

ARTIFICIAL INTELLIGENCE APPLICATIONS TO
COMPUTER-ASSISTED INSTRUCTION

Project Progress Report No. 3:

THE DESIGN OF A PROTOTYPE AI/CAI COURSE EMPLOYING
A RULE-BASED TUTORIAL STRATEGY

Jesse M. Heines

Systems Based Courseware
Educational Services Development & Publishing

Tim O'Shea

The Open University
Milton Keynes, England

May 3, 1983

d	i	g	i	t	a	l
---	---	---	---	---	---	---

Educational Services
Bedford, MA 01730

1.0 ABSTRACT

This report presents a specification for a prototype AI/CAI course to teach the ReGIS P, V, and C commands using a rule-based tutorial strategy à la that developed by Tim O'Shea. It is the third in a series of reports on my AI/CAI research. Previous reports were:

1. "Basic Concepts in Knowledge-Based Systems", and
2. "Where AI Can Fit in CAI".

The purpose of my research is to investigate the applicability of artificial intelligence to improving:

- the quality of CAI courseware and/or
- the efficiency of the CAI course development process.

The course described in this report addresses the first of these goals by using AI production rules to control branching between a large number of independent course modules. The full set of controlling rules defines the course's tutorial strategy.

The prototype AI/CAI course defined by this document has four basic parts: a set of teaching operations, a task analysis, a student model, and means-ends guidance rules. The teaching operations are the different instructional activities that the CAI program can present. These may be CAI lessons of the type currently produced by Educational Services, exercises directed at reinforcing specific instructional objectives, "laboratory" sessions in which students try out graphics commands in a controlled environment, or tests. The task analysis is a hierarchical list of the skills needed to master the material being taught. It is represented as a directed graph that defines the prerequisite relationships between each skill. The student model is a representation of the student's knowledge in terms of the task analysis and a history of the student's interactions. The means-ends guidance rules relate states defined by the student model to sets of teaching operations. These rules determine which teaching operations the AI/CAI program will present next for different student states.

The report describes each of the parts listed above and the ways in which they interact. It describes how the prototype AI/CAI course might be programmed, a task that I began during the last month of my assignment at The Open University. I hope to continue my efforts to implement on the prototype course into FY84.

2.0 TABLE OF CONTENTS

1.0	ABSTRACT	2
2.0	TABLE OF CONTENTS	3
3.0	THE STRUCTURE OF A RULE-BASED COURSE	4
4.0	WHY BUILD AN INTELLIGENT TUTOR ON REGIS?	5
5.0	THE VOCABULARY OF TEACHING OPERATIONS	6
5.1	Expository Demonstrations	6
5.2	Directed Exercises	6
5.3	ReGIS Laboratory	9
5.4	Formal Tests	9
6.0	THE TASK REPRESENTATION	11
6.1	List of Component Skills	11
6.2	Prerequisite Relationships: Directed Graph	11
6.3	Prerequisite Relationships: Production Rules	16
7.0	THE STUDENT MODEL	22
7.1	Skill Status	22
7.2	Learning Rate and Learning Style	23
8.0	THE MEANS-ENDS GUIDANCE RULES	25
8.1	Representation of Left-Hand Sides	26
8.2	Representation of Right-Hand Sides	27
9.0	CRITIQUE AND CONCLUSIONS	29
10.0	RESEARCH EFFORTS TO BE PURSUED	31
11.0	REFERENCES CITED AND RELATED READINGS	32

3.0 THE STRUCTURE OF A RULE-BASED COURSE

O'Shea (1979) has argued that one of the most important aspects of any CAI program is its response-sensitivity: to assert that one teaching program is more response-sensitive than another teaching program is to claim that in some sense it is more adaptive to the individual learning needs of the students being taught than the other program.

O'Shea's work achieved response-sensitivity by adopting Hartley's (1973) framework for adaptive teaching programs. This framework consists of:

- a "vocabulary" of teaching operations,
- a representation of the task,
- a model of the student, and
- a set of means-ends guidance rules.

The teaching operations are the different instructional activities that the CAI program can present. In the course design described in this document, these operations take the form of on-line presentations of new material, exercises directed at reinforcing specific instructional objectives, "laboratory" sessions in which students try out graphics commands in a controlled environment, and formal tests.

The representation of the task is a detailed task analysis listing each component skill needed to master the material being taught. It is represented as a directed graph that defines the prerequisite relationships between each skill.

The student model is a representation of the student's knowledge in terms of the task analysis and a history of the student's interactions with the program. It can be thought of as a state vector that describes the student's degree of mastery for each component skill and various other pertinent student characteristics.

The means-ends guidance rules relate states defined by the student model to sets of teaching operations. These rules determine which instructional activities the CAI program will present next given different student states.

Each of these components is described in detail in the sections that follow.

4.0 WHY BUILD AN INTELLIGENT TUTOR ON REGIS?

It is my hope that the prototype AI/CAI course specified in this report could be used to test the appropriateness of Hartley's framework and the feasibility O'Shea's rule-based implementation for building CAI courses for Educational Services. The subject matter chosen for the prototype course is critical if this hope is to be realized. The subject matter must be:

- sophisticated enough to yield a non-trivial task analysis yet simple enough that the number of component tasks that can be handled within the implementation time frame,
- applicable to a wide range of students at various levels so that its response-sensitivity can be evaluated,
- suitable for the type of treatment to be implemented.

A course that limits itself to the ReGIS P, C, and V commands fits these needs precisely. Additional arguments for using ReGIS as the prototype subject matter are as follows.

- ReGIS is very simple, allowing all states of the language to be known completely.
- ReGIS is very easy to parse, allowing sophisticated student/computer interactions such as a graphics "laboratory" with detailed error messages to be programmed relatively easily.
- The results of student responses can be shown visually, providing exceptionally strong feedback.
- The student model can be defined in terms of the ReGIS programming skills that the student possesses.
- The subject is discrete and can be readily modularized.
- Graphics is a popular subject area not currently addressed by Educational Services' curriculum.

If the prototype course proves feasible, it might be possible to extend it into a full graphics course. The final product might teach PROFESSIONAL graphics or even the SIG GRAPH core standard. A major roadblock to these extensions, of course, will be the course's size in terms of computing resources. I believe that it would be difficult to implement an AI/CAI course of this scope on a very small system, but the effort to do so might provide valuable insight into optimization of the techniques being researched. The approach would probably prove most feasible, however, for products at the higher end, especially ones aimed at users who will vary greatly in their levels of sophistication.

5.0 THE VOCABULARY OF TEACHING OPERATIONS

The prototype AI/CAI course will attempt to implement four types of teaching operations:

- expository demonstrations,
- directed exercises,
- a ReGIS "laboratory", and
- formal tests.

Due to time considerations, not all of these teaching operations will be implemented for every module in the course. All four will be implemented for at least some of the modules, however, so that the full power of the course structure and its AI functionality can be demonstrated.

The characteristics of these operations are described in the sections that follow. Their interrelationships are shown in Figure 1.

5.1 Expository Demonstrations

Expository demonstrations are basically "press-RETURN-to-continue" slide shows that introduce concepts. Of all the teaching operations, they exhibit the lowest levels of interaction and response-sensitivity. Their purpose is more "telling" than "teaching", and they last about 3-5 minutes each. While one or two orientation questions may be incorporated into the presentation, the student's only real option is to interrupt the operation and call up a control menu.

These sections make heavy use of graphics to relate ReGIS commands to screen actions. It would be nice to tie the presentations to an audio track to get around the need to read large passages of explanatory text from the screen. This might eventually be possible using the Telephone Management System (TMS) on a PROFESSIONAL. It might also be possible to simulate TMS with a computer-controlled audio cassette tape or even the sound track of a videodisc, but the preliminary version will not try to take this possibility into account. It will use screen text to provide the needed commentary in our normal manner.

5.2 Directed Exercises

Exercises provide either computer-generated or prestored problems for students to solve. In the prototype course, student responses to these problems will usually take the form of ReGIS command strings. The CAI program will parse student responses to allow extensive error checking and provide detailed feedback.

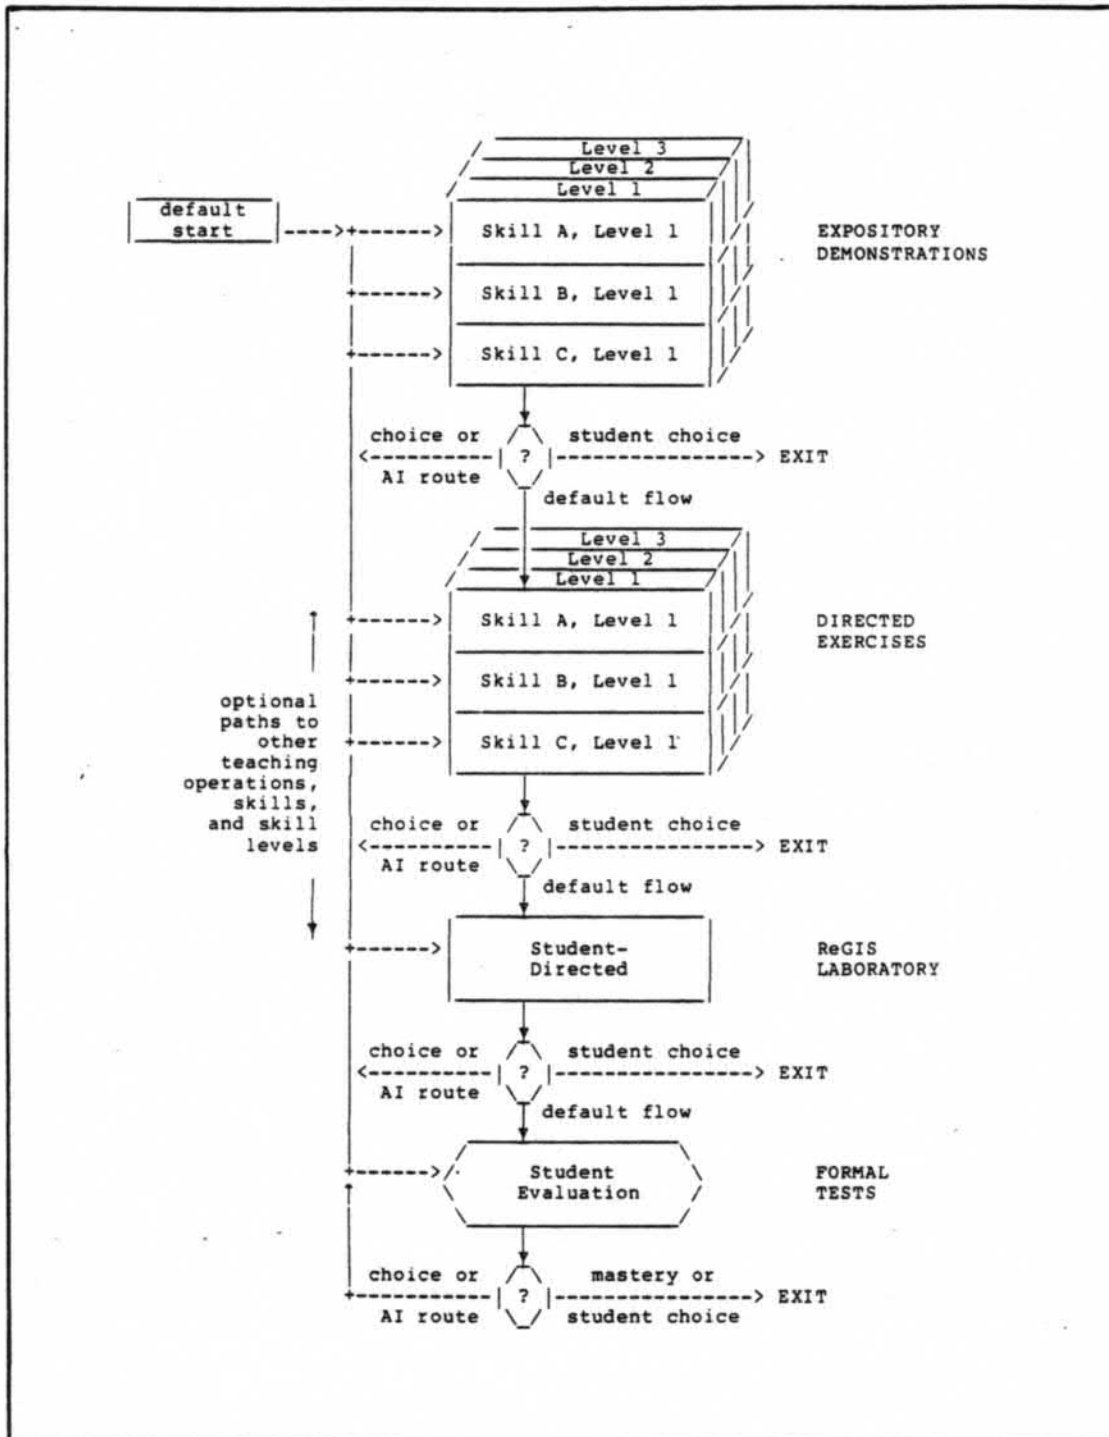


Figure 1

RELATIONSHIPS BETWEEN TEACHING OPERATIONS

AI comes into play in these exercises as student responses are parsed. Each skill demonstrated by the student is recorded in a response history file, and the student model is updated accordingly. This data can then be used to point out inconsistencies in the student's thinking if s/he makes an error using the same skill in a later lesson.

The "directed" nature of the exercises refers to using the student's performance data to determine the type and difficulty of problems presented. This feature contributes to the program's response-sensitivity. Students who do well will find that the problems get harder quite quickly. Weaker students will be led along more slowly, making sure that they possess each component skill before higher level skills are presented.

The initial exercise type and difficulty level depends on three factors:

- (1) whether the student has gone through the expository demonstration,
- (2) the reason the student is going through the exercises (by choice or as remediation caused by problems with higher level material), and
- (3) the student's own statement of what s/he thinks s/he knows.

Students will normally exit an exercise section by demonstrating mastery or non-mastery of the skills being practiced. (They may also exit by pressing an EXIT key or entering EXIT in response to a question.) If they demonstrate non-mastery, one of two things may happen depending upon their state as defined by the student model.

- (1) If the student has already demonstrated mastery on each of the prerequisite skills for this module, s/he will be branched to a set of secondary teaching operations on the current material. These teaching operations will represent the material in a more detailed manner or from another perspective. This rule implies that there will be some instructional modules that are only accessed for remedial purposes. Students who go through the course grasping concepts the first time around will never see these modules.
- (2) If the student has not demonstrated mastery on one of the prerequisite modules, s/he will be directed to study (or restudy) that module before attempting the current module again.
- (3) If there are more than one prerequisite modules on which the student has not demonstrated mastery, the system will

ask a number questions to try to identify which prerequisites are most like missing. This is done by analyzing the task representation in relation to the student model.

5.3 ReGIS Laboratory

The third teaching operation is intended to provide the greatest amount of response-sensitivity and demonstrate the course's most sophisticated AI aspects. These characteristics can be accomplished by providing students with a ReGIS "sketchpad" or "laboratory" environment in which they can enter ReGIS commands and see the results of these commands directly. The lab will be implemented in a dual screen format with the ReGIS code appearing on one side and the graphic output on the other.

The AI component here involves "watching" students as they enter commands, updating the student model for each skill demonstrated, and looking for cliché errors in their code. These techniques are similar to those of Burton and Brown (1982) and Shrager and Finin (1982) described in AI/CAI Project Progress Report No. 2. The prototype system would not attempt to offer any real "coaching" à la WEST (Burton and Brown, 1982), but would provide detailed error messages for syntactic and simple theoretical errors similar to those in the directed exercises.

The difference between the directed exercises and ReGIS lab is degree of computer control. If effectively implemented, measures of their instructional effectiveness should yield equivalent results. A major difference between the two styles, however, is that it is more difficult to identify missing prerequisites in a laboratory environment because there is no mechanism for asking direct questions. Failure to demonstrate mastery in this environment will therefore be coupled with more conventional directed exercises.

5.4 Formal Tests

When a student demonstrates a particular skill in either the directed exercises or ReGIS laboratory, the probability with which s/he actually possesses that skill is somewhat less than 1.0. For example, a student may "discover" defaults by leaving out an X or Y value for the P command in the ReGIS laboratory. No error message would be generated, and the student may not even realize what has happened until some time later. While this situation represents an excellent (and some would say the best) learning scenario, it is impossible to know for certain whether the student has actually internalized the concept s/he has just demonstrated.

Evaluation of actual learning requires a formal testing situation. The administration of formal tests could be very similar to the methods employed for directed exercises, with the stipulation that part of the feedback would be eliminated. Formal testing need not be limited to multiple choice and short answer responses.

6.0 THE TASK REPRESENTATION

6.1 List of Component Skills

Table 1 lists the component skills needed to master all aspects of the ReGIS P, C, and V commands. Each of these skills represents the course's smallest possible instructional unit. That is, the course's AI component will be designed to identify skills that a student lacks and route him or her to the specific teaching operations on those skills.

The component skills are grouped into modules for presentation to students going through the course for the first time. This does not mean that remedial work on one skill in a module necessitates redoing the other skills in that module. It simply means that the skills are presented together the first time for the sake of continuity and organizational convenience. The skill groupings for each module are shown in Table 2.

Skills 35 and 43 will not actually be taught in the course. These skills involve understanding angle measure in degrees and interpolation, respectively, and are what Robert Mager and Peter Pipe (1974) call "entry level objectives". They are required for students to master higher level objectives, but students are expected to bring these skills to the course rather than learn them from the course.

6.2 Prerequisite Relationships: Directed Graph

A number of prerequisite relationships exist between the 50 component skills. Such prerequisites indicate those skills a student must possess in order to master higher level skills. Figure 2 shows the prerequisite relationships represented by a directed graph. The graph should be interpreted as follows.

- Nodes inside shaded squares represent skills with no prerequisites. These are Skills 1, 25, 35, and 43.
- The skill hierarchy moves from the bottom upward. Lines connecting two skills indicate that the higher level skill requires the lower level skill below it as a prerequisite. For example, Skill 2 requires Skill 1 as a prerequisite.
- When two or more lines lead into a node from the bottom, all of the lower level skills are required as prerequisites. For example, Skill 5 requires both Skills 3 and 4 as prerequisites.
- When two or more lines lead out of a node from the top, the lower level skill is a prerequisite for each of the

Table 1
 LIST OF COMPONENT SKILLS

Number	Description
1	Recognizes screen as a rectangular dot matrix.
2	Can translate screen positions into (x,y) pairs.
3	Can translate (x,y) pairs into screen positions.
4	Knows absolute screen limits (767,479).
5	Knows that initial cursor position is upper left-hand corner of screen.
6	Understands the concept of current cursor position.
7	Can interpret the standard [x,y] address format.
8	Can specify absolute screen addresses in [x,y] format.
9	Understands defaults.
10	Given the current cursor position as [xc,yc], knows the meaning of:
11	[x] -> [x,yc]
12	[,y] -> [xc,y]
13	[] -> [xc,yc]
14	Understands relative addresses.
15	Given the current cursor position as [xc,yc], knows the meaning of:
16	[+x,+y] -> [xc+x,yc+y]
17	[+x] -> [xc+x,yc]
18	[,+y] -> [xc,yc+y]
19	[+x,y] -> [xc+x,y]
20	[x,+y] -> [x,yc+y]
21	Can use the basic P command to position the cursor.
22	Knows the current cursor position after execution of a P command.
23	Can use all addressing schemes in P commands.
24	Can work with P command aberrations such as P [x1,y1] [x2,y2] ... [xn,yn].
25	Can store addresses on the stack with (B).
26	Can pop addresses off the stack with (E).
27	Can use the S(E) command to erase the screen.
28	Can use the basic V command to draw a vector.
29	Knows the current cursor position after execution of a V command.
30	Can draw multiple vectors with a single V command.
31	Can draw closed polygons using (B) and (E) in V commands.
32	Knows the terms "open" and "closed" as applied to figures.

(continued...)

Table 1 (continued)
LIST OF COMPONENT SKILLS

Number	Description
31	Can use the basic C command to draw circles.
32	Knows the current cursor position after drawing a circle with the basic C command.
33	Can use the C(C) command to draw circles.
34	Knows the current cursor position after drawing a circle with the C(C) command.
35	Understands angle measure in degrees.
36	Can translate angle measures into screen angles.
37	Can translate screen angles into angle measures.
38	Can use the C(A<degrees>)[X,Y] command to draw arcs starting a [X,Y] and using the current cursor position as the center.
39	Knows the current cursor position after drawing an arc with the C(A<degrees>)[X,Y] command.
40	Can use the C(A<degrees>C)[X,Y] command to draw arcs starting at the current cursor position and using [X,Y] as the center.
41	Knows the current cursor position after drawing an arc with the C(A<degrees>C)[X,Y] command.
42	Can draw open curves with the C(S) command.
43	Understands interpolation.
44	Can use [] at the front of a C(S) command.
45	Can use [] at the end of a C(S) command.
46	Knows the current cursor position after execution of a C(S) command.
47	Can draw closed curves with the C(B) command.
48	Can use [] at the front of a C(B) command.
49	Can use [] at the end of a C(B) command.
50	Knows the current cursor position after execution of a C(B) command.

Table 2
GROUPINGS OF SKILLS INTO MODULES

Module	Component Skills
A	1 2 3 4
B	5 6 7 8
C	9 10 11 12
D	13 14 15 16 17 18
E	19 20 21 22
F	23 24 25
G	26 27 28 29 30
H	31 32 33 34
I	35 36 37 38 39
J	40 41
K	42 43 44 45 46
L	47 48 49 50

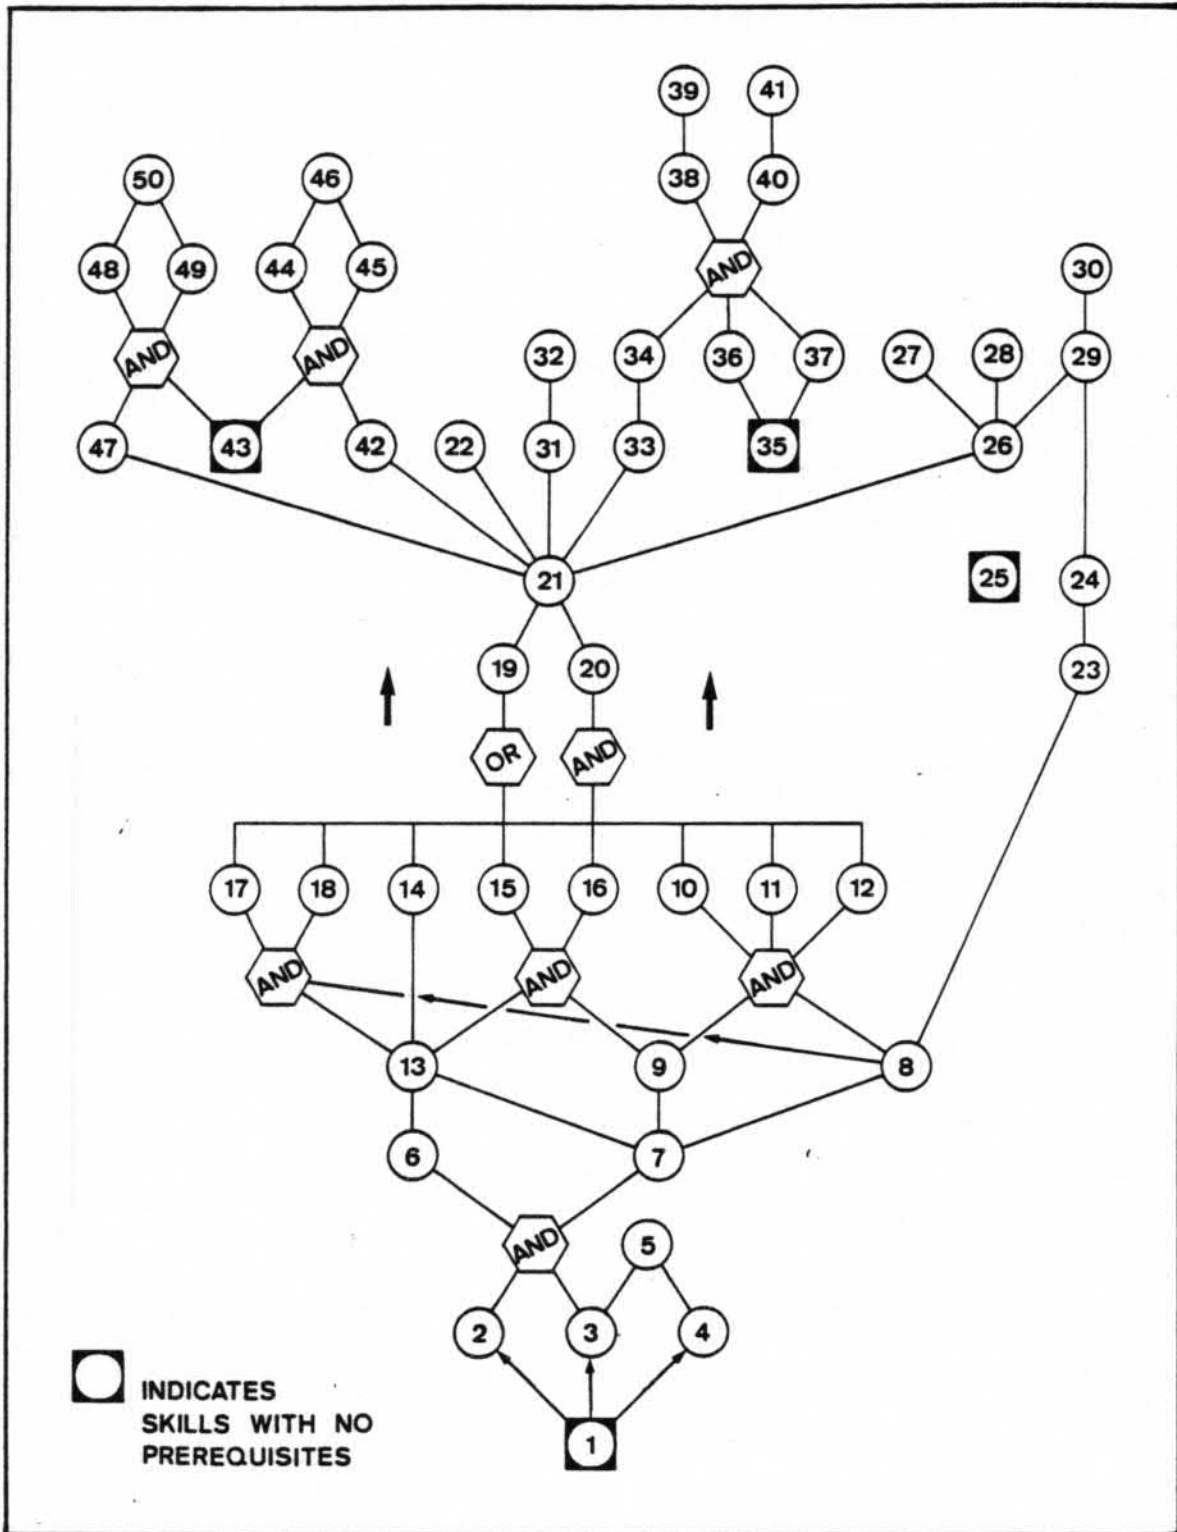


Figure 2

TASK MODEL REPRESENTED AS A DIRECTED GRAPH

higher level skills. For example, Skill 25 is a prerequisite for Skills 27, 27, and 28.

- The AND hexagon is not a node. It indicates that all of the lower level skills are required as prerequisites for each of the higher level skills. For example, Skills 6 and 7 each require both Skills 2 and 3 as prerequisites.
- The OR hexagon is not a node. It indicates that any one of the lower level skills is sufficient prerequisite for each of the higher level skills. For example, Skill 19 requires any one of Skills 10 through 18 as a prerequisite.

(The AND and OR hexagons were used to keep the graph readable by avoiding a large number of crossing lines.)

The prerequisite relationships of the component skills dictate the prerequisite relationships between the modules. These prerequisites are shown in Figure 3.

6.3 Prerequisite Relationships: Production Rules

The directed graph presented in the previous section can be defined by the facts listed in Table 3. This table should be interpreted as follows.

- NIL in the left-hand column indicates that no prerequisites are required for the corresponding skills in the right-hand column. For example, no prerequisites are required for Skills 1, 25, 35, and 43.
- If more than one skill is listed on a single line in the left-hand column, all of those skills are required as prerequisites for each of the skills listed in the right-hand column. (This is the AND function.) For example, Skills 3 and 4 are both required as prerequisite for Skill 5.
- If more than one skill is listed on a single line in the right-hand column, all of the skills listed in the corresponding left-hand column are required as prerequisites for each of the skills listed in the right-hand column. For example, both Skills 6 and 7 require Skills 2 and 3 as prerequisites.
- If a skill appears on more than one line in the left-hand column, that skill is required as a prerequisite for more than one higher level skill. For example, Skill 3 is required as a prerequisite for Skills 5, 6, and 7.

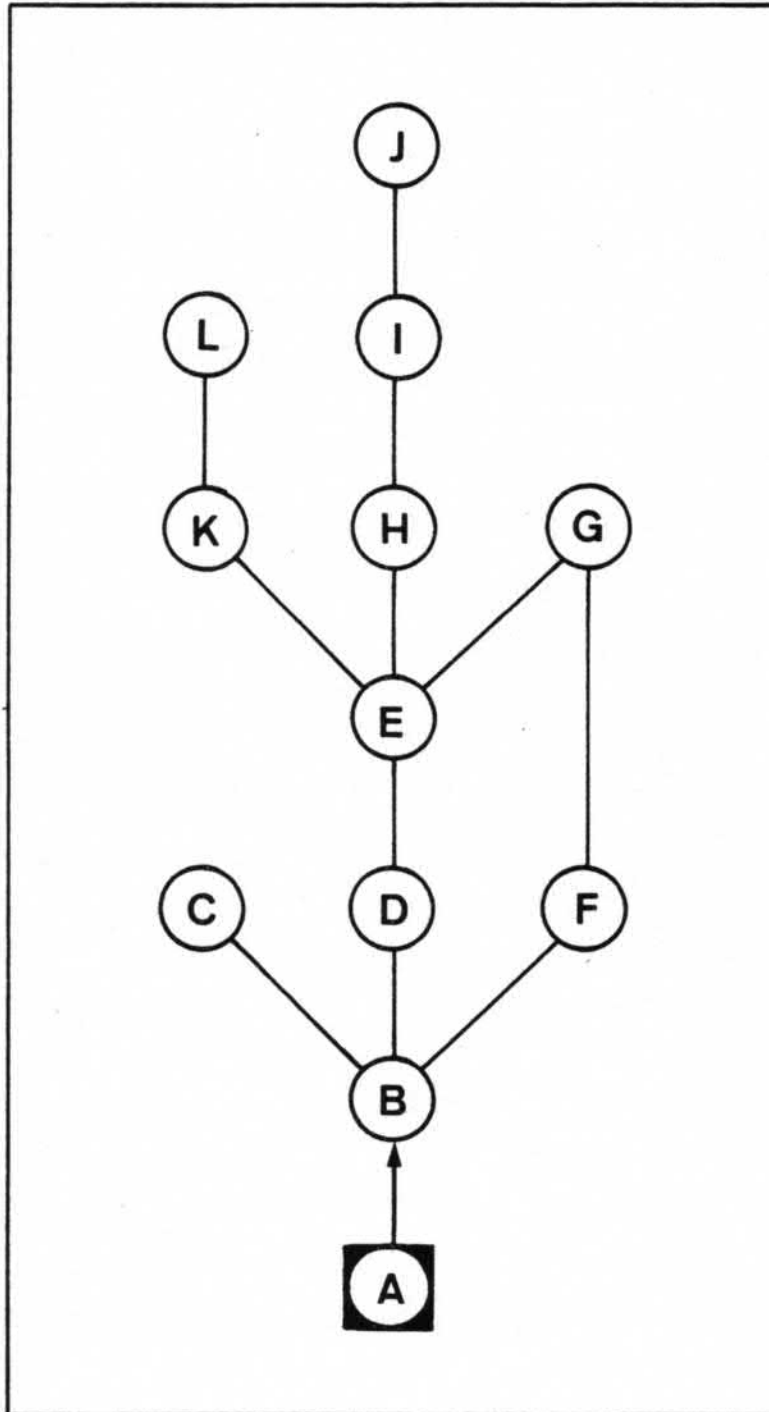


Figure 3

PREREQUISITE RELATIONSHIPS BETWEEN MODULES

Table 3

TASK MODEL REPRESENTED AS PRODUCTION RULES

if the student has mastered these skills ...	s/he has met the prerequisites for these skills ...
NIL	1 25 35 43
1	2 3 4
2 3	6 7
3 4	5
6 7	13
7	8 9
8 9	10 11 12
8 13	17 18
8	23
9 13	15 16
10 11 12 13 14 15 16 17 18	20
10	19
11	19
12	19
13	14
13	19
14	19
15	19
16	19
17	19
18	19
19 20	21
21	22 26 31 33 42 47
23	24
24 26	29
26	27 28
29	30
31	32
33	34
34 36 37	38 40
35	36 37
38	39
40	41
42 43	44 45
43 47	48 49
44 45	46
48 49	50

- If a skill appears on more than one line in the right-hand column, any of the corresponding left-hand columns provides sufficient prerequisites for that skill. (This is the OR function.) For example, Skill 19 requires any one of Skills 10 through 18 as a prerequisite.

Using these facts, the prerequisite relationships (and thus the entire representation of the task) can be defined by production rules. The full set of production rules defines the program's task model, which can be represented in LISP as a list of lists. For example, the first three rules are represented as:

```
(NIL (1 25 35 43))  
(1 (2 3 4))  
((2 3) (6 7))
```

One should take particular note of how easy it is to represent both the AND and OR functions in this formalism. For example, the rule:

```
((6 7) 13)
```

indicates that both Skills 6 and 7 are required as prerequisites for Skill 13 (the AND function). When taken together, the rules:

```
(10 19)  
(11 19)  
(12 19)
```

indicate that any one of Skills 10, 11, or 12 is sufficient as a prerequisite for Skill 19 (the OR function).

The production rule formalism makes it conceptually simple to identify both the skills for which the student has met the prerequisites and the prerequisites needed to study any particular skill. For example, suppose that a rudimentary student model consists of a simple list of the skills that a particular student has mastered. Such a list might be:

```
(1 2 3 7 8 9 12)
```

A function can then be written that steps down the list of task model rules, testing whether each left-hand side (LHS) is a perfect subset of the student model. If it is, the student has met the prerequisites for the skills listed on the right-hand side (RHS) of that rule. For the example list of skills shown above, this function would return:

```
(4 6 10 11 19 23 25 35 43)
```

For a fuller understanding of why this is so, compare this list to the graph in Figure 2. The student is ready for:

- Skill 4 because the list of skills mastered includes Skill 1,
- Skill 6 because it includes both Skills 2 and 3,
- Skills 10 and 11 because it includes both Skills 8 and 9,
- Skill 19 because it includes Skill 12 (only one of Skills 10-18 is required for Skill 19, but note that the student is not ready for Skill 20, because that skill requires all Skills 10-18),
- Skill 23 because it includes Skill 8, and
- Skills 25, 35, and 43 because these have no prerequisites.

The discussion thus far has concerned moving up the directed graph to answer the question: "given a specific set of mastered skills, which skills is the student now ready to study, i.e., for which skills does the student now possess the prerequisites?" The beauty of the production rule approach is that the same representation can be used equally well to move down the directed graph and answer the converse question: "given a specific skill, which prerequisite skills must the student possess to be ready to study it?" This characteristic is crucial to achieving the AI diagnostic qualities of the directed exercises and ReGIS laboratory discussed in Sections 5.2 and 5.3.

Suppose, for example, that the student hadn't really studied all of the skills specified by the list:

(1 2 3 7 8 9 12)

Instead, s/he may have actually only studied Skill 12. By virtue of demonstrating mastery on that skill, the system's AI component would update its student model by marking the student's mastery of all the skills prerequisite for Skill 12 as "assumed". To determine which skills to mark, a LISP function can be written that tests the RHS of each rule. If the skill just mastered is a member of the list of RHS skills, mastery of each of the skills listed on the LHS is assumed.

In practice, the function would be called with the number of the skill just mastered and a list of skills representing the student model. The function would then return a list of RHS skills that are not already members of the student model. If, for example,

the student model list is empty (NIL), calling this function with "12" as an argument would return the list:

(1 2 3 7 8 9)

If the student model already indicates mastery on Skills 1 and 3, calling it with "12" as an argument would return:

(2 7 8 9)

7.0 THE STUDENT MODEL

7.1 Skill Status

The basic purpose of a student model is to represent a student's current knowledge state. In its simplest form, this state might be defined by the student's status on each of the skills in the task model. As indicated in the previous section, one rudimentary way to do this is to maintain a list (or state vector) of those skills on which the student has demonstrated mastery. The utility of this list can be greatly improved, however, by letting the status of each skill take on a number of values. The student model employed in the prototype course will use the following six values.

- 3 NON-MASTERY DEMONSTRATED on a test. The student has demonstrated that s/he does not possess this skill by failing a test that covered it. This is the strongest assertion of non-mastery that the system can make.
- 2 NON-MASTERY ASSUMED due to incorrect usage in lab. The student is assumed not to possess the skill because s/he has failed a test that covers a skill prerequisite to the one in question.
- 1 NON-MASTERY ASSUMED due to incorrect usage of a prerequisite skill. The student is assumed not to possess the skill because s/he has either demonstrated non-mastery on or used incorrectly a lower level skill for which this skill is a postrequisite. (Note that the skill in question may be more than one level removed from the lower level skill on which the student is actually working.) This is the weakest assertion of non-mastery that the system can make.
- 0 NO DATA. The student has not studied this skill, has not demonstrated mastery on any skill for which it is a prerequisite, and has not demonstrated non-mastery on any skill for which it is a postrequisite.
- 1 MASTERY ASSUMED due to correct usage of a postrequisite skill. The student is assumed to possess the skill because s/he has either demonstrated mastery on or used a higher level skill for which this skill is a prerequisite. This is the weakest assertion of mastery that the system can make.
- 2 MASTERY ASSUMED due to correct usage in lab. The student is assumed to possess the skill because s/he used the skill in either the ReGIS laboratory or the directed exercises.

- 3 MASTERY DEMONSTRATED on a test. The student has demonstrated mastery of this skill by passing a test that covered it. This is the strongest assertion of mastery that the system can make.

The student model value for each skill is initialized to 0 when the student registers. As s/he works through the course, one of the non-zero values is assigned to each skill on which the system has or can infer data. These values add a level of complexity to the functions discussed in the Section 6.3, in that analysis of the production rules cannot be done simply by testing for the presence of a skill number in a list. The complication is not extreme, however, and should present no serious implementation problems.

It is also possible to express a student model in terms of procedures rather than a state vector. See Self, 1974, for a discussion of this technique.

7.2 Learning Rate and Learning Style

In addition to a student's skill status, the student model can also maintain two simple and rudimentary representations of the student's learning rate and learning style. Learning rate is a measure of the student's ability to assimilate new material quickly. Learning style is a measure of the manner in which the student prefers new material to be presented.

The student's learning rate would govern the speed, depth, and amount of repetition and reinforcement in initial presentations of new material. Fast students would receive fast presentations extending to considerable depth before going into the directed exercises as reinforcers. Slower students would be presented with more detailed introductions to new material at lower levels, and would find more repetition in the presentations as well as more frequent reinforcement via the directed exercises. Learning rate might be expressed as one of the following five values:

- VERY-FAST
- FAST
- AVERAGE
- SLOW
- VERY-SLOW

The student's preferred learning style might be represented by one of three values:

- EXPOSITORY -- the student prefers to go through the full expository demonstration before doing exercises.
- EXERCISE -- the student prefers to dive right into the directed exercises.

- LABORATORY -- the student prefers to try things out in the ReGIS laboratory after a short explanation of pertinent concepts and commands.

A number of techniques exist for assessing learning style, but a finesse is also feasible: simply ask students which style they prefer. Students would be allowed to change their learning style preference as the course progresses, as well as override the default selection for any particular module. The system might monitor the number of overrides and change the default learning style when this number becomes significant.

8.0 THE MEANS-ENDS GUIDANCE RULES

Means-ends guidance rules relate states defined by the student model and student history to specific teaching operations and determine which instructional activities the CAI program will present next given different student states. A rudimentary student history could be as simple a list of all responses entered by the student. In practice, this history might also flag responses to non-subject matter queries as "choices" made by the student, e.g., his or her selections when presented with a number of options on a menu.

Sample means-ends guidance rules (in plain English format) might be as follows:

1. If the student is entering a module for the first time, make the initial subject matter presentation in the form specified by the LEARNING-RATE and LEARNING-STYLE elements of the student model.
2. If the student is reentering a module s/he has already studied and done well on, query him or her as to what skills s/he wishes to study and in what learning style. (The response to these queries would be recorded in the student history as "choices".)
3. If the student is reentering a module s/he has already studied but done poorly on, make subject matter presentations in EXPOSITORY style on all skills for which the student model indicates NO-DATA or NON-MASTERY, and make these presentations as if the value of LEARNING-RATE was SLOW or VERY-SLOW. (This would force more repetition and reinforcement.)
4. If the student demonstrates non-mastery on a specific skill and the student model indicates that there are no prerequisites for that skill on which mastery has not been demonstrated or assumed (that is, mastery has been demonstrated or assumed on all the prerequisites), branch to a secondary teaching operation for that skill if one exists. (Secondary teaching operations are ones that are designed for remediation only, not initial presentation.) If no secondary teaching operations exist for the skill in question, apply Rule 3 above. (The bulky negative wording in the IF clause was used to make this rule consistent with Rules 5 and 6.)
5. If the student demonstrates non-mastery on a specific skill and the student model indicates that there is only one prerequisite for that skill on which mastery has not been demonstrated or assumed, apply Rule 1 to the module containing that prerequisite skill.

6. If the student demonstrates non-mastery on a specific skill and the student model indicates that there are more than one prerequisites for that skill on which mastery has not been demonstrated or assumed, apply Rule 1 to the module containing the prerequisite s/he is "most likely" lacking. (The system would determine which skill is "most likely" lacking by analyzing the student model values for other skills with the same prerequisites.)

8.1 Representation of Left-Hand Sides

The left-hand sides (LHSs) of these rules (the IF parts) represent specific patterns to be matched against data derived from the student model and student history. For Rules 1, 2, and 3, these data would include the status of the module the student has chosen to study and the statuses of each of that module's sub-skills. The LHSs of these rules might therefore take the form:

1. ((STATUS-OF-MODULE-CHOSEN EQUAL 0)
(ALL-SUBSKILL-STATUSES EQUAL 0))
2. ((STATUS-OF-MODULE-CHOSEN (NOT EQUAL) 0)
(AVERAGE-SUBSKILL-STATUS GREATER-THAN 0))
3. ((STATUS-OF-MODULE-CHOSEN (NOT EQUAL) 0)
(AVERAGE-SUBSKILL-STATUS LESS-THAN 0))

The function calls (STATUS-OF-MODULE-CHOSEN EQUAL 0) and (STATUS-OF-MODULE-CHOSEN (NOT EQUAL) 0) would return TRUE if the status of the module chosen is equal to or not equal to 0, respectively. The function call (ALL-SUBSKILL-STATUSES EQUAL 0) would operate on sets of skills, and return TRUE if each of those skills has a status value equal to 0. Likewise, the function calls (AVERAGE-SUBSKILL-STATUS GREATER-THAN 0) and (AVERAGE-SUBSKILL-STATUS LESS-THAN 0) would return TRUE if the average subskill status value is greater than or less than 0, respectively. When the values of all function calls on the LHS of a rule are TRUE, that rule fires.

For Rules 4, 5, and 6, the patterns to be matched would include the status of the particular skill just studied and the statuses of each of that skill's prerequisite skills. The LHSs of these rules might therefore take the form:

4. ((SKILL-STATUS LESS-THAN 0)
(NO-PREREQ-SKILL-STATUSES LESS-THAN 0))
5. ((SKILL-STATUS LESS-THAN 0)
(ONLY-ONE-PREREQ-SKILL-STATUS LESS-THAN 0))

6. ((SKILL-STATUS LESS-THAN 0)
(MORE-THAN-ONE-PREREQ-SKILL-STATUS LESS-THAN 0))

8.2 Representation of Right-Hand Sides

The right-hand sides (RHSs) of the rules (the THEN parts) could be expressed as a TEACH function with the form:

```
(TEACH (MODULE-ID
        SKILL-ID
        LEARNING-RATE
        LEARNING-STYLE
        SEARCH-STRATEGY))
```

where

- MODULE-ID represents the module to be entered, as identified in Table 2.
- SKILL-ID represents the skill to be taught, as identified in Table 1.
- LEARNING-RATE represents the amount of repetition and reinforcement to use in presentations.
- LEARNING-STYLE represents which of the three types of teaching operations to use.
- SEARCH-STRATEGY represents the manner in which this teaching operation was selected, e.g., REMEDIAL or MOST-NEEDED.

Using this format, the RHSs of the sample rules could be expressed as follows (the asterisk is a wild card that matches any value of the function argument list in the corresponding position):

1. (TEACH (MODULE-CHOSEN
*
LEARNING-RATE
LEARNING-STYLE
*))
2. (TEACH (MODULE-CHOSEN
QUERY
FAST
QUERY
*))

3. (TEACH (MODULE-CHOSEN
(SKILLS-WITH-STATUSES (LESS-THAN OR EQUAL) 0)
SLOW
EXPOSITORY
*))

4. (TEACH (MODULE-CONTAINING-SKILL
SKILL
LEARNING-RATE
*
REMEDIAL))

OR
(TEACH (MODULE-CONTAINING-SKILL
SKILL
SLOW
EXPOSITORY
*))

5. (TEACH (MODULE-CONTAINING-PREREQUISITE-SKILL
*
LEARNING-RATE
LEARNING-STYLE
*))

6. (TEACH (MODULE-CONTAINING-PREREQUISITE-SKILL
*
LEARNING-RATE
LEARNING-STYLE
MOST-NEEDED))

The RHS for Rule 4 has two TEACH functions to cover the case in which no secondary teaching operations exist for a specific skill. This representation makes the rules more complex, but provides explicit definition of what to do if a TEACH function request cannot be filled. Another way to tackle this problem is to use only one TEACH function in each rule but check whether the rule succeeds after the RHS fires. If the rule does not succeed, the system must continue looking for another rule whose LHS matches the data in the student model and student history.

9.0 CRITIQUE AND CONCLUSIONS

The rule-based tutorial described in this document has not undergone the test of implementation. However, the approach described here is a conceptually clean extension of a working computer tutor that uses production rules in the teaching of quadratic equations (O'Shea, 1979). On implementation, some of the details of the formalism described here would probably have to change to ensure computational efficiency and to maintain reasonable response time in the particular interactive computer environment adopted.

For example, test implementations in MacLISP on a DECsystem-20 have shown that while the production rule formalism is highly efficient for expressing the prerequisite relationships as shown in Table 3, use of this formalism to update the student model after each response is relatively inefficient. The main inefficiency stems from processing the rules repeatedly to find all of the prerequisites or postrequisites for the skill on which the student is currently working. This problem can be attacked by computing all of the prerequisites and postrequisites when the course is installed and storing these as lists in a simple array. This approach allows the student model to be updated much more quickly without sacrificing the elegance of the rule-based strategy.

Test implementations have also shown that the production rule approach can be applied to interpreting student responses. The parser that interprets command strings entered by the student in the ReGIS laboratory can be made to return specific patterns, which can be related to the skills in the task model via production rules. For example, consider the command:

```
P[250,100]
```

to position the cursor 250 pixels from the left-hand margin and 100 pixels down from its current position. Given this command, the parser returns:

```
((COMMAND POSITION) (TYPE POINT)
 (X-VALUE-TYPE ABSOLUTE) (X-VALUE 250)
 (Y-VALUE-TYPE RELATIVE) (Y-VALUE-SUBTYPE +)
 (Y-VALUE 100))
```

This result can be related to the skills in the task model with the following rules (these are only a subset of the full set of address diagnostic interpretation rules):

```
(( (X-VALUE-TYPE ABSOLUTE) (Y-VALUE-TYPE ABSOLUTE)) 8 19)
(( (X-VALUE-TYPE ABSOLUTE) (Y-VALUE-TYPE DEFAULT)) 10 19)
(( (X-VALUE-TYPE DEFAULT) (Y-VALUE-TYPE ABSOLUTE)) 11 19)
(( (X-VALUE-TYPE DEFAULT) (Y-VALUE-TYPE DEFAULT)) 12 19)
```

```
((X-VALUE-TYPE RELATIVE) (Y-VALUE-TYPE RELATIVE)) 14 19)
((X-VALUE-TYPE RELATIVE) (Y-VALUE-TYPE DEFAULT)) 15 19)
((X-VALUE-TYPE DEFAULT) (Y-VALUE-TYPE RELATIVE)) 16 19)
((X-VALUE-TYPE RELATIVE) (Y-VALUE-TYPE ABSOLUTE)) 17 19)
((X-VALUE-TYPE ABSOLUTE) (Y-VALUE-TYPE RELATIVE)) 18 19)
```

For each rule whose LHS is a perfect subset of the result returned by the parser, the program updates the student model by:

- (1) assigning a value of +2 to each the skill listed on the rule's RHS, and then
- (2) assigning a value of +1 to each of the prerequisites for each of those skills.

Thus the rule-based tutorial might employ several sets of rules for different functions. The advantage of this approach is that all such sets are easily changed to "tune" the course and enhance its response-sensitivity.

Production rule programming is not a trivial task, and one line of development being pursued is a rule-based authoring system which makes it easy for educational designers without substantial programming skills to enter and change sets of course related production rules (O'Shea et al., 1983). However, we contend that even without access to such a system it is still more effective to build CAI programs by identifying tutorial rules than to use conventional CAI authoring languages, because latter typically exhibit restrictive orientations toward automating programmed learning texts via the clumsy apparatus of frames, branches, and multiple choice questions.

In conclusion, we have shown how the rule-based approach can be usefully applied to the design of a course that includes exposition, directed exercises, a simulated laboratory, and tests. In the resulting course, the various teaching operations are modular and distinct, as are the production rules used in the student model for response-sensitivity and as means-ends guidance rules for scheduling the presentation of teaching operations. It is therefore possible, for example, to integrate new teaching operations into the course while maintaining the general level of response-sensitivity by adding new production rules. Likewise, any increase in response-sensitivity achieved in the student model will be applied to all teaching operations. We believe that these rule-based techniques represent an efficient and elegant approach to the task of designing and implementing CAI tutorials.

10.0 RESEARCH EFFORTS TO BE PURSUED

The next step in my research is to try to implement at least part of the course described in this document. The functionality described here is extremely ambitious, and its implementation would require considerably more effort than is warranted for a prototype. I have no intention of trying to implement the entire course, but the prototype should get far enough to encounter and tackle all of the major implementation problems and to demonstrate functionality at a reasonable level of sophistication.

The prototype course will be implemented initially in LISP. I have already implemented part of the ReGIS Laboratory under MacLISP on the DECsystem-20 at The Open University, and I hope to be able to get this software running on VAX by the end of June, 1983. Implementation of the full prototype will probably take about half of next fiscal year.

Once the prototype is up and running, I intend to take a look at the problem of converting the program to a small system such as the PRO-350. Such a conversion may prove extremely difficult, but we must find out whether trade-offs can be made that make implementation of these techniques at all feasible on the IVIS-type systems we see as the major courseware delivery vehicles of the future.

11.0 REFERENCES CITED AND RELATED READINGS

- Burton, Richard R., and John Seely Brown, 1982. "An Investigation of Computer Coaching for Informal Learning Activities." In Intelligent Tutoring Systems, ed. D. Sleeman and J.S. Brown, pp. 79-98. New York: Academic Press.
- Hartley, J.R., 1973. "The Design and Evaluation of an Adaptive Teaching System," International Journal of Man-Machine Studies, 5(2).
- Mager, Robert F., and Peter Pipe, 1974. Criterion-Referenced Instruction: Analysis, Design, and Implementation. Mager Associates, Los Altos Hills, CA.
- O'Shea, Tim, 1979. "Self-Improving Teaching Systems." Ph.D. Dissertation, University of Leeds. Basel, Boston, Stuttgart: Birkhauser.
- O'Shea, Tim, R. Bornat, B. du Boulay, M. Eisenstadt, and I. Page, 1983. "Tools for Designing Intelligent Computer Tutors," in Human and Artificial Intelligence, ed. A. Elithorn and R. Banerjii, London: North-Holland.
- Self, John, 1974. "Student Models in Computer-Aided Instruction," International Journal of Man-Machine Studies, 6:261-276.
- Shrager, Jeff, and Tim Finin, 1982. "An Expert System that Volunteers Advice." Proceedings of the 1982 National Conference on Artificial Intelligence (sponsored by the American Association for Artificial Intelligence), Pittsburgh, Pennsylvania, August 18-20, 1982, pp. 339-340. Los Altos, CA: William Kaufmann, Inc.