the system verifies that the user has not finished all the new content and that there remains some new content records 41.11. If there is new content, then a course designed especially for new content is executed. The quiz and
10 notebook modules may be configured by the user, instructor, or content provider.

Generally, a module for new content contains basic question types. The settings for error response are set for a thorough revision to give the user the opportunity to learn. For example if the user makes an error the error
15 response may include a textual and audio display of the correct answer, followed by the user having to then type in the correct answer, and then redoing the question again. The QQP (*Quiz Question Pattern*) might be set so each content record is viewed through each question type, before moving to the next record. However, the actual configuration could be set many
20 different ways according to the overall purpose of the course, and to the content type.

When the notebook and quiz course is finished, the ICC may prepare a similar course for the errors the user made 42.3. In this case, the error content module is loaded, along with the error content notebook and quiz modules. The error
25 content is always dynamically generated with an SQL that, for example, searches for errors made in that session. As was described previously, the new content unit is calculated to occupy a full session.

When the user starts another session, the content bandwidth would be such that no new content would be added 41.9. Instead, the ICC systematically goes through each CDR content module from the lowest CDR to Ω -CDR, with the corresponding notebook and quiz modules. The content bandwidth is
5 set such that the user will finish all the μ -CDR content in one session, including redoing error content.

In another future session, when the Ω -CDR contains enough content, the Ω -CDR modules will be executed in a similar way to the previous modules, except that the content of the modules will be different. At any time, as set by
10 the configuration module, a series of review modules may be triggered at the end of a session 42.5.

The ICC mechanism rotates through the various modules both systematically, and according to the flow of content from CDR-null to Ω -CDR. The actual combination of modules is therefore dependent on the users performance and
15 hence, different for each user. The ICC mechanism performs the task of orchestrating and co-ordinating the various configurable elements into a full course that takes the user from their first session to their last.

Second Embodiment

The second embodiment, in general terms, fulfils the same function as
20 described in the first embodiment. However, in this embodiment, the training process may be partially or fully automated to increase the effectiveness of each training session.

Automated learning spans the entire spectrum from manual, semi automated, to full automation. There are several different types of automation as follows:
25 content selection, Notebook, quiz, timer, configuration, and course. Content selection automation refers to the selection of content whereby the user plays no role. As soon as the overall goals of the course are specified, i.e. content type, content range, width, depth, et, then ICC automatically selects content so

as to optimise content flow (the rate at which each content record moves from CDR 1 to Ω -CDR). This involves determining the rate of new content to be added, the total amount of content between CDR 1 & Ω -CDR at any given time, the frequency for each CDR level, and much more.

5 Notebook automation, as opposed to manual notebook, refers to the ability of the system to display information on each content record in a specified manner, and order, in a timed sequence. In other words, rather than simply selecting to view say a word and its definition for a set of content, and then pressing a button to view the next content record, automation may firstly
10 choose the association type according to the users performance on each content record. For example, with one content record the user may have erred many times on the synonym and definition association type, on the next it may be antonym and word. The associations would be selected and displayed accordingly.

15 Furthermore, automated notebook specifies the order that an association is displayed, and it's timing. For example, a word may be displayed for 5 seconds before the audio of it's definition is played, and then after a lapse of 6 seconds, the text of the definition could be displayed for 5 seconds before moving to the next content record. This feature has yielded significant results
20 for focusing concentration and maximising the APH.

Quiz automation refers to the selection of question types that are relevant to the content type, or to the skill type the user may want to focus on.

Automated timer refers to the automatic setting of the timer according to user performance.

25 Automatic configuration is the ability of ICC to automatically configure any part of the system that the user does not configure, using the user profile and the ICC algorithms to determine the most appropriate setting for a particular user.

Lastly, automated course is the automation, or combination of all the above features. In other words, all aspects of the system are configured automatically to the specifics of each individual user. The goals for the optimised configurations may be to maximise learning speed, or for detailed
5 diagnostic measurements depending on the purpose set at the beginning of a course (usually by an expert, or by loading a module).

Third Embodiment

This embodiment, in general terms fulfils the same functions as described in the first and second embodiments. However, in the present embodiment the
10 system is adapted to train a user to develop a manual or sensory skill, for example manual dexterity and colour recognition, respectively. As described above, prior to being used to develop knowledge and/or skills, the system of the present invention may be used to ascertain the current level of skills or knowledge possessed by the user. Therefore, the present embodiment of the
15 invention may be used to diagnose existing problems or deficiencies of the user in areas of manual or sensory skills; for example colour-blindness. Therefore in the present embodiment the system is adapted to operate with content type other than language or knowledge, or other similar content type, as described in the first embodiment.

20 Such skill based content is computer generated stimulus that does not represent words or language. For example, exercises designed to improve colour perception, or timing perception. In these cases, the stimulus is a colour, or series of colours, a sound, or series of sounds.

The generation of content by the computer is different for each content type,
25 and uses methods standard to the industry. However, the process of selecting the content to be generated is the same for all content types. Specifically, with each skill type, a skill is targeted and stimuli, or questions are generated in such a way that is easily responded to by the user, and then gradually

becomes more difficult, until it is beyond the ability of the user to respond to them meaningfully. The system changes the values of the parameters used to generate the content, systematically and in a linear sequence, one by one, within a range that is determined by the user's responses.

5 Visual Perception

Figure 43 shows 1 colour displayed in the question position 43.1. The colour itself is generated by a combination of numeric values for red, green, and blue that compose a colour. The answer choices the 1st 43.2, the 2nd 43.3, et, are also generated numerically. The user may be asked to find the most similar /
10 dissimilar colour, hue, saturation, or brightness.

The output stimulus start with distinctly different answers making them obvious to most users. They gradually become less distinct, and hence more difficult, until the user is unable to make meaningful choices. One example of a question involves choosing the most similar hue from a series of colours.
15 For the first question, each of the answer choice colours are separated by a hue of 50°. If the question colour is blue, then 1 answer choice would also be blue, but of a slightly different hue. The other answer choice colours would be different from the correct answer by a hue of at least 50°. A hue difference of 50° is obvious to most users according to previously conducted
20 experiments. However, if a user has difficulty with 50°, the difference may easily be increased to a point that the user can easily distinguish the correct answer.

The saturation and brightness of the colours may remain constant in this case, while the questions become increasingly difficult as the hue difference of the
25 answers are cut 5° per question. 5° is an arbitrary amount, and may be set by the user, or by the ICC. Eventually the user will reach the limits of their colour perception. If they are able to perceive hue differences of just 5°, then the differences may be set to 1°.

Once the user reaches his limit, and is no longer able to make meaningful selections, the value of the question colour hue is systematically incremented or decreased by a pre-set amount, and another series of questions start, again moving from easy to difficult by hue differences of 5° .

- 5 When the user reaches his limit, and is no longer able to make meaningful selections, the value of the question colour saturation is systematically incremented or decreased by a preset amount, and another series of questions start, again moving from easy to difficult by hue differences of 5° .

- 10 Further, when the user reaches their perceptive limit within each series of questions, spanning the full range of saturation values, then the brightness may be changed. Changing the brightness systematically by pre-set amounts, and then starting the entire range of questions again at the first saturation value, and the first hue value.

- 15 The different qualities of colour: hue, saturation, and brightness are thus combined to generate a wide variety of questions (over 16,000,000 possible combinations). The user may then be led through questions that target brightness perception in a similar manner. The ICC will generate questions that systematically move through all the colour qualities testing for brightness perception. The same process is again repeated when saturation is targeted.

- 20 In each case, the parameter values of the user's ability are saved in the CDP (*Customised Diagnostic Profile*). Each context (in this case the values of the hue, saturation, and brightness) contains the degree to which the user was able to answer meaningfully for each of the target qualities (hue, brightness, and saturation). Figure 55 shows the data structure of the user profile for
25 computer generated colour content. The hue 55.1, saturation 55.4, and brightness 55.7 for each series of questions is recorded at the point the user's responses become meaningless. The hue- Δ 55.2, saturation- Δ 55.5, and

brightness- Δ 55.8 show the degree each factor, or colour quality is changed each sequence.

The hue sequence 55.3, saturation sequence 55.6, and the brightness sequence 55.9 show the order in which the questions are systematically generated. For example, a sequence of hue 1, saturation 2, and brightness 3 means that the hue of the question is rotated at a degree determined by hue- Δ after each question series. Once all the hue values have been used, then the saturation is changed by a degree determined by saturation- Δ , and once again the hue is changed for each question series until all the saturation values have been used. Finally the brightness is changed by a degree determined by the brightness- Δ , and each saturation value and hue is used until all the values have been systematically generated as a question.

The *question and answer* difference is the degree to which the target colour quality (i.e. hue) differs between the question and the answer. The *answer and answer* difference is the degree to which the target colour quality of the answers differ. The values recorded for these 2 difference quotients is the degree where the user's responses become meaningless. These values mark the limits of the users colour perception ability within a defined context. The Question type 57.12, stores which question type was used. Examples of different question types are shown in figures 43 to 54. Target quality 57.13, stores the quality of colour that is the focus of the question.

Systematically going through all the possible combinations until the users responses become meaningless, generates a detailed user profile that will be used in future courses. Future courses use the user profile to generate questions focused on the areas just before the user starts having difficulty, and slowly moves in small increments beyond their ability. This may be repeated any number of times for each series of questions (as set by the user or ICC) before moving to the next series of questions.

The algorithms that determine the RGB values for the colours are as follows:

Hue is a value between 0 and 360. Each hue value has a corresponding RGB value which is stored in a data table shown in figure 63. The hue value 63.1 has a value for red 63.2, a value for green 63.3, and a value for blue 63.4.

5 Each hue is set at saturation and brightness values of 100%.

Brightness may have a value of between 0 and 100%. The brightness value is then multiplied by the RGB values derived from the hue table shown in figure 63.

Red = Red * Brightness

10 Green = Green * Brightness

Blue = Blue * Brightness

Saturation values are between 0 and 100%. Saturation is the degree to which a colour is pure and not mixed. A pure colour may be comprised of two values such as red and green, and still have a saturation of 100%. When
15 saturation drops, the value of the two lowest RGB values rise according to the following algorithm.

Highest RGB = Highest RGB

Low RGB ₁ = (Highest RGB - low RGB ₁) (1 - saturation) + (low RGB ₁)

Low RGB ₂ = (Highest RGB - low RGB ₂) (1 - saturation) + (low RGB ₂)

20 With the RGB values, colours and hence questions may be generated.

Alternatively, the numerical values of red, green, and blue may be displayed instead of the colour. The user would then match the colour to the RGB values.

Figure 44 is another example of a question type comprised of colour content.

25 The structure of the content and the content generation do not change, only the display of the content. The question 44.1 is a colour and is comprised of three values, red, green, blue. The answers 44.2, 44.5, et, are each comprised

of three colours, 44.2, 44.3, and 44.4. The user must select which combination of colour will create the question colour when combined.

Figure 45 shows a question type that is opposite to figure 44. The question colour is comprised of three colours, 45.1, 45.2, and 45.3. The user must
5 select from the answers 45.4, 45.5, et, the colour that is the combination of the question colours.

Figure 46 is another question type for colour content. The question is comprised of 4 separate colours 46.1, 46.2, 46.3, 46.4, which may be combined in a logical sequence. The sequence may be anything, for example,
10 they may range from light to dark. By dragging the colours into 46.5, the user may arrange them into a sequence from dark to light. Any of the qualities of colour may be used as the condition for determining the sequence.

Figure 47 is an example of a question type that tests the user perception of colour combinations. One colour pair 47.1 is comprised of 2 colours 47.2 and
15 47.3. There are 4 colour pairs each with a different relationship. The user may be asked to select the highest contrast, or the smallest difference in saturation, etc.

Figure 48 is similar to figure 47, except that there exists a question pair 48.1. The user must find the most similarly contrasted pair, or the pair with a
20 similar saturation, etc.

There are a large number of possible question types; however, in each case, the method of systematically generating content that starts at a level the user is able to answer, and gradually becomes too difficult for the user to answer, enables the program to develop a detailed profile of the users ability. The
25 structure of the profile data is similar in most cases, and is easily modified to suit the needs of the content type and the question type. The ICC configurations determine how thoroughly each colour quality is tested and may easily be adjusted. When an answer is incorrect, the error response

setting will determine whether or not to display the correct answer, and for how long. The ideal course structure for colour may be different from language based content; however, by merely changing the ICC parameters, an ideal course structure for colour may be developed.

5 **Pattern Recognition**

Figure 49 shows a basic example of a pattern recognition problem. The three circles represented by 49.1, 49.2, 49.3, illustrate the change in state of one circle over a period of time. The first circle 49.1 is black representing an off state. The next circle 49.2 is white representing an on state. Finally, the last circle 49.3 is black, and therefore represents the off state again. The duration which the on state lasts can be the object of the question. As in previous question types, these may be multiple choice, where the user would select the correct answer out of many. The user may also type in the correct length of time. The degree that the user must be correct determines the difficulty of the question. If all the answers are different by over 5 seconds then that would be easier than if they were only different by 0.05 seconds. The degree of the users accuracy may be measured.

Figure 50 shows the same basic features as figure 49; however, one new dimension of complexity has been added. The state of the circles may continue switching on and off for an indefinite number of times. 50.4 shows the circle on again, and then 50.5 off again. The duration of each on state may vary, making the question inherently more complex and difficult. When the states continue changing back and forth numerous times, the exact number of times may be the object of a question.

Further, figure 51 shows that the duration of the off state 51.2 is now the object of the question. The first state 51.1 remains constant. Figure 52 shows an added complexity, both the on state 52.1, and the off state 52.2, may have variable durations throughout a set number of state changes.

Figure 53 introduces a further type of state 53.4. Now, instead of just on/off, there exists 3 possible states, black, white, and grey. The number of possible states may be further increased. However, the user's ability to answer questions meaningfully over a set number of states will in practice limit this number.

Each of the above mentioned question parameters, may be merged into many question types of varying difficulty. Each question, similar to the colour content, targets one specific quality, and tests the users ability to respond correctly as that particular quality becomes more and more difficult to perceive. For example, at a set number of states, and state changes, with only one state whose duration changes, can the user select the pattern that is the same. If yes, then the next question becomes more difficult, by decreasing the difference between the duration of the correct answer and the incorrect answers. At some point the user will not be able to perceive the difference, and thus the answers will be meaningless.

When the user starts to make errors on a particular target quality, then the context, or the other qualities are changed systematically by a set amount until each quality is tested in a multitude of different contexts. A highly detailed profile of the user is thus generated.

Figure 60 shows the data structure of the user profile that stores the user's results, and enables the system to generate question types of appropriate complexity. As with the colour content, when a user starts to make errors in a particular series of questions, the question settings at that point is stored in the profile. The user then starts another series of questions. The data that is stored includes the number of states 60.1, and the combination type 60.2. The combination type determines the order the states are displayed when there are more than two. Number of flashes 60.3 is the total number of times all the states are displayed in a question. State 1 duration 60.4 is the length of time

the state is displayed for. State 1 pattern 60.5 refers to any one of a group of preset patterns that determine how the duration of state 1 changes in the question. For example, the duration may slowly become shorter, or faster by a set amount. The patterns are mathematically generated and may be as
5 complex or simple as necessary. The above information is also applied to the 2nd state. As more states are added, the same information is again recorded individually.

Q&A # of flashes difference 60.10a, is the difference in the number of flashes between the question and the answers. The Q&A combinations difference
10 60.11 is the difference in the combination type between the question and the answers. All the Q&A differences 60.10a to 60.16 store the amount of difference between the question, and the answers. When targeting one quality, usually that quality difference is above 0, and all other quality differences are 0.

15 Similarly, the A&A differences 60.17 to 60.23, store the degree that the possible answers differ from each other. The smaller the difference, the more difficult the question. Question type 60.24 is the question type used for the particular series of questions that the user is answering. Target quality 60.25 is the quality that the question is targeting.

20 An example of a different question type that uses the same basic user profile is shown in figure 54. A string of circles with set positions change state one at a time according to a constant, or varying time. Each state changes from black to white, and then back to black again as the next circle is reached. At certain points, 54.5, 54.6, and 54.7 the state changes to grey instead of white.
25 At those points, the user must press a key on the key board. The degree of accuracy may then be measured and captured. As one cycle finishes, it may repeat itself as long as necessary.

At any point, the flashes may be replaced with sound, or both may be used simultaneously. Both may also be removed, leaving the user to continue pressing the keys in time, with no clues. The user's accuracy may again be measured and recorded.

- 5 Numerous drills and questions may be generated for a wide variety of stimulus formats. These may include, but are not limited to content involving shapes (including their geometry and size), distance motion, quantities, estimations, etc. The mechanism of generating the question types, and using a detailed user profile to dynamically create a course customised to each
10 individual user remains unchanged.

Fourth Embodiment

- This embodiment in general terms fulfils the same functions as described with respect to previous embodiments. However, in this embodiment the system is arranged to analyse the profiles of a group of users as opposed to a single
15 user. This is achieved through the use of an Education Decision Support System (EDSS), as will now be described.

EDSS (Education Decision Support System)

- Using network connections, the user profiles (CDP) may be centrally stored in memory, and accessed by the EDSS. The EDSS determines the relative
20 strengths and weaknesses of each user and provides instructors, or program supervisors with graphs representing the findings. The data that the EDSS produces includes: a list of the users sorted in order of learning rate. The learning rate is $CDR-\lambda$ which is calculated by the ICC. Alternatively the EDSS may produce a list of users ordered by the amount of content remaining
25 in CDR-null. This would show which user was closest to finishing a course, and what content remains still to be learned.

Any of the information used by the ICC may be gathered and analysed collectively for a group of users. Further useful information would be a list of

content in order of the difficulty experienced by the users. The content which proved most difficult would be listed first, and the content that was least difficult would be listed last.

5 Beyond user and content comparisons, the EDSS also analyses the ICC configurations of each user. By comparing the ICC settings and the actual user results, the EDSS may determine which ICC configuration settings proved most effective. The EDSS then reconfigures each user's ICC, and awaits the results. After a set amount of time, determined by the program managers, or by the EDSS, the user results are again compared and checked
10 for relative improvements over the previous period. These measurements are calculated for each user, and for the group as a whole.

The EDSS also provides a mechanism for finding the optimum settings of different programs. Each of the possible configuration values is given a start value. Then, by systematically changing the values, one at a time, and
15 measuring the overall user performance for each change, the system is able to determine optimum configuration settings. The system changes each value by a predetermined amount, the Δ value, as determined for each configuration setting. Boundary limits set the minimum and maximum values for each setting, which may be preset by the program managers, or by the EDSS. This
20 feature of the invention enables the optimum configurations for each different content type, and for the different goals set by each course, to be determined automatically.

The user profile is collected regularly and the results are compared to check for changes in the overall learning speed. If gains are made, then those setting
25 are saved, and the remaining settings are again systematically checked for any further improvements in the optimal configuration.

The EDSS is flexible to allow for variable ICC values. The systematic approach to testing all the possible configurations, and comparing there results remains unchanged regardless of the total number of ICC values.

5 It is desirable that the group of users may be able to use the system on a network, such as a local area network or the Internet, as networking plays an important role in rapid knowledge acquisition. By connecting users together they can see how their peers are doing and compare their own performance. When multiple users are connected, only information regarding the relative position of the user, the top user, the last user, and the average user. This
10 information serves mainly as a motivational tool. For some users, competition can create a strong sense of urgency that, in experiments has yielded substantial increases in the APH. This is not relevant to all types of content.

From an instructors perspective networking provides features which enable
15 and facilitate group management. This includes managing multiple content types, multiple students, on differing embodiments of the invention.

Fifth Embodiment

This embodiment in general terms fulfils the same functions as described with respect to previous embodiments. However, in this embodiment the way in
20 which the stimulus is presented to the user is modified in order to minimise the user's TPA (*Time Per Association*).

This is achieved by stimulating the user to answer questions quickly, by creating a sense of urgency, using a Stimulation Drive Interface (SDI). Ordinarily, a user must generate a significant amount of energy to motivate
25 themselves, and remain motivated to respond to each question as quickly as possible. This requires a high degree of concentration and motivation which is focused not on the content itself. SDI stimulates the user in a way that

naturally generates a sense of urgency without the user having to waste any energy creating that sense of urgency on his own.

SDI also creates an environment where the user's attention may be occupied by things other than just the content of the question. The degree to which the user's attention is occupied by things other than the content is a value that may be adjusted according to the user's level of familiarity with the content. In other words, the more familiar the user is with the question content, the less they need to concentrate to respond correctly. At a certain level of familiarity, it is said the content has become "second nature" to that user. This embodiment applies a systematic method of gradually introducing factors that force the user to focus a part of their concentration on things unrelated to the content. The SDI gradually nurtures the users ability to respond reactively and fluently to stimulus. This is the highest degree of "knowing" that the ICC strives to lead the user to achieve. Figures 56, 57, 58, and 59 show examples of the SDI.

Figure 56 shows a display for displaying and responding to a question. The question 56.1 and the answers 56.2 may be displayed anywhere on the screen 1.1 as long as they are clear and easy to see. Each answer choice is represented by a numbered lane 56.5. The user must select the lane that represents their answer choice by steering into that lane with a steering wheel/joystick 57.7, or with the arrow keys found on a typical keyboard, or other suitable input device. An object 56.6 shows the user's present lane, this could be any graphical object such as a car. The car moves forward, as represented by the movement of the lane division marking, or features such as tunnel 56.4, approaching the foreground. The user must make their selection by the time the tunnel 56.4 reaches the position object 56.6. In effect, this is a timer. The faster the tunnel approaches the position object, the less time the user has to answer. The user's answer is registered when the position object passes through one of the tunnel entrances 56.3.

The answer response which determines the user's choice to be correct or false, is the same as was previously described, as is the error response which may display the correct answer if the user's selection is wrong. Depending on the user's response, the "car" may accelerate, or decelerate, enabling the user to
5 view the correct answer, or to proceed to the next question. For example a correct answer causes the "car" to accelerate and an incorrect answer causes it to decelerate.

The degree to which the user must focus his attention on things unrelated to the content is determined by the difficulty of the input method. For example,
10 in figure 56 the road is straight and there are no obstacles. The user may not stray from the road, therefore, most, if not all of their attention may be focused on the content. If the road becomes convoluted, and obstacles such as other cars that must be avoided are added, and the position object is allowed to veer off the road, then the user will be forced to focus more of his attention
15 on the means of selecting the answer, rather than on the content itself. This becomes an exercise that trains the user to use content with only a fraction of his concentration, a skill that is most commonly associated with fluency. Therefore, as the user begins to master the content with which he is working, such distractions and obstacles may be introduced at an increasing rate.

20 The degree to which the SDI forces the user to concentrate on the means of answering a question, rather than on the content is called the degree of distraction. If the user answers a group of questions correctly, and in a short period of time with a high degree of distraction, then he is determined by the ICC to have reached a relatively high degree of fluency with that content.
25 Fluency is not limited to language based content, in this case it may apply to any content type. The questions which are displayed for the user are generated according to the means previously described.

A further example of SDI (*Stimulus Drive Interface*) is shown in figure 57. Instead of a driving a car on a track, the user must navigate through a maze. A question 57.1, and answers 57.2 are displayed on the screen 1.1, or played through speakers 1.2, or whatever the output medium is. Doors which
5 correspond to each possible answer selection 57.5 and 57.3 must be selected by the user. A "runner" 57.6 is controlled by an input device such as a joystick, a keyboard 1.8, or a mouse 1.9. The user must find their way out, by answering correctly. If the user does not answer within a set period of time, then they loose "life points". If they answer quickly, they are rewarded with
10 "life points" as indicated by 57.7. As with all the questions, if the user answers incorrectly, the correct answer may or may not be displayed according to the ICC settings.

The degree of distraction may be adjusted by adding obstacles such as holes in the floor, which must be avoided, or objects which attack, from which the user
15 must defend himself. Any object which forces the user to focus more attention on the means of answering the question rather than the content, increases the degree of distraction.

Figure 58 shows another version of the SDI. A question 58.1, and answers 58.2 are displayed, with moving "meteors" representing the possible answer
20 choices. The user guides the pointing object, in this case a space ship 59.5, to their selection. The input device may be a joystick, a keyboard 1.8, a mouse 1.9 or any appropriate device that enables the user to steer the "ship" into the answer. The speed of the ship is controlled by the user response so the SDI level is customised to the users ability, using the methods described
25 previously. The degree of distraction may be adjusted by making the targets move, and therefore more difficult to hit.

Lastly, figure 59 shows another version of the SDI. A question 59.1, and answers 59.2 are displayed, with moving targets 59.3 which represent the

possible answer choices. The user makes his selection by firing the gun 59.5 at the selection of his choice. If the "bullet" 59.4 hits the correct target, then the next question may be displayed. The degree of distraction may be adjusted by increasing, or decreasing the speed at which the targets move, and hence the time to answer. The targets themselves may be made larger, or smaller. The user may use any input device that enables him to aim the gun 59.5, including a joystick, a keyboard 1.8, a mouse 1.9, etc.

The ICC may determine the degree of distraction applied to SDI by increasing it by a predetermined amount distraction- $\alpha\Delta$ when the user answers correctly, and within a set time. The degree of distraction is lowered when the user answers incorrectly, or after the set time, by a set amount, distraction- $\beta\Delta$. Alternatively, the user may set the degree of distraction himself.

Sixth Embodiment

This embodiment, in general terms fulfils the same functions as described in previous embodiments. However, in the present embodiment the system forecasts the process which is likely to be made by a user on the basis of previous use made by others.

Forecasting is an element of ICC that serves several purposes. The first is for motivation. ICC uses the previous scores to determine the most likely score and time on a quiz. The information is then displayed at the beginning of a quiz as a goal to beat, then again at the end of a quiz to view the relative performance. The time and score forecasts are calculated using the following algorithms:

Forecast time = [(average time of the selected question types) + (average time of the selected content type) * (number of questions)] / 2

Forecast score = [(average score of the selected question types) + (average score of the selected content type)] / 2

Further application of forecasting includes determining the duration, depth, width, and content quantity of a course. At the beginning of a course the user may determine any of the above factors. The ICC then calculates the feasibility of the settings and displays the results, allowing the user to adjust them as necessary to generate achievable goals. The first forecasts are based on the performance of an "average user". From time to time the ICC updates the forecasts based on dynamically generated user performance. As the course advances, the forecasts necessarily become more and more accurate, since they are based on user averages.

The duration of a course depends firstly on the quantity of content the user selects to learn. All other factors remain unchanged, as content quantity increases, so does duration. Depth and width (see figure 38) determine the average number of associations necessary for content records to reach Ω -CDR, thus affecting duration. The TPA (*Time Per Association*) is a necessary component of the duration forecast algorithm, as is efficiency. Efficiency is the ratio of total session time to association time. Association time is the time actually used to make associations. When the user configures a course, or pauses the course to take a break, or simply loses concentration and starts daydreaming, then they are not making associations. Therefore, when forecasting the number of associations a user will make, efficiency must be considered.

Efficiency = association time / total time

TPA = average time of all user associations

Therefore,

Duration = $(\Omega\text{-CDR}-\lambda) * (\text{content quantity}) * (\text{TPA}) * (\text{efficiency})$

The duration necessary to learn a set amount of content is thus determined. $\Omega\text{-CDR}-\lambda$, as previously explained, is the average number of associations necessary to master 1 content record. Depth and width, as well as the average

number of errors, and the loss rate are included in its calculation. Content quantity is determined by the user, otherwise the entire content will be used for the calculation. The first session, ICC will have no user profile, so an average user profile is used initially, until the user generates enough data for their own forecasts.

Duration is thus calculated in hours. If the user determines a desired time per session, and the number of sessions per week, then ICC may also forecast the end date of the course. Alternatively, ICC may calculate the amount of content that the user will master, to a predetermined degree, within a specified date.

Content quantity = Duration / $[(\Omega\text{-CDR}-\lambda) * (\text{TPA}) * (\text{efficiency})]$

Changes the user makes to the depth and width are calculated according to the $\Omega\text{-CDR}-\lambda$ algorithm used to explain figure 38.

Forecasts give the user an idea of the amount of content he can master in a given period. The ICC adjusts the course accordingly, so that it fits the needs of different users. Some users may want a thorough course regardless of the amount of time it takes, while others may prefer a quick review before an exam. All scenarios may be accommodated according to the user's preferences.

Instructors, or training program supervisors may use forecasting to determine the necessary length of a course, or even who should participate. The total cost of a course can be figured with the forecasted number of hours and the cost per hour.

Examples of how the invention may be used in different applications are described below:

Example 1: Rapid Language Acquisition

In this scenario, the content is Japanese from English. There are 2,500 question types for each word, ranging from beginner, to advanced, as well as 2,500 Japanese words, each containing 6 association types: word, definition, synonym, antonym, homonym, and sample sentence. All association types include 5 media types: native language text & audio, target language text & audio, and graphics.

On the first display screen, there appears a wizard, the ICC wizard. ICC stands for (*Intelligent Course Customisation*), and through a series of short questions, it will optimise the course settings for the user.

10 1. The first question relates to duration. How much time do you wish to study per session? The user inputs a time span, say 1 hour.

How many sessions per week? "5" times per week.

How many weeks do you wish to study? 7 weeks.

The program then calculates the end date of the course: "July 14th".

15 2. Next, the ICC wizard asks: "Select the amount of content you wish to include in your course". A scroll bar enables the user to select anywhere between 1 and 2,500 words. "You may also select specific lesson numbers." The user in this scenario selects the first 1,000 words.

20 3. The ICC wizard continues: "Select the content width. Width is the number, and type of associations that will be included in the course." The user selects word, definition, sample sentence, and synonym for a total width of 4.

25 4. Now, select the depth to which you wish to learn the content. Another scroll bar appears allowing the user to choose anywhere between 0% (no recognition) and 100% (fluency). "100% fluency?" . The user selects 100%.

5. Next, the ICC wizard asks: "How familiar are you with the selected content?" Another scroll bar appears to enter the selection: 0% (not at all) to 100% (very). The user may opt for a short diagnostic test to determine his familiarity. A short, but representative set of questions appear 1 at a time on the screen in multiple choice format. ICC tabulates the results and automatically ICC sets the familiarity to 15%.

The ICC wizard, having calculated the users' settings, responds. "According to the average student's learning speed with this learning tool, it will take you 14 weeks to finish this course. Would you like to continue anyway, manually reset the course, or have ICC automatically configure a course for you?" The user may choose to manually modify the depth setting which he sets to 65% and adds an extra week to the duration.

This time ICC responds, "ICC will use your settings to generate a customised course for you. As ICC collects detailed data on your performance it will generate a CDP (*Customised Diagnostic Profile*) and use it to make regular adjustments to your course settings."

Before starting, please read the following instructions to enhance your learning speed:

1. The number of associations you make per hour (APH) is an important factor in determining learning speed. Most people can make an association, and respond to it, in 1 to 5 seconds, or 3,600 to 720 associations per hour. It is more effective to make many associations in a short time than it is to view each association for a long time. When you answer questions, don't spend too much time on them. If you make a mistake, then that same word will come back at a higher frequency. Your APH (*Associations Per Hour*) will always be displayed on the screen for you to monitor.

2. Efficiency is the amount of time you spend productively learning, i.e. making associations, versus the total session time. In other words, once you have started the program, any time not used to make associations will lower your efficiency.

5 Would you like to start now?"

The user presses a "Start" button. The Automated Notebook (Auto Note) appears. There are 4 different Notebook views, card view, word form view, list view, and association view. ICC selects association view which starts off with one Japanese word on the left, its pronunciation underneath it in roman
10 letters, an audio button beside them, and the English word on the right. On the lower right, a flashing button that looks like the start arrow of a CD player lures the user to click it.

Auto Note starts. The words disappear for a second, and then the Japanese word reappears, followed shortly by an audio recording of the same word. As
15 the audio ends, the pronunciation appears under it for about the time it would take to repeat the word once. Then the English word appears, completing the association. This sequence is repeated 50 times, once for each word in lesson 1 and 2. At 12 seconds per word, it is bellow average, but most of this content is new to the user, so the speed is justified. If at any time the user feels
20 uncomfortable with the speed, he may adjust it manually to suite his own pace. In any case, The user just completed 50 words in 10 minutes.

ICC continues with the Auto Note, this time adding a slight variation: the Japanese word still appears on the left, the audio also plays, but this time it skips the pronunciation, and goes directly to the English word. In place of the
25 pronunciation, there is a text input box for the user to type in the pronunciation. After hitting the "Enter" key, the Pronunciation displays briefly above the input box. If the user's response is incorrect, the audio re-plays, and he must type the pronunciation again, this time while looking at the

correct pronunciation still on the screen. This time, it takes an average of 17 seconds per word, for a total of 14min 16sec.

The user is now given the option to repeat the same Auto Note drill, but instead of typing the pronunciation, he can say it into a microphone. Using
5 voice recognition the program determines if the spoken response is correct or incorrect. The pronunciation as well as may be measured for accuracy. However, the user prefers to leave pronunciation practice for another day. He chooses to continue with the regular course; however, at anytime, he may choose to use the microphone to input his responses rather than the key board.

10 Meanwhile, ICC has been recording the time spent on each word, the errors that were made, and also the exact association types that have been viewed so far with each word.

Next, ICC automatically changes to the quiz window, and selects 5 easy
15 question types each using a combination of the association types just viewed in Auto Note - Japanese word, text, pronunciation, audio, and the English word. These are selected from a possible 2,500 question types that range from very easy, to very difficult.

The first question is true and false, the Japanese word audio plays, and the
20 pronunciation is displayed on the screen. Is it the correct answer, true or false? The timer is counting down from 8 seconds, but it's fairly easy, and The user answers in 3 seconds. Next, the question type changes to the Japanese word in text as the question, and 2 English words as possible answers. He's just seen these, so again, he answers in about 3 seconds. The
25 next question is the Japanese word audio, and again 2 English words as possible answers.

These questions are all based on the same content, i.e. the same word. All 5 question types will appear with each word, before changing to the next question word and starting the questions again. ICC has determined, that

these being new words, it is best to concentrate clusters of questions on each word before moving on. Other possibilities include doing one question type with each word, before moving to the next question type, or rotating through a question type, and a word at the same time. Each of these patterns have their purpose in the context of a long term course - in this case, an 8 week course.

The user answers the question in 4 seconds. He continues at this pace until he makes his first error on the 26th question, or the 6th word. Immediately, the correct answer is displayed in large red letters, concurrently with the audio and his incorrect answer. The user must type the answer to continue. He does so having used a total of 14 seconds, but the impression left in his mind is significant.

There are over (25) different error response types from which ICC can chose according to a number of different criteria. Some require only a 2 to 3 second break from the question flow.

After the 132nd question, The user notices that as he answers correctly the timer of those question types starts slightly lower, at about 7 seconds for some, and the question types that he has erred most on, now provide him with about 9 seconds to complete the question. When he doesn't answer in time, the same thing happens as when he makes an error. The correct answer is displayed, the audio played, and he must type in the answer before continuing.

The user continues until he finishes all 250 questions. His time is 27min 08sec, an average of 6.5 sec per question. Double the average, but this is still his first session. ICC meanwhile, has recorded each error, the time of each response, and much more. The question types that proved most difficult will occur at a higher frequency in the future, as will the erred words which have all been categorised into over 15 different groups. At this stage, no word has reach the highest group where ICC deems it to have been learned. It will

require that each of the association types, as determined in the width setting, is tested correctly, before that category (the Ω -CDR) is achieved.

With only 8min 36sec remaining in the 1 hour session, ICC suggests a quick review of all today's errors. Knowing the exact number of errors, and the users' APH (*Associations Per Hour*) ICC prepares a review to fit the time.

5 The user continues to answer questions, this time only the ones he responded to incorrectly. If he made an error on the audio and English word association for a word, then that same association would reappear. The words that proved most difficult in this session are then viewed again in the Auto Note leaving

10 just 1min 17sec in the session by the time he completes the review.

The user's average response time dropped to a respectable 4.5 sec per question in this last section, allowing him to review 97 problem areas. The APH for the first session is 397.

ICC now has data to do some preliminary calculations on the users performance. The user will never see these, but ICC will determine the appropriate course structure that will enable him to finish the entire course within the 8 weeks he set for the course.

15

ICC then alerts The user that the session time is up. "Would you like to continue?"

20 "No" he enters.

"Would you like to print out the content you studied today?"

The user clicks "OK".

"All the content, or errors only?"

He selects errors only.

25 "You made 64 errors today on 23 words. Print the top xx% of errors?"

The user chooses to print 100% of his errors today.

Example 2: Rapid Knowledge Acquisition

In this scenario the content is anatomy. It contains over 3,000 body parts organised into lessons such as Muscular structure, Vascular network, Pulmonary system, Nervous system, Muscular structure, and more. Similar to
5 Language content, each anatomical part consists of 7 association types: 1) anatomical part, 2) function, 3) physical location, 4) component parts, 5) related systems 6) a sample sentence using the anatomical part and 7) description. During the first session, The user set the width to 7 so as to include all the association types in his course. These also include text, audio,
10 and graphic representations of each association type.

According to the users' study program, he is ahead of schedule, and could finish 2 weeks early if he continues at the same rate. He starts the program and enters his password. It's important nobody else uses the program under his name, otherwise the profile would be rendered useless. Content he was
15 already familiar with would appear, and content he did not know yet would show only rarely.

The user is now well into the 6th week of study, and the system has accumulated a detail profile of his knowledge base. The course that the user does now is quite different from the early ones as a result of the profile.

20 The first window displays, and the user immediately clicks on the start button. There will be no new content added today, there are 400 content records which he has already started to learn, but has not yet mastered. ICC has determined that is the maximum amount of content for the user to be working on at one time to achieve an optimum learning speed. Therefore, today's
25 session will cover only content that was previously introduce. When the user masters a portion of this content, also determined from his profile by ICC, then new content will be introduced.

The quiz window opens and a blinking start button beckons. The user starts the quiz. One by one, questions based on the content that the user is least familiar with begin to appear. There are 17 levels of familiarity in total. As the user answers correctly, or incorrectly, the content moves up and down through the different levels. When new content is introduced, it starts at 0 (α -CDR), from there, it may go up or down. 15 (Ω -CDR) is the highest level, and content is considered "learned" when it reaches this far. However, sometimes content can fall back down again even from level 15 (Ω -CDR). The rate at which content falls from Ω -CDR is called the "loss rate", and the number of times it must be viewed again before returning to the Ω -CDR is called the "degree of loss"

The user's loss rate is only 3%, and the degree of loss is 2. ICC determines this is far too low. 10% loss rate usually leads to a faster overall learning speed. The user has been "over learning", which means he has spent more time than was necessary on each content for it to reach Ω -CDR. Once content has been learned to a certain degree, it is effective to view it at lower and lower frequency intervals. Whereas new and unfamiliar content requires multiple viewings each session, or high frequency intervals.

Each person must view content in order to learn it. Some must view it at a high frequency for a long period before they are able to retain it, and this also differs from content to content. Other people seem able to retain content more quickly and can spend their time more effectively covering more content at lower frequencies. It is different for each individual, and it takes time for ICC to gather enough information on the student to determine an optimum frequency for each CDR level that the content passes through on the way to becoming learned.

The user's APH (*Associations Per Hour*) is also fast at 740, and the expected drop in score when he is given more difficult question types is less than

average. ICC will therefore lower the overall frequency per content, and increase the general difficulty of the question types. It will do this in small increments until the degree of loss exceeds its pre-set maximum limit, or until The user's learning rate peaks, and begins to fall again - which ever comes
5 first.

The content for the first quiz is CDR-(-1). This is the content that the user is least familiar with. Previously ICC generated 5 easy question types for this category, now it will use 3 easy question types, and 1 slightly more difficult.

The first question is multiple match: 10 anatomical part graphics on the left
10 side, and the equivalent in text on the right side, in random order. The user must match these correctly in under 80 seconds. As he completes the first multiple match, the system prepares the next set of content until all the content has been viewed in multiple match. The next question type is also multiple match, but this time with the text of the anatomical part on the left,
15 and the physical location in text on the right. This is slightly more difficult, so instead of 10 rows of content at a time, ICC prepares only 6. Again, the user answers all the CDR-(-1) content before moving on to the next question type.

When the user finishes his first set of 6 rows, his errors are highlighted in red,
20 and he is given an opportunity to correct them. If they are still wrong, the program will display the correct answer and may, or may not require him to type the errors in a text input box before continuing to the next set of 6 rows.

Following two more question type sets, the user finishes the CDR-(-1) content, and prepares to move on to CDR-0. By this point he has already
25 finished 120 questions in 9min 40sec. As he completes each CDR level the question types get more and more difficult; however, his familiarity with the content is increasing at the same pace. Therefore, as he becomes more familiar with the content, the question types become more challenging.

An hour and a half into the two hour course, The user is on CDR-12 and has already completed 1,200 questions. He is aware of his APH (*Associations per Hour*) and also of his efficiency measurements on the bottom of the screen.

5 The user answers the next question in 4 seconds, but the next question is unexpectedly difficult, a graphic picture of the anatomical part and a text box to type in the function. The answer is correct, but it takes a full 7 seconds.

The next question displays the names of 5 anatomical parts, one of them is part of the pulmonary system, and the rest are part of the vascular network. The user must chose the odd one out. These question are more challenging
10 than usual, and The user's APH and score may suffer as a result. However, since there will be no new content today, the overall APH will probably increase.

By the end of the 2 hours, the user has finished all 400 content records. He chooses to do a quick 5 minute review of the errors this session. 60 simple
15 True False questions which he zips through. He finishes a total of 1,660 questions in 2 hours 15 minutes.

This session, over 35 content records reached Ω -CDR. That leaves 355 content records which the user has started to learn, but hasn't mastered yet (μ -CDR). Next session ICC will not add any new content because each new
20 content unit is 75 records, and at this point 75 records would increase the μ -CDR (content which has been introduced, but not yet mastered) to beyond its maximum limit (in The user's case, 400). At the present learning rate, it will take 2 more sessions before new content is once again added. By this time the μ -CDR will be 295.

25 The user quits the program and heads to class where his professor who has access to each students profile through the school network, checks the strengths and weaknesses of each student. He modifies his planned lecture to cover the areas of weakness in more detail. He also prints a pop quiz

customised to each student. With each student profile, the professor can print a quiz using a ratio to determine the mix of CDR-1 to Ω -CDR content. Having just finished his own session, the user scores perfect on the quiz.

Example 3: Rapid Scenario Acquisition

5 The content used in this scenario is Banking procedures. Each procedure contains the following association types: 1) procedure name, 2) procedure purpose, 3) procedure results, 4) procedure steps - (a) initiation 5) procedure steps - (b) content, and 6) procedure steps - (c) conclusion. The procedure steps a, b, c, are a series of animations which demonstrate the procedure
10 visually. For example, if the procedure involves filling out forms, then the animation would first display the form, highlight, one by one, the information that needs to be added in the appropriate sequential order, and then display the next form when there are more than one. If the procedure involves a computer, then the animations would run through the steps necessary to
15 complete the task. If the procedure involves the assemblage of parts, then the animations would play a sequence which demonstrated the sequence, and method of assembling each individual part. If the procedure involves the operation of a machine, or vehicle, then the animations would demonstrate the proper operation of the machine, or vehicle, step by step, in different
20 circumstances.

Not all procedures contain the same number of steps, therefore this number is variable to meet the requirements of different content. In the Bank's case, all procedures have been classified into 3 categories - (a) initiation, (b) content, and (c) conclusion. Each of the above mentioned association types may be
25 combined in ways similar to the content in scenarios 1 & 2.

The ICC is pre-configured for this course. However, as the user nears the end of her course, the program has optimised itself to enable her to achieve Rapid Scenario Acquisition.

At this point, the user is already familiar with all the content, meaning that all the content has reached Ω -CDR. ICC does not have to use multiple questions for each content record, as it has already been learned, and just requires periodical "refreshing". Learned content is sometimes forgotten; however, it usually quickly returns to memory once the association has been re-made. The main objectives of ICC with the Ω -CDR content is two pronged, 1) to enable the user to view as much content as possible in a short period of time. 2) to increase the difficulty of the question types, and expand the variety of association types.

10 The user starts the session. The first question is the scenario name as the question, and 4 scenario animations of the scenario content step, played one after the other. One of the scenario animations corresponds to the scenario name. Next, the scenario name appears again as the question, but this time, the answer includes 6 scenario animations. The user must combine 3 of those
15 animations in the proper sequence to produce the correct answer. The user has done these enough times that the real issue today, is not her score, but rather her time. Real life situations require that she responds to the possible scenarios quickly and accurately.

The user continues the course for 1 hour, only making 4 errors. These have
20 been replayed several times as a result. They will also appear more frequently next session.

Meanwhile, in another office on the other side of the Bank, The user's manager is reviewing the progress of all the users. The user's scores and times are average, but her learning rate is nearly 1 ½ times faster than the
25 average. Fast learners save the Bank time and money. Statistically, they are also superior performers. With the DSS (*Decision Support System*) the manager can forecast the time it will take for the user to achieve the necessary competency to fulfil the requirements of the new posting. The manager also

plugs in the user's hourly wages, and a few other costs associated with training, and determines the total cost of her training program.

Example 4: Rapid Skill Acquisition

5 The content in this scenario is entirely computer generated and is comprised of the following categories: 1) Colour Intelligence, 2) Timing Intelligence, 3) Motion Intelligence, 4) Shape Intelligence, 5) Spatial Intelligence, 6) Sound Intelligence, and 7) Estimation Intelligence. There are many more "Intelligence" types; however, these will be introduced when the user has
10 made sufficient progress on the initial categories mentioned above.

Colour Intelligence content is generated by the computer using a combination of values from 0 to 255 for R, G, B (*Red, Green, Blue*). For example, an RGB value of 0,0,255 will produce blue, and a value of 255,255,0 will produce yellow. Over 16 million colours may be produced using this method;
15 however, many display systems are only capable of producing 65 thousand colours, and some, even less. Anything less than 65 thousand colours would not be useful for the purposes of the course.

The first question the user sees, is a group of 6 colours, from which she must select the two most similar colours. The degree of difficulty is determined by
20 several different factors.

1) the relative difference between the similar colours, and the dissimilar colours is one factor. For example, if the two similar colours are both shades of green, and the dissimilar colours are red, blue, black & yellow, then for most people, the answer would be obvious. However, if the two
25 similar colours are green, and the dissimilar colours are also green, but only slightly more different by an imperceptible amount, then most people would find the answer un-obvious. In fact, because of the computers ability to generate colours from numeric values, and because of the sheer

number of colours that can be generated, it is possible to generate colours which are imperceptibly different. Starting from an obvious difference in colours, and systematically diminishing the difference until it becomes imperceptible is one way to determine the degree of colour perception of an individual.

5

2) The degree of similarity of the similar colours is another factor which determines the difficulty of a question. If the two similar colours are in fact identical, then the answer is obvious to most people even when the other colours are relatively close. * If the similar colours are quite dissimilar, then the answer is less obvious.

10

3) The number of colours available to choose from is another factor. The more colours that must be compared before answering generally increases the difficulty of the question. It also usually increases the time needed to answer the question.

15

4) Colour theme is a factor which determines the general colour of the entire question. Examples of this include where all the colours are variants of cyan only, or are all within a certain brightness range, etc. Different levels of colour sensitivity are affected by the theme colour. Some people may be above average in determining colours within the cyan range, but below average with a different colour theme.

20

5) Background colour is similar to colour theme in effect; however it deals not with the colours within the question itself, but rather the background colour on which the question colours are viewed.

25

Each of these factors can be numerically controlled enabling the program to create a course structure based on a systematic approach, similar to the above mentioned scenario types. By first measuring the user's colour perception range in a multitude of themes and backgrounds the system can generate a course aimed at gradually increasing the user's range of colour perception.

Beyond determining similar colours, there is a wide range of possible question types based on the above principles. Examples of these include the following: select the most dissimilar colour, select colours with the highest / lowest contrast, select the sharpest / dullest colour, and so on with each category of colour quality.

As the user answers each question they get progressively more difficult. Inevitably, the user makes an error. Immediately, the correct colour is displayed giving the user a feel for how to correct her perception. Since this is not a knowledge based exercise, it's not simply a matter of the user remembering the correct answer, it is a skill she needs to develop overtime, through practice. Over time her perception of the different qualities of colour will become sharper.

For now, ICC records the user's responses and creates a detailed profile. From the profile, ICC will determine where the user started to have difficulty perceiving the target colour attribute. Then, it will start a series of drills to target that colour attribute within the range of the user's weakness. The first question starts from a degree of colour that the user can easily perceive, and then in increments which are also determined by the ICC, the questions become more and more difficult, until finally, she can no longer perceive the target colour attribute. At this point ICC will continue from the original starting point again and perform the same drill over again. After a certain number of times, determined by the ICC, the theme, background colour, question type, or some other variable is changed, and a new target is defined.

Next, pairs of colours appear on the screen, and the user is asked to choose the odd one out. Each pair has the same hue, except for 1. The user chooses it, and moves on to the next question. The difference in the hues soon becomes imperceptible, and the user starts the exercise again. Her scores have already started to improve since her first session.

- Today, the user will start a more advanced drill, designed, not only to improve colour perception, but also to work on colour logic. The first drill includes two colours as the question, and 4 colours as the multiple choice answers. The user must determine which colour is the result of mixing the 2 question colours. As with the previous questions, the answers are obvious at first, and gradually become more and more difficult until she is no longer able to make any meaningful answers. Now 3 colours are added. ICC continues changing all the parameters of each question type, collecting a detailed profile of the user, and, at the same time, developing her colour logic and perception.
- There are hundreds of possible question types categorised by there level of difficulty. ICC determines when to move on to new question types according to the user's past performance; however, at any time, these settings may be overridden by reconfiguring the ICC and the goals of the overall course. If the user wishes, she may pick and choose only the drill types she enjoys.
- When the user is don with the colour drills, she moves on to Timing Intelligence drills. These work according to the same basic principles as the colour intelligence drills, mainly by producing drills of varying difficulty by adjusting a number of parameters. For example, the basic timing question type is a light (or sound) that lights up for a eertain amount of time. The answer is comprised of 4 other lights that light up for different lengths of time. The point is to choose the most similar / dissimilar, etc. Timing patterns can be made more complex by turning the light on and off at set intervals, and further, by gradually shortening or lengthening the intervals. Generating intervals of random length, and also the length of time the light is on creates the most complex pattern, that, by extending it through time, eventually becomes impossible for any human to answer.

Other adjustable parameters include the brightness of the light. The changing brightness may also be used to create a pattern from which a number of questions may be asked.

5 As with Colour Intelligence, ICC determines the user's perceptual limits and creates drills which focus on the narrow band between what she can and can't perceive. Gradually The user's Timing skills and perception improve, and overtime she will start more advances and sophisticated timing drills.

Example 5: Psychometric testing

10 The content for this scenario is psychometric test questions. Psychometric tests have been around for a long time, and have evolved into highly sophisticated and accurate means of determining an individuals potential, among other things. Thousands of different psychometric tests exist, with new ones being developed yearly. Some companies even develop their own with the help of consultants to meet the very specific needs of their
15 organisation. The high degree of custom tailoring possible with psychometric tests solves the specific problems of the test users, but, at the same time, produces a new set of problems for those same users. Essentially, "which test do I use?" It not only requires an expert to create and develop these tests, but also one to determine which test is appropriate for the specified needs.
20 Further, the actual administration of the tests is a time consuming process that in many cases also requires an expert.

The content of a psychometric test is inherently different from the content of the previous scenarios in one essential aspect: there are no incorrect answers. Psychometric tests are set up not to determine a "score" but rather to
25 determine tendencies. As such, each answer choice has meaning that, when combined with other answer choices, generate profiles which are interpreted by more experts to yield sought out information about those being tested.

The cost of producing such tests is substantial as it requires research, experimentation, and usually the combined efforts of many experts. One challenge is to create a test which is applicable to a wide audience, while still being effective for each individual. While tests have been generally successful for large populations, they necessarily work on a statistical basis, and lack the sensitivity to individual cases that may skew overall results. Ideally, a test would be designed specifically for each individual, with an expert to administer it in an interactive way. Essentially, an interview, or several interviews to determine the specifics of each case, would allow for a more accurate and detailed survey. The ability to follow a line of questioning to check in greater detail for potential irregularities if certain "flags" are triggered is the advantage of the person to person testing model.

The supervisor has chosen a package that responds to many of these issues. He is now preparing the system for a company that has specialised needs due to the nature of their work environment. He answers a series of questions designed by experts to determine the most appropriate test out of a possible several thousand. Each test has its own specific use, and in many cases their usefulness overlap. The questioning leads the supervisor to 1 specific test, as well as a group of related tests that are peripherally pertinent. He saves the results for reference to avoid having to go through the questions again next time the same company makes a request.

The supervisor observes as a user answers a series of general questions designed to determine his needs and thereby select a relevant test. The system suggests a number of tests to further explore the user's aptitudes. There is a match between the tests recommended for the user, and those for the company.

This is only the first step, and does not mean much, other than the user is a potential candidate for the companies' test. It is significant; however, that the

process to this point has been automated. If the user had found that his initial general profile also matched several other companies', then the matches would be listed according to the degree of match, from which The supervisor, or the user himself could select the appropriate direction. If, for example, the user had been a highly creative artist, looking to find an outlet for his creative energy, and the company was looking for someone with a penchant for performing repetitive and monotonous tasks, then there would not be a match.

In this case, there is a match, and the supervisor selects to have the user try the test most suitable for the company. As with all questions in the scenarios, each response is recorded and analysed by the ICC according to criteria established by the test developers and saved as "ICC rules". The user responds to each question one at a time until he finishes the entire test.

As it turns out, ICC's analysis shows that the user is suitable; however, the test also revealed a few danger spots which need to be tested further. The specific tests are listed, and the user is given the choice to continue. On completion of these tests, it is determined that the initial warning was a false alarm, and that the user is in fact ideally suited for the position.

Example 6: The following scenario demonstrates the SDI (*Stimulation Drive Interface*) and it's application within the context of the invention

The user is now nearing the 8th week of his Japanese course. He has improved steadily since starting, and now, after having learned all the content (meaning that all the content has reached Ω -CDR), he is ready to start the finale stage. Ω -CDR content review usually entails a rapid succession of questions aimed primarily at covering a large amount of content. In this way, ICC can quickly locate content that has been "loosed", or forgotten. "Loosed" content usually does not require much review once it has been located, since it has already been "learned" once. In any case, ICC will generate a higher frequency with these words as a precaution to make sure they are not loosed again.

The only way to maximise learning speed at this stage is to increase the APH. A drop of 2.5 seconds from an average of 7 seconds per question to 4.5 seconds per question leads to an actual increase of 286 questions per hour. Per session, this is a significant increase, per course it is even more significant.

When the user makes a determined effort to answer each question as fast as he can, his speed sometimes doubles for a short period. However, he inevitably slows down and falls back into his average pace. The mental energy required to generate an artificial sense of urgency necessary to answer each question as fast as possible is exhausting. In many cases, The user simply forgets about speed as his mind is fully engaged in answering difficult questions. This will change now that the user is ready for SDI.

SDI (*Stimulation Drive Interface*) is an interface designed to generate a natural sense of urgency, and thereby spur faster responses to the questions. The underlying ICC, and question generation remain unchanged, it is the interface that is transformed along with the answer input mechanism.

The first drill that appears on the users screen is multiple choice; however, the screen displays a car on a city road. The question displays on the top of the screen, it is the Japanese word. The answer choices appear at the mouth of 4 separate roads. Using the arrow keys on the key board, or a joystick with a steering wheel, the user "drives" the car to the correct street. Meanwhile, the next question already appears on the top of the screen, and the speed has increased slightly. After a few more correct answers, The user is struggling to answer each question as fast as his reactions will allow him. If ICC simply increased the speed indefinitely, it would eventually be beyond The user's ability to respond in time, even if he new the correct answer. This is not ICC's objective, so, by testing the user's reaction limits, ICC keeps the speed just within the user's ability to answer. Essentially, if the user responds as

fast as he can, then he will have enough time. The leeway ICC provides the user is also dependent on the difficulty of the question itself.

The user's APH climbs significantly. He continues to drive and takes a wrong turn which is his first error. The road immediately becomes rough, and he is
5 forced to slow down. At this time the correct answer flashes so that he understands his error. The next question displays without delay, and the user is back in the race. Sometimes the answer choices are represented by different lanes on a highway, whereby the correct answer leads straight ahead, and the error choices lead to off ramps, detours, and other temporary delays that
10 provide just enough time to see the correct answer.

The user is passed by another car. The system is networked, and now there is another user who is doing the same drill. The user may compete with other users, and thereby increase their APH.

It will be recognised that the above described embodiments are merely
15 examples of the invention, and that the features of various embodiments may be used in different combination from those described specifically above. Moreover, the skilled reader will recognise many modifications and substitutions which may be made without departing from the present invention. Accordingly, any and all such modifications and substitutions are
20 to be regarded as forming part of the invention. Protection is sought for any and all combinations of subject matter which may be disclosed herein.

CLAIMS:

1. An apparatus for training or analysis comprising processing means, a memory store for storing at least one first and at least one second data item and a plurality of relationships, output means and input means;

the apparatus being arranged to output said first data item and to accept the input of said second data item, said second data item being associated with said first data item by one of said plurality of stored relationships selected by said apparatus.

10

2. An apparatus according to claim 1, wherein said first data item is output in the form of a question defining said relationship, the answer to which comprises said second data item.

15

3. An apparatus according to claim 1 or claim 2, wherein said apparatus is arranged to output said data item in more than one medium, said relationship determining the medium through which the output is made.

20

4. An apparatus according to claim 3, wherein said medium is audio.

5. An apparatus according to claim 3 or claim 4, wherein said medium is video.

25

6. An apparatus according to claim 5, wherein said medium comprises text.

7. An apparatus according to claim 5, wherein said medium comprises graphical representation.

8. An apparatus according to claim 7, wherein said graphical representation is in the form of an animated sequence.

5 9. An apparatus according to claim 1, 2 or 3, wherein said relationship determines the format in which the question is made.

10 10. An apparatus according to claim 9, wherein said output format is multiple match.

11. An apparatus according to claim 9, wherein said output format is multiple choice.

15 12. An apparatus according to claim 9, wherein said output format is dictation.

13. An apparatus according to any of claims 1 to 12, wherein said relationship is randomly selected from said plurality of relationships by said processing means.

20 14. An apparatus for training or analysis comprising processing means, memory means for storing quiz data comprising corresponding question and answer data, input means comprising a manually operable aiming device and output means comprising a monitor screen;

25 the processing means being arranged to output a video arcade game interface to said monitor screen, the processing means being further arranged periodically to output a question to said output means, the apparatus being arranged to receive a response to said question through said input means, the

processing means being arranged to compare the response with the stored answer data corresponding to the question.

5 15. An apparatus according to claim 14, wherein said processing means is arranged to increase the difficulty of the arcade game if responses match the corresponding stored answer data and/or to decrease the level of difficulty of the arcade game if one or more responses do not match the corresponding stored answer data.

10 16. An apparatus according to claim 15, wherein the level of difficulty of the arcade game may be varied by adding or removing features which serve to distract the player of the game.

15 17. An apparatus according to any one of claims 14 to 16, wherein said manually operable aiming device comprises a gun.

 18. An apparatus according to any one of claims 14 to 16, wherein said manually operable aiming device comprises a steering wheel.

20 19. An apparatus according to any one of claims 14 to 16, wherein said manually operable aiming device comprises a joy stick.

 20. An apparatus according to any one of claims 14 to 16, wherein the question is output to said monitor.

25

 21. An apparatus according to claim 20, wherein the response to said question is presented as a selectable feature in the environment of the game.

22. An apparatus according to claim 21, wherein the response to said question is selectable from a plurality of selectable features.

23. An apparatus according to any one of claims 14 to 22, wherein
5 the processing means is arranged to receive said response only during a predetermined time after outputting the question.

24. An apparatus for training or analysis comprising a memory
store for storing a plurality of data items, each said data item being associated
10 with at least one further data item, output means and input means;
the apparatus being arranged to output a series of said data items and
to accept corresponding input data items during a subsequent predetermined
time period, the apparatus being further arranged to compare each input data
item with the further data item which is associated with the corresponding
15 output data item, and to vary said predetermined time period in dependence on
said comparison result.

25. An apparatus according to claim 24, arranged to increase said
predetermined time period when one or more input data items do not
20 correspond to said associated further data item.

26. An apparatus according to claim 24, arranged to decrease said
predetermined time period when one or more input data items do correspond
to said associated further data item.

25

27. An apparatus for training or analysis comprising processing
means, memory means, input and output means; the memory means being
arranged to store a set of question data, the processing means being arranged
to repetitively generate questions by selecting a subset of said set of question

data and outputting said questions through said output means; the processing means being further arranged to receive answers to said questions through said input means, and to determine whether the received answers are correct; the processing means being still further arranged to calculate the proportion of the question data which is correctly answered and to control said proportion by introducing new question data to said subset and/or removing existing question data from said subset.

28. An apparatus according to claim 27, wherein the processing means is arranged to vary the size of said subset.

29. An apparatus for training or analysis comprising processing means, memory means, input and output means; the memory means being arranged to store a set of question data, the processing means being arranged to repetitively generate questions from said set of question data, and to output said questions through said output means; the processing means being further arranged to receive answers to said questions through said input means, and to determine whether the received answers are correct, wherein the processing means is further arranged to monitor question data used in questions incorrectly responded to, and to vary the frequency of occurrence of said incorrectly responded to question data being incorporated in future questions.

30. An apparatus according to claim 29, wherein the frequency of occurrence of each said incorrectly responded to question data is increased.

31. An apparatus according to claim 29 or claim 30, further comprising means for allocating a difficulty weighting to said incorrectly responded to question data.

32. An apparatus according to any one of claims 29 to 31, wherein the processing means are further arranged to select a format, from a plurality of stored formats, in which to output each question, the processing means being further arranged to record the formats used in questions incorrectly
5 responded to, and to increase the chance of said formats being incorporated in future questions.

33. An apparatus according to claim 32, further comprising means for allocating a difficulty weighting to each said format.

10

34. An apparatus for training or analysis comprising memory means for storing a set of data records, and input and output means; the apparatus being arranged to repetitively generate stimuli for output to a user from said set of data records in accordance with a pre-defined method; the
15 apparatus being further arranged to receive corresponding responses from the user, to determine the validity of said responses, and to determine the rate of learning of the user therefrom.

35. An apparatus according to claim 34, arranged to receive
20 training requirements input by a user, and to generate a set of preferred teaching parameters for the user.

36. An apparatus according to claim 34 or 35, being arranged to
25 compare the rate of learning with an estimated rate of learning and to vary said pre-defined method in accordance with said comparison.

37. An apparatus according to claim 34, 35 or 36, arranged to independently train a plurality of users, comprising means to calculate an average preferred set of said teaching parameters.

38. An apparatus according to claim 34, arranged to compare the validity of the user's responses to repeated questions in order to determine the rate at which the user's knowledge is lost.

5

39. An apparatus according to claim 38, arranged to adjust said method to maintain said loss rate at a predetermined level.

10

40. An apparatus according to claim 34, arranged to periodically calculate the average number of repetitions required for the user to achieve a predetermined level of learning of the information comprised in the selected data records.

15

41. An apparatus according to claim 40, arranged to calculate a time period required for the user to achieve said predetermined level of learning.

20

42. An apparatus according to claim 40 and 41, arranged to adjust said method in dependence on said calculation.

25

43. An apparatus for training or analysis comprising memory means, processor means for generating questions and validating corresponding answers, and input and output means; said memory means being arranged to store a plurality of question and answer data, a plurality of relationships, associating each question data with at least one answer data, a plurality of question and answer formats, and a plurality of media types through which each said question and answer may respectively be output and input; the processing means being arranged to generate a question by selecting one of said question data and one of said relationships in order to determine

said corresponding answer data, and one of said formats and one of said media types to determine the question to be output and the correct answer.

44. An apparatus for training or analysis comprising memory
5 means for storing question and answer data, input and output means, processor means for generating a series of questions to be output and validating corresponding input answers; the processing means being arranged to measure a user's familiarity with said series of question and corresponding answer data; the processing means being further arranged to detect an error in
10 an input answer and to provide one or more corrected outputs in response to said error.

45. An apparatus according to claim 44, arranged to vary said one or more corrected output in dependence upon said familiarity.

15 46. An apparatus according to claim 45, wherein said corrected outputs comprise one or more further questions associated with said question, the response to which was an error, said apparatus being arranged to vary the difficulty of said one or more further questions in dependence upon said
20 measured familiarity.

47. An apparatus for training or analysis comprising input and output means, a memory means for storing stimulation data and a stimulation generator for generating a series of stimuli comprising two or more outputs to
25 a user, the apparatus being arranged to allow the user to discriminate between said two or more outputs of each stimulus using said input means; the apparatus being further arranged to increase the similarity between the two or more outputs of each successive stimulus until the user is no longer able to discriminate between said outputs.

48. An apparatus according to claim 47, wherein the series of stimuli comprise visual outputs.

5 49. An apparatus according to claim 48, wherein the user may discriminate between said two or more outputs of each stimulus on the basis of a difference in any one of brightness, hue, saturation, shape, size, number, movement or pattern.

10 50. An apparatus according to claim 47, wherein the series of stimuli comprise audible outputs.

15 51. An apparatus according to claim 50, wherein the user may discriminate between said two or more outputs of each stimulus on the basis of a difference in any one of pitch, rhythm, duration or volume.

52. A method of controlling the difficulty of a video arcade game comprising the steps of:

20 periodically generating a question, and outputting said question via the medium of said arcade game;

 receiving a response to said question through the selection of a feature in the environment of the game by a player of the game;

25 comparing the received response with a stored answer, and increasing the speed of the arcade game if the received answer matches the stored answer, or decreasing the speed of the arcade game if the received answer does not match the stored answer.

53. A method of implementing an educational or analytical quiz, comprising the steps of:

determining an initial quantity of question data;
repetitively, generating questions utilising the available quantity of
question data;
measuring the proportion of the correct responses; and,
5 varying said quantity of question data in dependence on the proportion
measured.

54. A method according to claim 52, wherein the said quantity of
question data is increased if the measured proportion is above a predetermined
10 level, and decreased if the quantity of question data is below said
predetermined level.

55. A method according to claim 52, further comprising the steps
of:
15 recording the questions which are incorrectly responded to;
recording the questions which are correctly responded to;
increasing the chance of question data, incorrectly responded to in
previous questions, being incorporated in future questions and decreasing the
chance of question data, correctly responded to in previous questions, being
20 incorporated in future questions.

56. A method according to any one of claims 53, 54 or 55, further
comprising the steps of:
randomly varying the format in which the questions are generated;
25 recording the number of questions of each format which are
incorrectly responded to;
recording the number of questions of each format which are correctly
responded to;

increasing the chance of utilisation of question formats where previous questions utilising those formats have been incorrectly responded to, and decreasing the chance of utilisation of question formats where previous questions utilising those formats have been correctly responded to.

5

57. An apparatus substantially as herein described with reference to the accompanying drawings.

58. An method substantially as herein described with reference to
10 the accompanying drawings.