

AD-A178 268

THE ROLE OF COGNITIVE SIMULATION MODELS IN THE
DEVELOPMENT OF ADVANCED TR. (U) MICHIGAN UNIV ANN ARBOR
TECHNICAL COMMUNICATION PROGRAM D E KIERS 15 JAN 87

1/1


UNCLASSIFIED

TR-87/DNR-23 N00014-85-K-8138

F/G 5/10

NL

													1/1 A-1 D-10




AD-A178 268

12

The Role of Cognitive Simulation Models in the Development of Advanced Training and Testing Systems

David E. Kieras

University of Michigan



Technical Report No. 23 (TR-87/ONR-23)

January 15, 1987

This research was supported by the Personnel and Training Research Programs under Contract Number N00014-85-K-0138, Contract Authority Identification Number NR 667-543. Reproduction in whole or part is permitted for any purpose of the United States Government.

Approved for Public Release; Distribution Unlimited

DTIC FILE COPY

72 3

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1a REPORT SECURITY CLASSIFICATION Unclassified		1b RESTRICTIVE MARKINGS	
2a SECURITY CLASSIFICATION AUTHORITY		3 DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release: distribution unlimited.	
2b DECLASSIFICATION / DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S) TR-87/ONR-23		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a NAME OF PERFORMING ORGANIZATION University of Michigan	6b OFFICE SYMBOL (If applicable)	7a NAME OF MONITORING ORGANIZATION Cognitive Science Office of Naval Research (Code 1142CS) 800 N. Quincy Street	
6c ADDRESS (City, State, and ZIP Code) Technical Communication Program Ann Arbor, MI 48109-1109		7b ADDRESS (City, State, and ZIP Code) Arlington, VA 22217	
8a NAME OF FUNDING / SPONSORING ORGANIZATION	8b OFFICE SYMBOL (If applicable)	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-85-K-0138	
8c ADDRESS (City, State, and ZIP Code)		10 SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO 61153N	PROJECT NO RR04206
		TASK NO RR04206-0A	WORK UNIT ACCESSION NO. NR667-543
11. TITLE (Include Security Classification) The Role of Cognitive Simulation Models in the Development of Advanced Training and Testing Systems			
12. PERSONAL AUTHOR(S) David E. Kieras			
13a TYPE OF REPORT Technical Report	13b TIME COVERED FROM _____ TO _____	14. DATE OF REPORT (Year, Month, Day) January 15, 1987	15 PAGE COUNT 38
16 SUPPLEMENTARY NOTATION To appear in Frederiksen, Glaser, Lesgold, and Shafto (Eds.), <u>Diagnostic Monitoring of Skill and Knowledge Acquisition</u> . Hillsdale, N.J. Erlbaum			
17 COSATI CODES		18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	Training, task analysis, intelligent tutoring systems	
05	08		
19 ABSTRACT (Continue on reverse if necessary and identify by block number)			
<p>To be most efficient, an intelligent tutoring system must be based on a thorough analysis of the task that is being trained, and the knowledge, both declarative and procedural, that the learner must acquire in order to perform the task. This paper explores the thesis that constructing a cognitive simulation model of the task is an effective approach to characterizing this content. Because the simulation model must be explicitly stated in order to carry out the task properly, the model acts as a specification of (1) the exact nature of the user's task in terms of goals and subgoals; (2) the exact procedural knowledge required to accomplish the goals in the task setting; (3) the exact declarative knowledge required to support the procedural knowledge in carrying out the task. Such specifications could be used to both specify the content of instruction and the content of diagnostic tests.</p> <p style="text-align: right;">(continued on reverse)</p>			
20 DISTRIBUTION / AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21 ABSTRACT SECURITY CLASSIFICATION	
22a NAME OF RESPONSIBLE INDIVIDUAL Susan Chipman		22b TELEPHONE (Include Area Code) (202) 696-4318	22c OFFICE SYMBOL

19. Abstract (continued)

Two examples of the applications of this approach are presented, both assuming that the procedural knowledge is represented in the form of production rules. In the first example, by determining what production rules were actually stated by the training materials for a word processing system, it became clear that certain procedures were not correctly stated. The second example concerned the training and testing materials used in experiments in which subjects learned the fictitious inner workings of a simple control panel system, and then had to infer how to operate the control panel to accomplish a simple goal under different situations. By comparing the content of training and testing to the exact production rules in a simulation model for the tasks, it was clear that both the training and testing materials were far less precise in their content than had been intended.

These examples show that the content of a simulation model can be used to evaluate existing training and testing material in terms of the adequacy of its content. Unanswered questions concern whether a simulation model could be effectively used to prepare training and testing material, and whether the approach can be applied on a practical scale. Although constructing the model would be extremely time-consuming, it is reasonable in the context of an intelligent tutoring system, since such systems require a detailed specification of the content of training, and this is a substantial portion of the content of the model.

ABSTRACT

To be most efficient, an intelligent tutoring system must be based on a thorough analysis of the task that is being trained, and the knowledge, both declarative and procedural, that the learner must acquire in order to perform the task. This paper explores the thesis that constructing a cognitive simulation model of the task is an effective approach to characterizing this content. Because the simulation model must be explicitly stated in order to carry out the task properly, the model acts as a specification of (1) the exact nature of the user's task in terms of goals and subgoals; (2) the exact procedural knowledge required to accomplish the goals in the task setting; (3) the exact declarative knowledge required to support the procedural knowledge in carrying out the task. Such specifications could be used to both specify the content of instruction and the content of diagnostic tests.

Two examples of the applications of this approach are presented, both assuming that the procedural knowledge is represented in the form of production rules. In the first example, by determining what production rules were actually stated by the training materials for a word processing system, it became clear that certain procedures were not correctly stated. The second example concerned the training and testing materials used in experiments in which subjects learned the fictitious inner workings of a simple control panel system, and then had to infer how to operate the control panel to accomplish a simple goal under different situations. By comparing the content of training and testing to the exact production rules in a simulation model for the tasks, it was clear that both the training and testing materials were far less precise in their content than had been intended.

These examples show that the content of a simulation model can be used to evaluate existing training and testing material in terms of the adequacy of its content. Unanswered questions concern whether a simulation model could be effectively used to prepare training and testing material, and whether the approach can be applied on a practical scale. Although constructing the model would be extremely time-consuming, it is reasonable in the context of an intelligent tutoring system, since such systems require a detailed specification of the content of training, and this is a substantial portion of the content of the model.



Availability Codes	
Dist	Avail and/or Special
A-1	

THE ROLE OF COGNITIVE SIMULATION MODELS
IN THE DEVELOPMENT OF ADVANCED TRAINING AND TESTING SYSTEMS

David E. Kieras

University of Michigan

The thesis of this paper is that cognitive simulation modelling is a way to obtain a specification of the knowledge required to do a task, or to evaluate the adequacy of a proposed specification of the body of knowledge. The actual representations of the knowledge contained in a simulation can be the input for intelligent tutoring and other advanced training and testing systems, and can assist in preparing higher quality traditional paper-based systems as well.

This paper is organized as follows: First will be presented a brief discussion of cognitive simulation modelling, with a focus on what kind of information is contained in a cognitive simulation. Second will appear a discussion of how constructing a cognitive simulation can assist in developing training and testing procedures for new tasks. The paper will conclude with an extended discussion of a couple of examples of potential applications of this approach. The first example concerns the acquisition of procedures for operating equipment. The second example concerns learning and using a mental model for a piece of equipment. A concluding section discusses some of the potential problems with this approach.

COGNITIVE SIMULATION

A cognitive simulation model is defined here as a computer program that realizes a theoretical idea about mental structures and processes. The computer program contains explicit representations of proposed mental processes and knowledge structures. Changes in the internal state of the model are supposed to represent changes in the internal state of the human mind, at some level of analysis. These programs most commonly involve symbol manipulation, rather than numeric calculation, and use languages such as LISP and programming techniques originally developed in artificial intelligence.

Goals of Cognitive Simulation

Traditionally, cognitive simulations have been developed in order to turn a vague set of ideas about mental processes into well-defined ideas, whose completeness, adequacy, and consistency could be confirmed by determining whether the program produces the correct behavior. Thus, constructing a simulation model is a

way to evaluate and make more specific a set of theoretical ideas about cognitive processes.

An important feature of simulation modelling is that it fits in well with the very detailed methodology that is currently available in computer-based psychological experimentation. The deep level of detail in a simulation model corresponds well to the level of detail in data collected with these modern methods. There has been a greater tendency in recent years for simulation models to be tested for empirical accuracy against data at this greater level of detail.

Thus, the simulation itself provides an explicit and detailed theoretical statement that can summarize a large body of data. A fairly new assertion about the simulation approach is that it can provide theory-based predictions and evaluations in practical situations. See Kieras (1985) for more discussion on this topic.

Cognitive Architecture

The linkage between the psychological theory and the simulation program is best when the theory and the program are both based on definite assumptions about the architecture of cognition, rather than consisting of an arbitrary collection of proposed processes in the psychological theory, and correspondingly arbitrary data structures and code in the simulation program. By working within the structure imposed by an explicit cognitive architecture, the simulation modeler is afforded some degree of protection against accounting for behavior simply by large quantities of ad hoc programming. By adhering to a cognitive architecture, the simulation program is forced to have a consistent and principled structure.

A currently popular cognitive architecture assumes that there is both declarative and procedural knowledge in the mind. Declarative knowledge is represented as either a set of propositions, or as a semantic network. Procedural knowledge is represented as a set of production rules, which are elementary IF-THEN statements. The production rules test for the presence of various conditions in the declarative knowledge representation, and either manipulate the declarative representation or produce behavior. Perhaps the best representative of this cognitive architecture is Anderson's work on cognitive skill (Anderson, 1976, 1983). This architecture has the advantage of supplying a uniform modular notation for both declarative and procedural knowledge; the proposition is the unit of declarative knowledge; the production rule is the unit of procedural knowledge. These units can be counted, to yield quantitative predictors of performance (e.g., Kieras & Bovair, 1986).

The Information in a Cognitive Simulation

Once a cognitive simulation model has been constructed, and appears to behave correctly in the task of interest, there is in fact a considerable amount of information built into the simulation model and in the detailed specifications of under what task conditions the model is supposed to work.

In order to specify a cognitive simulation model, the modeler must provide a detailed task analysis. This analysis is typically far more detailed than what is often meant by the phrase "task analysis". At a global level, before a simulation model can be made to work, the modeler must lay out in specific detail what goals and subgoals the system is supposed to accomplish. In addition to this goal structure, the task analysis provides a detailed and quite explicit description of what the person must do in response to each individual specific aspect of the task situation. The implications of each individual relevant feature of the stimulus must be specified, along with what impact this is supposed to have on the person's cognitive processes and behavior. This detailed analysis is necessary in order to permit the simulation program to be written. Thus a successful simulation model is associated with a thorough analysis of the task.

Assuming the cognitive architecture mentioned above, a working simulation model contains an explicit description of the declarative knowledge required to do the task. This means that the modeler must lay out, again in detail, exactly what facts about the domain that the person must have in order to accomplish the task. In addition, the simulation model contains an explicit description of the procedural knowledge required to do the task; the model must contain the relevant rules, procedures, and heuristics.

As an example, consider a hypothetical simulation model that is able to carry out electronics trouble-shooting in a psychologically realistic fashion. Such a model will contain in it the answers to the following questions: (1) What is the goal structure of the trouble-shooting task? (2) What are the critical facts about electronics and the specific electronic system being repaired? (3) What are the rules for inferring the state of a component from various observations? (4) What are the strategies and heuristics used for isolating the malfunctioning component? If the simulation model is actually able to carry out the electronics trouble-shooting task successfully, then we can have some confidence that we have accurately and completely characterized the critical knowledge.

Using the Information in a Simulation

The task analysis and knowledge specifications from the simulation in fact characterize the task itself. This information has several potential uses. First, it can be used to evaluate the design of the system that a trainee must interact with. For example, if a system requires that the user have an extreme amount of knowledge before being able to operate it correctly, there must be something wrong with the design of the system. An improved design would simplify the training process. The approach to cognitive complexity in human-computer interaction taken by Kieras and Polson (1985) is based on this idea.

The second potential use of information in the simulation model is that it could be used to compare tasks in terms of the underlying subtasks. For example, if there were two related tasks, each with a simulation model, the specific knowledge and procedures contained in the simulations could then be compared directly to each other to precisely isolate what the tasks have in common. Thus instead of comparing tasks in terms of intuitions about their relationship, tasks could be compared in terms of their precise knowledge requirements. Such comparisons could be used to "optimize" the set of tasks assigned to individuals in particular jobs. That is, if how a set of tasks are related to each other could be precisely characterized, then it should be possible to pick combinations of tasks to assign to an individual that draw on the same body of knowledge. This would make effective training easier, and also would allow the individual to become more expert by being able to concentrate on a specific body of knowledge.

A third potential way to use the information is that the contents of the declarative and procedural knowledge representations in the simulation model provide an explicit specification of the minimum required knowledge to accomplish the task. Thus these representations specify what knowledge is really important, as opposed to optional elaboration or detail. Training and testing materials can be inspected to see if they actually contain this information. The examples below show how intuitively prepared material often seriously lacks critical information.

A fourth use is that the knowledge representations can be directly used in computer-based systems such as intelligent tutoring systems. Since the cognitive simulation models are typically constructed using the same AI techniques and concepts as are used in intelligent tutoring systems, it should be possible to make direct use of the declarative and procedural representations. Thus the effort spent in characterizing the task for purposes of cognitive simulation will be of direct use for the intelligent tutoring or testing system.

SOME EXAMPLE POTENTIAL APPLICATIONS

The Acquisition of Procedures

If one starts with the assumption that a good way to represent procedural knowledge is in terms of production rules that have a specified syntax and execution properties, then the psychological properties of a body of procedural knowledge should be revealed very directly by examination of the production rule representation for the knowledge. In work reported in Kieras and Bovair (1986), and Polson and Kieras (1985), it seems clear that such a detailed characterization of procedural knowledge in terms of production rules does carry considerable empirical content. In particular, the training time and amount of transfer of training can be accounted for with great precision, at least in some experimental paradigms, by considering the number of production rules that have to be learned, which depends on how many previously learned production rules can be applied in the current procedure being trained. Thus, training time and transfer effects can be predicted with some precision with this explicit representation. The focus of the present discussion however, is that the same characterization could function as a tool to look at training materials more carefully.

In work done by Kieras and Polson (1985, Polson and Kieras, 1985), production rule simulations of a person using a word processor have been constructed and compared against data. As a result, we have a fairly clear picture of the procedural knowledge required to operate a word processor. A natural step is then to examine training materials for a word processor to see if they specify a correct and running simulation. A couple of examples appear in the training materials for a commercial word processor. It should be noted that by current industry standards, these are extremely high-quality materials. Table 1 presents a portion of a procedure for deleting material on this word processor. Following each sentence is an informal translation of that sentence into production rules. Even though this is an informal translation, it is clear from the Kieras and Polson work how the rules could be made to actually run in a simulation.

If the reader is supposed to be acquiring production rules from this text, then a good training text should be one that conveys a correct and functional set of production rules as directly as possible. Notice in Table 1 that the sentence If the wrong characters are highlighted, press CODE+CANCL and try again corresponds very nicely to a production rule in terms of its sentence form. But, in examining the other production rules stated by this portion of the text, one will see that the action of "trying again" does not match the condition of any other

Table 1

A Word Processor Deletion Procedure

Each instruction sentence followed by informal production rule

Making Deletions

IF goal:delete text X THEN do step 1
...

The first step to delete text is to move the cursor under the first character to be deleted. This tells the system where the deletion starts.

IF step:1
THEN add goal: move cursor to first of X, do step 2
...

The next step is to press the DEL (Delete) key.

IF step:2 THEN press DEL, do step 3

When the system prompts, "Delete what?", type the last character of the text to be deleted.

IF step:3 and prompt="Delete what?"
THEN type last of X, do step 4

The cursor moves to the last character.

All the text from the first character through the last character is highlighted. You can see exactly what is going to be deleted before it is deleted.

If the wrong characters are highlighted, press CODE + CANCL and try again.

IF step:4 and X not highlighted
THEN press CODE + CANCL, add goal try again

When the text you want to delete is highlighted, press the ENTER key.

IF step:4 and X highlighted THEN press ENTER
...

The highlighted word is deleted, and the remaining text on the line moves over to take its place.

production rules in the vicinity. This means that after this rule is fired, there will be no other rules that match the "try again" situation, and so the procedure will grind to a halt. These training materials in fact did not say what "trying again" would consist of. Thus, rather than simply studying the material, the reader will have to infer the correct procedure, probably by trial and error experimentation with the word processor.

This lack of accuracy and completeness is apparently not an isolated problem. The second example, shown in Table 2 shows a more serious situation in the same section of the same manual. The reader is lead to believe that this paragraph will provide a summary of how to make deletions from the text. Following the text is an attempt to translate it into production rules. As can be seen by examining these rules, the text is in fact quite inaccurate and incomplete in specifying the whole procedure. There are many missing pieces of information. If those rules were put into a simulation, they would fail to carry out deletions correctly.

The point of these examples is that even if the reader is a perfect information processor, the instructional material in fact does not explicitly contain the complete and correct procedures that the learner is supposed to acquire. Instead of simply learning the procedures as stated, the learner must engage in a considerable amount of problem solving and experimentation in order to compose a correct procedure. Thus, it would be fair to say that the training materials do not assist the person in simply learning the procedures, but in fact require the learner to engage in problem solving to invent the procedures.

Of course, it could be argued that learning by solving problems is more effective than simply absorbing explicitly stated procedures, but this has not yet been demonstrated. It could be that some of the apparent superiority of learning by exploration is simply due to ordinary training materials being badly defective even as rote training materials.

Thus, a first speculation is that better training would result if the complete and correct procedure was explicitly presented in a form that the learner could easily translate into production rules. Attempting to translate the materials into a running production rule set is a way to evaluate the adequacy of the materials. Alternatively, the production rules in a running model could be used as the starting point for the training materials. The deletion procedure on the word processor is probably short and compact enough to present in complete detail, but apparently it is possible for readers to suffer from "overload" effects in attempting to acquire procedures from text (Kieras and Bovair, 1986). Thus there is apparently a pedagogical problem of how to build a complex procedure without

Table 2

Summary Description of Deletion Procedure and Its Production Rule Translation

Instructional Text

To Make Deletions

A. Position the cursor under the first character to be deleted.

B. Press the DEL key.

C. When the prompt, "Delete what?", appears, type the last character to be deleted. If you are deleting a single character, do not move the cursor.

D. Press ENTER.

Informal Production Rule Translation (Comments in brackets)

IF goal:Delete text X THEN do step A

IF step:A THEN add goal:move cursor to first of X, do step B
[another set of rules satisfies the move cursor subgoal]

IF step:B and cursor at first of X THEN press DEL, do step C
[this rule waits until the subgoal has been satisfied]

IF step:C and prompt="Delete what?"
THEN type last of X, do step D
[correct only if last character of X is unique in X]

IF step:C and prompt="Delete what?" and X is single character
THEN delete goal:move cursor
[the goal deleted here was not added by previous rules; this
is not correct - the step C rule should be specialized into
two rules for the one-character and string cases]

IF step:D THEN press ENTER, goal of deleting text X satisfied

overloading the learner. This question could be addressed with considerable precision by using the production rule characterization of the to-be-learned procedures.

A second speculation is that one could more efficiently test a learner for knowledge of the procedure by using the production rule characterization. Test problems could be devised that use paths that traverse all rules within the procedure. If the learner could successfully demonstrate knowledge of each one of the production rules, one could be confident that all of the components of the procedure were present.

The Acquisition and Use of a Mental Model

A good example of the use of a cognitive simulation to examine a testing and training situation appears in the work that we have done on the use of a mental model in learning how operate a simple piece of equipment (Kieras and Bovair, 1984). In this situation, people are acquiring not only a set of procedural skills, but also some declarative knowledge about the organization of a piece of equipment.

The task and the phenomena. Subjects learned how to use a simple piece of equipment whose front panel is illustrated in Figure 1. Shown in Figure 2 is the diagram of the fictitious system that subjects were told underlies the control panel. They were told that this system is the control panel for a "phaser bank" on the starship "Enterprise", and studied a few pages of material that essentially explained the diagram. Subjects studied the material, took a test on it, and were required to study further until they answered all questions correctly.

Two kinds of experimental tasks were used. In the first, subjects were explicitly trained on the procedures for operating the device in various situations, including those in which some fictitious internal component of the device was malfunctioning and subjects had to determine whether it was possible to compensate for the malfunction and if so to do so. After training, subjects were tested by being put into the situations repeatedly, without any feedback. The basic measurements are the usual learning and performance criteria. The second type of experimental task consisted of the subject having to infer how to operate the device in various situations without any explicit training in the procedures. The subjects were informed what the goal of the operation was, and then were free to attempt to get the device to operate. The performance measures consist of how many actions with controls subjects had to perform before they succeeded in getting the device to work, or before they could draw a conclusion that the device could not be made to work.

Briefly, compared to a group that had no training in the mental model for the device, the group which had the mental model

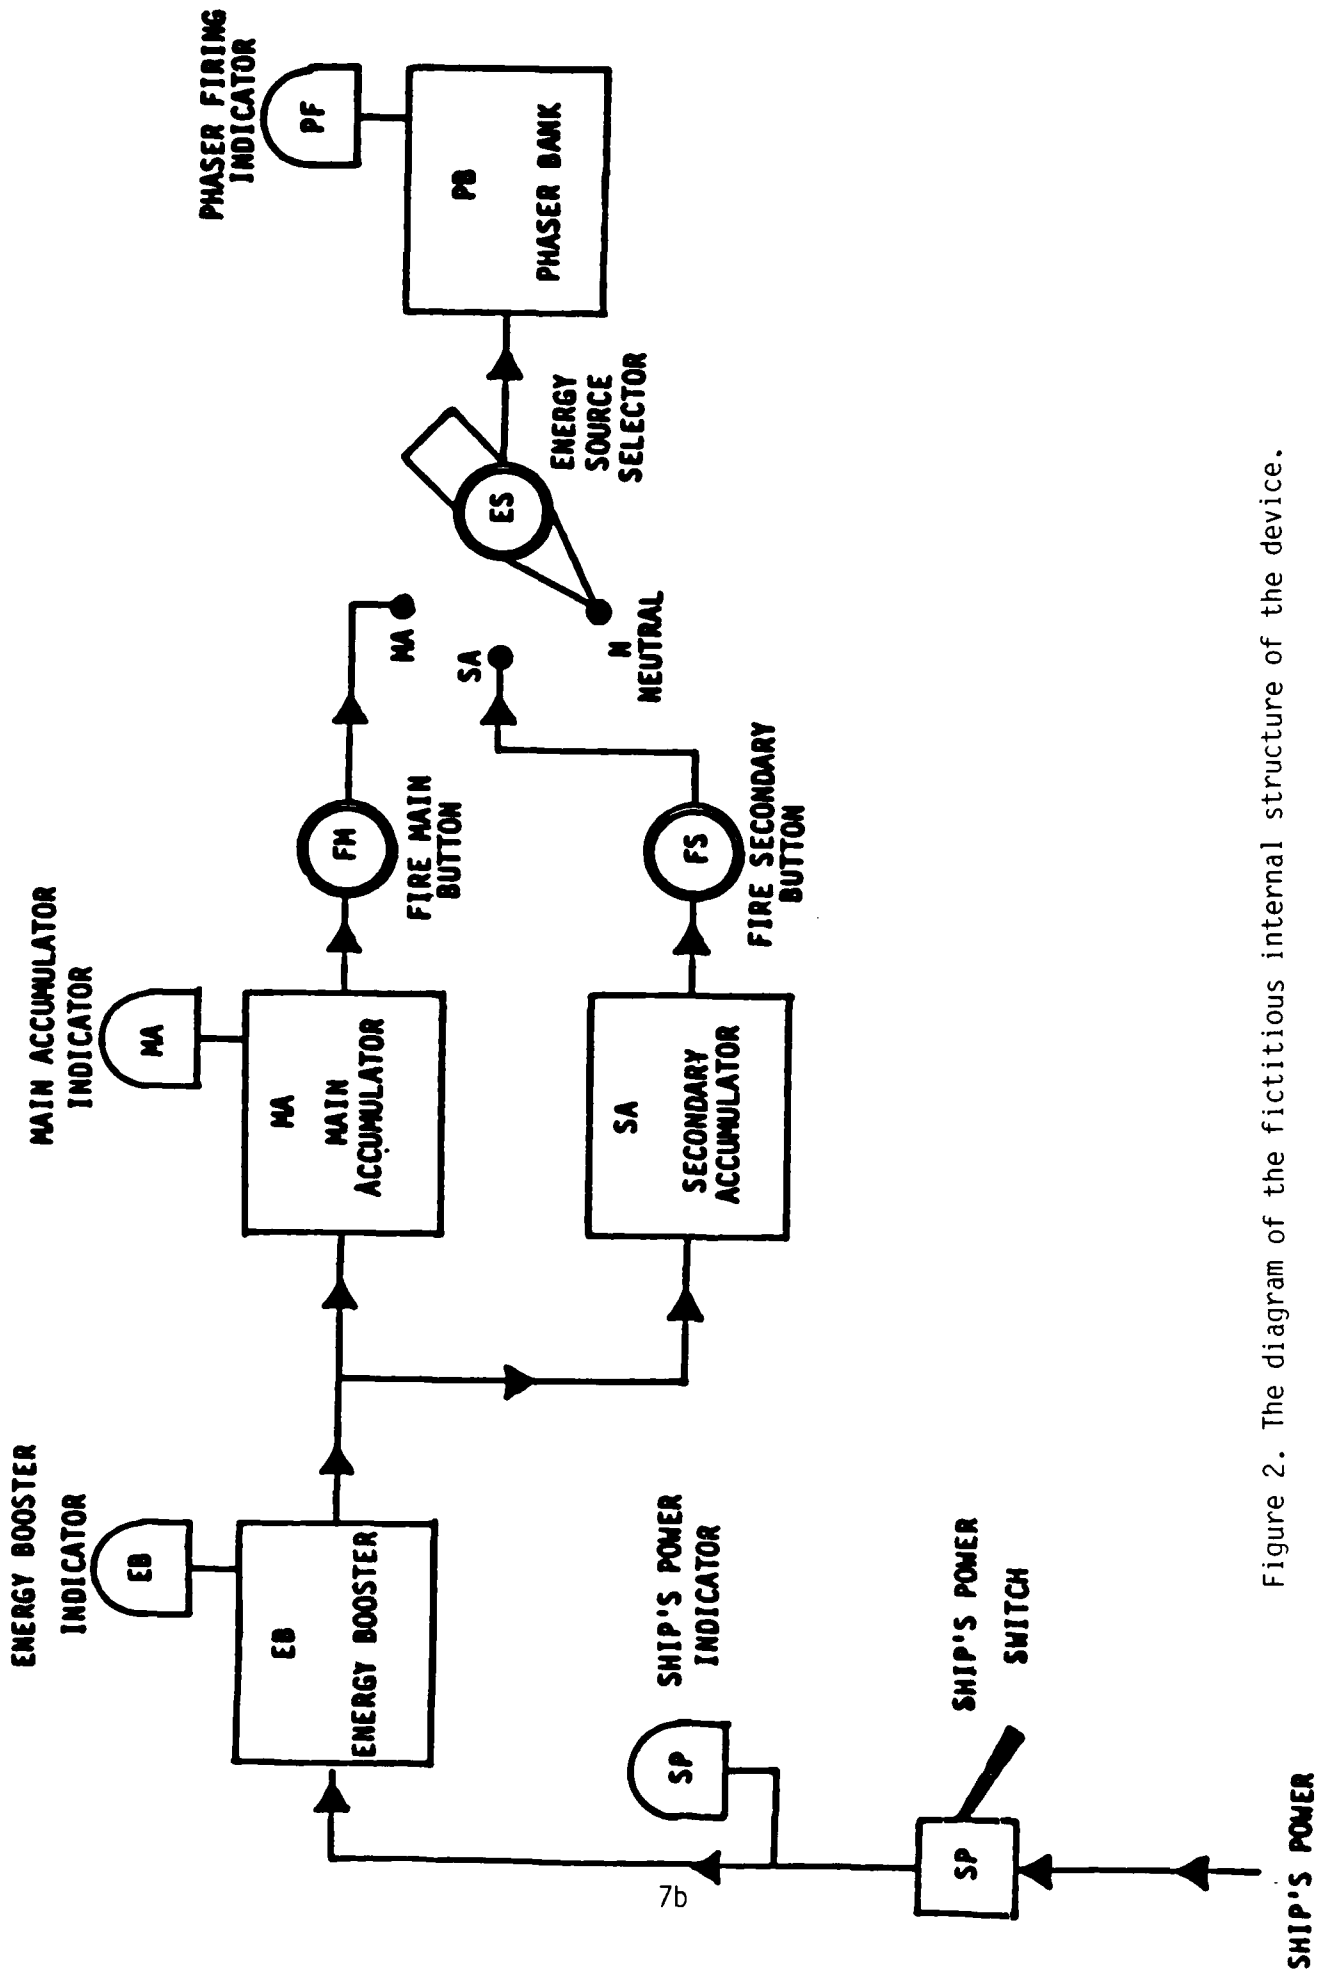


Figure 2. The diagram of the fictitious internal structure of the device.

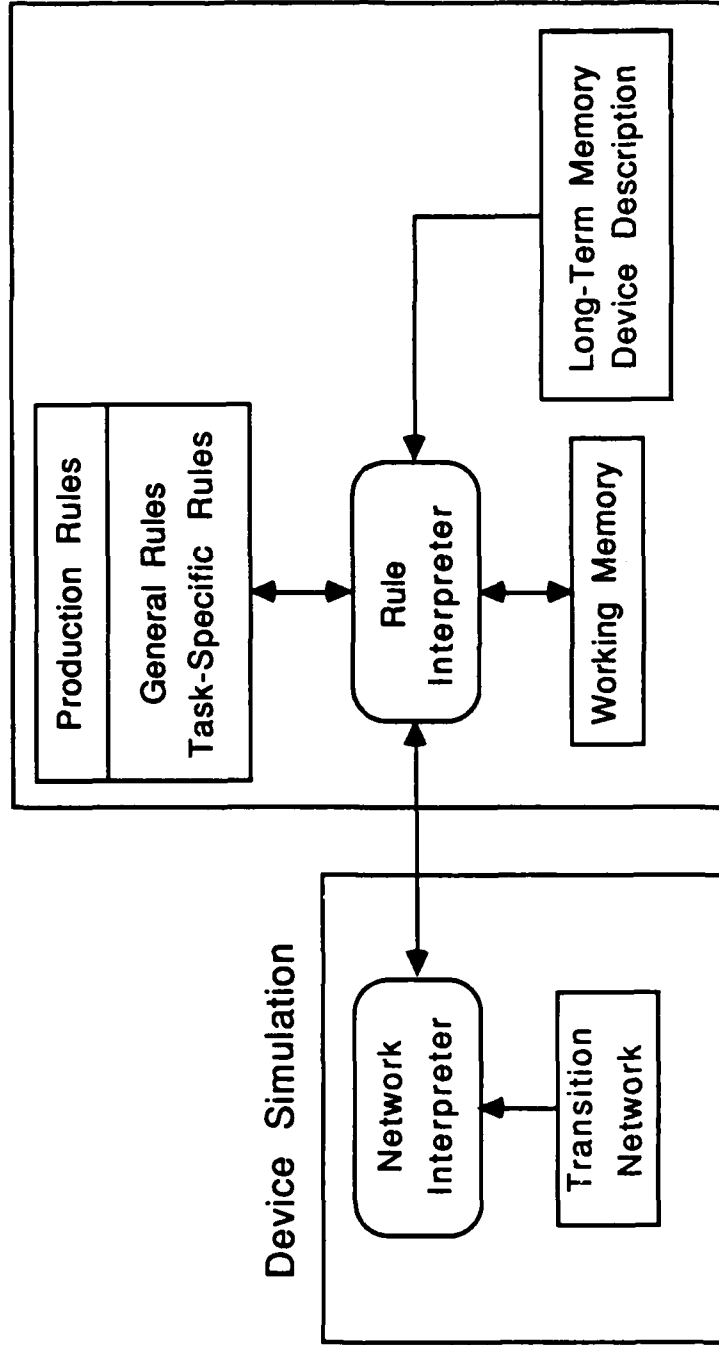
learned the procedures substantially faster, executed them more quickly, and retained them better, even after one week. In the inference task, subjects who knew the mental model were able to infer how to operate the device essentially immediately, with little or no trial and error search. In contrast, the subjects without the mental model essentially performed a systematic trial and error search, which eventually succeeded due to the simplicity of the device, but took substantially more operations on the device. By manipulating the training materials, we concluded that the important content of the mental model material was the information about which components were attached to which controls and indicators and to which other components. In other words, the topological information shown in Figure 2 was the key information. The fantasy content, or even the supposed explanations of how each of the components worked, was not relevant to this task.

When we devised the training materials, we were guided by intuition; we tried to explain the device, but not to "give away" the procedures for operating it. The test for mastery of the mental model was also prepared intuitively. The comparison of these materials with the simulation makes it very clear what we were actually teaching and testing.

A simulation of procedure inference. A simulation model was developed for how subjects could infer how to operate this device given knowledge of the system topology and of a set of rules for reasoning with the topological knowledge. Many such models are possible. The model was built using an intuitive analysis, as are most simulations, based on a characterization of the logical requirements of the task, and observation of the sequence of operations subjects performed. The model can in fact infer the procedures for operating the device. It generates sequences of actions that correspond to the majority of subject's systematic behavior, and the temporal predictions of the model account reasonably well for the latencies between individual actions. A basic constraint on the model was that it was to be as general as possible so that similar devices with different topological arrangements could also be operated by the model. The model that will be described is an updated version of the one reported in Kieras (1984).

The overall structure of the model is shown in Figure 3. The model consists of two major sections. One is a simulation of the control panel device itself, which simply makes developing the simulation of the human user more convenient, and so need not be discussed further. The simulation of the user's cognitive processes was implemented in terms of production rules that operate on a declarative database which consisted of a description of the device topology, and ISA facts about each object in the device. This is essentially a propositional

User Simulation



8a

Figure 3. Overall structure of the simulation model for procedure inference.

paraphrase of the diagram in Figure 2. Table 3 shows a few examples of this information.

The goal of the model is to operate the controls so as to get the PF indicator to light (this corresponds to the "phasers" being fired). The model has a top level control structure that divides this task into first starting up the device by switching it on, and perhaps setting the selector switch in response to a command as to which "accumulator" should be used. Then the model simulates the internal state of the device, resulting in a set of propositions in working memory that reflect where energy is assumed to be in the system and what the states of some of the internal components are. Then the model constructs a plan for operating the device by starting with the goal of getting energy to the phasers, and working backwards through the diagram until it finds an energy source. Once this plan is constructed, it is then executed by a process that represents stereotypical knowledge of what order controls should be operated in. For example, push buttons are normally operated last. If the goal state is reached, that is the PF indicator flashes, the model then signals a successful attempt. If not, the model updates its simulation of the internal state of the device, applies some heuristics to diagnose what the source of the problem is, which will result in an additional component being labelled as malfunctioning, and then attempts to construct a new plan for operating the device. If this new plan cannot be successfully constructed, or fails to succeed and no other plan is available, then the model signals that the device cannot be successfully operated.

Example production rules. It is not necessary to go into any additional detail in this paper concerning the model's behavior or how it works. What will be given is some samples of the production rules to illustrate how this model specifies pieces of knowledge that the user should need to know.

Table 4 shows some examples of production rules. The first rule is an example of a rule that propagates energy forward along a connection. The name of this rule is InferConnectionEnergy, and it specifies that if the goal is to simulate the internal state of the device, and in working memory we have that energy is at a certain point, indicated by the value of the variable ?T1, and we know from the LTM representation of the device that ?T1 is connected ?T2, which is a terminal, and there is not already energy at ?T2, then we will add to working memory the proposition that there is energy at the point ?T2, and also a note that we have made some progress in simulating the device. This rule captures the piece of inferential knowledge that necessarily underlies diagrams of the type shown in Figure 2. Namely, the lines shown in the diagram mean that if there is energy at one point in the system, then it can be inferred that this energy is also at the next point "downstream" in the system. This rule,

Table 3

Excerpt from Long-Term Memory Description of Device Structure

(LTM ISA EB COMPONENT)
(LTM ISA MA COMPONENT)
(LTM ISA SPI INDICATOR)
(LTM ISA EBI INDICATOR)
(LTM ISA EB-OUT TERMINAL)
(LTM ISA MA-FM TERMINAL)
(LTM ISA SP SWITCH)
(LTM ISA FM BUTTON)
(LTM ISA SHIP-POWER POWER-SOURCE)
(LTM ISA SWITCH CONTROL)
(LTM ISA BUTTON CONTROL)
(LTM ISA SELECTOR CONTROL)
(LTM SETTING SP SP ON)
(LTM SETTING ESS-MA ESS MA)
(LTM SETTING FM FM PUSH)
(LTM ASSOCIATED MA FM)
(LTM ASSOCIATED MA ESS-MA)
(LTM CONNECTION SHIP-POWER SP-IN)
(LTM CONNECTION SP-IN SP)
(LTM CONNECTION SP SP-EB)
(LTM CONNECTION SP-EB SPI)
(LTM CONNECTION SP-EB EB)
(LTM CONNECTION EB EB-OUT)
(LTM CONNECTION EB EBI)
(LTM CONNECTION EB-OUT MA)
(LTM CONNECTION MA MA-FM)
(LTM CONNECTION ESS-PB PB)
(LTM CONNECTION ESS-PB PB)
(LTM CONNECTION PB PFI)

Table 4

Examples of Production Rules from the Simulation

```
(InferConnectionEnergy
IF  ((GOAL SIMULATE DEVICE)
      (WM AT ENERGY ?T1)
      (LTM CONNECTION ?T1 ?T2)
      (LTM ISA ?T2 TERMINAL)
      (NOT (WM AT ENERGY ?T2)))
THEN ((Add WM AT ENERGY ?T2)
      (Add WM SIMULATION IN PROGRESS)))

(FollowControlBack
IF  ((GOAL FIND PATH)
      (WM GOAL ENERGIZE ?T3)
      (LTM CONNECTION ?T2 ?T3)
      (LTM CONNECTION ?T1 ?T2)
      (NOT (WM GOAL ENERGIZE ?T1))
      (NOT (WM AT ENERGY ?T3))
      (LTM ISA ?T2 ?C2)
      (LTM ISA ?C2 CONTROL)
      (LTM ASSOCIATED ?C3 ?T2)
      (WM COMMAND USE ?C3)
      (NOT (WM STATE ?C3 BAD)))
THEN ((Add WM GOAL ENERGIZE ?T1)
      (Add WM FINDING IN PROGRESS)
      (Add WM PLAN OPERATE ?T2)
      (Delete WM GOAL ENERGIZE ?T3)))
```

together with several others, are used by the model to update and maintain a representation of where energy must be present in the control panel system.

The second rule in Table 4 shows a small piece of planning knowledge. This rule, called FollowControlBack, starts with a goal to find a path through the system, more specifically to get energy to a certain point, ?T3. It finds two additional points in the system, ?T2, which is a control of some sort, and ?T1, which is the point of the system "upstream" from the control. It also identifies a component ?C3 that is associated with this control, and checks to be sure that this component is both supposed to be used, as specified by the command, and is not known to be bad. If so, this rule makes a note that our new goal is to energize the point upstream from the control, ?T1, and that part of our plan for operating the device will be to operate the control ?T2. Thus, this rule represents a specific piece of the knowledge required for determining how to get energy to a particular point in the system. Namely, if the pathway for the energy involves a switch or control of some sort, then that control will have to be turned on before the overall operation goal can be accomplished.

The set of rules shown in Table 5 illustrate how inferences can be made about the states of the internal components in the system. The first rule, InferIndicatorEnergy, states that if an indicator light, ?T2, on the device is on, then one can infer that the point at which that indicator is connected has energy present. The rule InferComponentGood says that if it is known that energy is at a component, then that component must be good. This rule, working in conjunction with InferIndicatorEnergy, will result in the system knowing that a component is good if the indicator light attached to it is on. The rule InferComponentBad does the opposite inference. If there is an indicator that is attached to a component that has energy at its input, but the indicator light is not on, then the component must be bad. Thus if it is known that there is energy going into a component, then the failure of the attached indicator light to be on must mean that the component is bad.

The last rule shown in Table 5 is InferComponentBad2. This rule is used to diagnose why the device failed to operate after the plan was executed. If there is a component that is known to have energy at its input, and for which part of the plan was to operate a control associated with the component, and there is apparently no energy coming out of the component, then the conclusion is that the component is bad. This rule is actually a heuristic rule, in the sense that even with this simple device it will not always apply correctly. But it does represent the idea of blaming the most "upstream" component in a situation in which it is not possible to determine positively which component is bad. For example, referring to Figure 2, if an attempt was made

Table 5

Some Inference Rules from the Simulation

```
(InferIndicatorEnergy
IF ((GOAL SIMULATE DEVICE)
    (DEVICE ?T2 ON)
    (LTM ISA ?T2 INDICATOR)
    (LTM CONNECTION ?T1 ?T2)
    (NOT (WM AT ENERGY /T1)))
THEN ((Add WM AT ENERGY ?T1)
      (Add WM SIMULATION IN PROGRESS)))
```

```
(InferComponentBad
IF ((GOAL SIMULATE DEVICE)
    (LTM CONNECTION ?T1 ?T2)
    (LTM CONNECTION ?T2 ?T3)
    (LTM ISA ?T1 TERMINAL)
    (LTM ISA ?T2 COMPONENT)
    (LTM ISA ?T3 INDICATOR)
    (WM AT ENERGY ?T1)
    (NOT (DEVICE ?T3 ON))
    (NOT (WM STATE ?T2 BAD)))
THEN ((Add WM STATE ?T2 BAD)
      (Add WM SIMULATION IN PROGRESS)))
```

```
(InferComponentGood
IF ((GOAL SIMULATE DEVICE)
    (LTM ISA ?T2 COMPONENT)
    (WM AT ENERGY ?T2)
    (NOT (WM STATE ?T2 GOOD)))
THEN ((Add WM STATE ?T2 GOOD)
      (Add WM SIMULATION IN PROGRESS)))
```

```
(InferComponentBad2
IF ((GOAL DIAGNOSE DEVICE)
    (LTM CONNECTION ?T1 ?T2)
    (LTM CONNECTION ?T2 ?T3)
    (LTM ISA ?T1 TERMINAL)
    (LTM ISA ?T2 COMPONENT)
    (LTM ISA ?T3 TERMINAL)
    (WM AT ENERGY ?T1)
    (WM PLAN OPERATE ?C2)
    (LTM ASSOCIATED ?T2 ?C2)
    (NOT (WM AT ENERGY ?T3))
    (NOT (WM STATE ?T2 BAD)))
THEN ((Add WM STATE ?T2 BAD)
      (Add WM DIAGNOSIS IN PROGRESS)))
```

to use the secondary accumulator (SA), and the phasers did not fire, this could be due either to a defective phaser bank (PB) or to a defective secondary accumulator. Because the secondary accumulator does not have an associated indicator light, the situation is ambiguous. A reasonable heuristic seems to be to assume that first component upstream must be the defective one, in the absence of any other information. Thus the rule InferComponentBad2 will apply and will label the secondary accumulator as bad. When the new plan is formed, it will conclude that the main accumulator should be tried.

Evaluating training materials. The purpose of this example is to show how the detailed characterization of the task contained in the simulator could be used to examine the instructional materials and testing materials more closely.

At a very general level, one useful result of having developed the simulation is that it shows how much is actually needed for a usable mental model. The characterization coming from the experiments was simply that a mental model should contain the content necessary to allow the user to infer how to operate the device. It seemed to consist mostly of information about the system topology, along with some inference rules for making use of this topology. However, what we see by looking at the simulation model is that the rules of inference for making use of the topology are many in number, quite specific in their content, and have to be supplemented by additional rules for the overall strategy of performing the task. Beyond this general characterization, we can look in some detail at the instructional materials to see what they actually contain. For this comparison, I will be assuming that the mental model simulation described above is in fact an adequate or correct characterization of the reasoning that the typical user would do. The truth of this statement is another matter; here the model can just be used to illustrate the approach.

Examination of the materials (presented in full in Kieras, 1984) made it clear that the instructional materials did not state the overall strategies or even the general rules used in the simulation. Instead, the instructional materials contain only limited, highly specific information. For example, the only thing in the material that has any relationship to the rules shown in Table 4, such as InferConnectionEnergy, is a single sentence that states The arrows on the diagram show how power flows through the system. Elsewhere, the material only states information that power can flow from one specific component to another. There is in fact no explicit statement of the general rule that power can flow between connected components in a downstream direction. Apparently the learner is being asked to infer that this is a property of the system.

Likewise, the only general statement about the controls is the sentence, The switch, selector, and push buttons control the flow of power. Notice that the material does not state the general rule that when a switch or control is set to the on state, then power can flow through the control to the connected points. Instead, what these materials present is each specific case. The bulk of the information about the individual selectors and buttons is always stated only in the specific form and sometimes in a slightly misleading fashion. For example the instructional material says that When the selector is set to MA, then power can flow from the main accumulator. When the selector is set to SA, then power can flow from the secondary accumulator. This material is misleading in that setting the selector is not enough to get the power flowing from the accumulator; rather, the push button must also be pressed as well. The material presents this fact a few sentences later, but one wonders whether the subject had at least for a short time an incorrect representation of the relationship of the selector to the accumulators.

The only obvious case in which the materials present a general rule similar to those in the simulation is a sentence describing the indicator lights. This sentence is The indicator will only light if the component that it is connected to is both receiving power and working properly. This rule is very similar to the rules in Table 5, but it is easy to see that this stated rule does not directly correspond to a single one of the simulation rules for inferring that a component is good (this may actually be a defect in the simulation). Furthermore, the opposite rule, that a component is bad if its indicator is not on when it should be, is being left up to the reader to infer.

The rest of the training material about the indicators presents the inference for each indicator individually, and in an order that is in some sense the reverse of what the subject would actually use. For example, the material states that The main accumulator indicator will light if the main accumulator is receiving power from the energy booster, and the main accumulator is working properly and putting out power. This rule states a special case of the general rule discussed above, and it does so in a backwards direction. That is, it seems that the subject would start with the perceptual information that the indicator is on and will want to know what conclusion can be drawn. Thus, the reader will have to take the information in the training materials sentence, and translate it into a form more suitable for their needs.

The training materials contain no apparent information on how to diagnose a malfunction in the system, or how to construct a plan for operating the device. Of course, these materials were deliberately designed not to contain any procedural information, but it was not clear to us at the time that we had required the subject come up with so many of their own inference rules.

In summary, the training materials lacked explicit specification of: (1) the overall strategy for doing the task; (2) general rules for making inferences about power flow or the state of components; (3) a general procedure for planning how to operate the controls given the inferred state of the device; (4) rules for diagnosis of failures in the system. On the other hand, the training materials did clearly include specific facts about power flow and the effects of the controls, and some specific inferences that could be made from the indicator lights.

An obvious question is what is the correct characterization of the materials? Notice that the materials definitely lack key facts and procedures, but the simulation is complete (it can perform the task) and is reasonably consistent with subject's behavior. Thus, if the model is taken as a description of the subject's knowledge, it is clear that the subjects had to infer a considerable amount of knowledge beyond that presented in the training materials.

Evaluating testing materials. In the experiments, we always insisted that subjects demonstrate a knowledge of the material that presented the mental model before proceeding to the rest of the experiment. This was primarily for methodological reasons; we suspected that many of the experiments intended to demonstrate the effects of mental models had not done an adequate job of ensuring that the subjects actually knew the mental model information. However, examining the testing materials for information that was critical in the simulation model produced something of a surprise. Rather than testing for the ability to reason about the system, the test questions focused mostly on the system topology, and most of these could be answered by simply paraphrasing the instructional text.

For example, one question is Where does the secondary accumulator get its power from?, with the multiple choice alternatives being (1) Directly from the ship board circuits (2) From the phaser bank (3) From the energy booster. This question can be answered almost directly from two sentences appearing in the training materials, Starting on the lower left of the diagram, you can see that power comes in from the ship board circuits. ... Power from the energy booster then flows into both accumulators. Likewise, questions about the indicator lights could also be answered almost directly from the instructional material.

Only two of the twelve test questions are more demanding. These are presented in Table 6. These two questions essentially ask the subject to simulate the operation of the phaser system, based on the diagram (which was present during training and testing) and their knowledge of how to make inferences. In terms of the simulation model, by simulating where the energy is in the

Table 6

Test Questions that Require Some Inference

Assume that the phaser control system is in full working order, that the PS is on, and that the selector is set to MA.

Now, what will happen if the M button is pressed?

- (1) The Main Accumulator will send power to the Phaser bank.
- (2) The Phaser bank will receive power from the Secondary Accumulator.
- (3) The Phaser bank will receive power directly from the energy Booster.

Assume that the phaser control system is in full working order, that the PS is on, and the selector is set to MA.

Now, what will happen if the S button is pressed?

- (1) Nothing. The selector must be set to SA for power to flow to the Phaser bank when the S button is pressed.
 - (2) The Main Accumulator will send power to the Phaser bank.
 - (3) The Secondary Accumulator will send power to the Phaser bank.
-

system, the subjects should be able to choose the correct alternative answer. The instructional material contains several sentences that are relevant to answering this question, but unlike the other questions, these sentences do not provide the answer by paraphrase. Rather, at least a little bit of inference is required. Thus, these two questions might actually be testing for whether subjects had acquired not only a correct understanding of the system topology, but also the rules for making some inferences from it.

However, these two questions are not very strong tests of inferential ability; the tested inference process is only one of several types of inference that the subject must be able to perform in order to operate the device. For example, according to the simulation, some other inferences that have to be made are how to operate the device given the simulated state of the system, and how to determine whether a component is defective, especially if the indicator information is not adequate. But knowledge of how to make these inferences was not addressed by the test at all. Other information missing from the test is an overall strategy for executing the task, and how to operate a device given a plan for operating it. For example, as mentioned above, for devices of this general type, push buttons are conventionally operated last in a sequence of operations.

Questions presented by the analysis. In the experiments, the mental model group still had an advantage over the rote learning group, even though, as we have seen, the training materials and the test questions were far from adequate in presenting a complete mental model and testing for the acquisition of it. This raises a question about whether a group that knew a mental model would have an even greater advantage if a fully usable model had been explicitly presented and tested before subjects began trying to use it. Furthermore, with the materials that we did use, it appears that subjects would have had to perform a lot of inference and transformation of the materials before they could use them to reason about the device. It should be possible to capture some of these processes at work.

Given that the materials actually presented only specific information about the device, as contrasted to the rather general inference rules in the simulation model, another question is whether subjects who were explicitly taught the more general model whether they could operate the device better, or show superior performance in transferring to a different device.

Another question concerns the formal properties of the relatively specific material that was presented to subjects. Would it be possible to start with the specific rules and descriptions contained in our training materials, and construct a simulation model which could successfully infer how to operate the device? This would be a way to determine just how seriously

defective our training materials were. Presenting subjects with these specific rules of inference may have in fact worked fairly well, although we would predict that they would have trouble in transferring such specific knowledge to a different device.

Concerning the test itself, notice that the test questions could have been very different. For example, we could have tested for the acquisition of the general rules for reasoning about the device instead of the specific ones. More specific topology information could have been tested for. The simulation model works on the basis of the individual point-to-point connections in the system; the test could have tested for similar point-to-point connection knowledge. The test questions also could have exercised the specific inference processes that users needed to have. For example, the test questions could have asked people to draw an inference about which components were bad from a specific pattern of the indicator lights.

If subjects could pass a test containing such items, we could be confident that not only did they have the declarative knowledge of how the device was structured, but also the procedural knowledge for how to make the relevant class of inferences from the model. As it was, we were essentially leaving it up to subjects to construct these inferences as they went along, trying to operate the device. Thus, if we had test items that tested for the specific aspects of the mental model, and if passing the test is the criterion of when the mental model training is complete, we should see superior performance in operating the device. Additionally, given less thorough training, such a test should allow us to distinguish good from poor performers.

CONCLUSIONS

In this paper I have argued that a working cognitive simulation model contains a specification of the knowledge that a person must have in order to accomplish the task, and this specification can be used as the basis for training and testing materials.

A possible useful alternate approach would be to start with an existing set of instructional materials that are supposed to be good materials for the domain, and attempt to construct a simulation model using the materials as directly and as naively as possible. If the resulting model cannot perform the task, then it is clear that something is missing from the materials, and needs to be added. Once the model can perform the task, then the behavior of the model can be examined in some detail to determine which portions of the material are actually critical and which are not. Thus, rather than using the simulation as a source of specifications for to-be-written training material, the

simulation approach could be used as a way to evaluate existing materials.

Although the example analyses presented above are informal, it does appear that such model-based evaluations of testing and training materials can clarify the actual content and its adequacy relative to the model. The claim was made that materials that explicitly contain the same knowledge as the model will be superior to the normal intuitively prepared materials; but this has to be demonstrated experimentally.

Puzzles and Obstacles

The above discussion of using cognitive simulation to prepare and evaluate training material is obviously speculative, and many problems would have to be solved before such an approach could even be evaluated, much less implemented. However, at this point several obstacles and potential problems can be pointed out.

Level of detail. The first problem concerns the issue of level of analysis in the model. It should not be necessary for the cognitive simulation modelling to explicitly model all aspects of the cognitive processing involved in a particular task. We should be able to finesse some of the more complex processes.

For example, to model electronics trouble-shooting skill, it really should not be necessary to model the perceptual process by which schematic diagrams are perceived. Clearly an electronics trouble-shooter must have this skill, but we could analyze the skill of using schematic diagrams into roughly two components: one is the visual perception of the diagram itself, the other is the interpretation of the diagram. The interpretation process consists of interpreting the meaning of the symbols in the diagram, determining from the diagram what the relations between these symbols are, and making the appropriate inferences from the diagram about the structure and behavior of the system. Modelling the interpretation processes would of course be a major task; modelling the visual perception process involved in examining a schematic diagram would also be a major task, perhaps even more difficult than the interpretation problem.

The argument is that it should not be necessary to model the visual perception process simply because everyone who is likely to be trained in electronics trouble-shooting should be able to perceive lines, squares, circles, and so forth that are printed on paper. For example, every trainee should be able to determine that two circles are connected with a line. What has to be taught as part of electronics training is information such as that two circles with certain additional symbols inside represent

two different transistors, and that the line interconnecting specific places on these two circles indicates that the collector of one transistor is electrically connected to the base of the other transistor. This is part of the interpretation process, not a visual perception process. Thus, the modelling effort could finesse the perceptual process, and start with the input to the interpretation process.

Making decisions about what psychological processes can be finessed, and which processes have to be explicitly represented, would be critical to the success of an attempt to use simulation models to specify the knowledge that has to be learned. The puzzle is whether we can reliably identify, and then focus on, exactly the critical information in the complex tasks that we are interested in training.

Non-identifiability. A second problem is that it is impossible to know whether a particular simulation model is the correct representation of how a person actually performs a task, and so basing training on a single model may be misleading. This non-identifiability problem has been discussed at length elsewhere (e.g. see Kieras, 1981; Anderson, 1978). However, it should not discourage us from trying this approach. Notice that although there are many possible simulation models that could account for behavior in a task, only a subset of these will appear to the developer of training materials as being relevant, based on expert opinion on how such knowledge and tasks should be organized, presented, and used. For example, trouble-shooting strategies for electronics maintenance can be obtained from experts in the domain, and the processes and representations needed to implement these strategies will then be more highly constrained than if a purely AI or intuitive approach were used. Thus, non-identifiability will not prevent us from arriving at reasonable, intuitively sound, and effective cognitive models.

Cost-effectiveness. A third obstacle is that historically, constructing simulation models of complex processes has been a difficult and time-consuming occupation. It would not help an advanced training and testing effort to be more efficient if it required an extremely time-consuming and difficult knowledge analysis to be done beforehand. Thus, the question is whether cognitive simulations can be constructed cost-effectively. It should be for skills such as electronics troubleshooting, which are in great demand, and have a well-understood rational basis that can be used as the foundation for the modelling effort.

The answer to this problem is that we need to have a technology of cognitive modelling. The adoption of explicit cognitive architectures is a step in this direction, because the cognitive architecture represents a pre-packaged set of decisions and mechanisms; adopting an architecture frees the modeler from having to make these decisions and construct these mechanisms

from scratch for each model. But the basic bottle-neck in constructing cognitive simulations seems to lie primarily in the task analysis methodology. Apparently, cognitive psychologists know intuitively how to look at a task and break it down in a quick, informal way into its major components; this informal analysis is the first step in constructing a more specific simulation model. We have not codified task analysis methodology, but have always done it intuitively. Thus, we cannot convey task analysis methodology to other people except by long-term apprenticeship in a graduate training program. This limitation will clearly make it very difficult to develop large numbers of simulation models for complex tasks quickly. An effort to make explicit our intuitive analysis methodology would potentially be of great benefit.

Portability of representations. A final obstacle is whether the information in a cognitive model can be converted directly into the information needed in intelligent tutoring and advanced testing systems. Based on the examples given in Tables 3, 4, and 5, it seems clear that the information in the simulation model can be expressed in ways that are quite compatible with existing approaches and techniques in artificial intelligence. Thus once the model was developed, the information in it should be directly exportable as input to intelligent training and testing systems. This is basically a software engineering problem; both the simulation models and the intelligent tutoring and testing systems need to work with compatible knowledge representations, meaning that some kind of common architecture needs to be assumed. In some intelligent tutoring efforts this commonality of representation between the simulation model for the cognitive processes and the intelligent tutoring system is in fact present.

References

- Anderson, J. R. (1976). Language, memory, and thought. Hillsdale, N. J.: Lawrence Erlbaum Associates.
- Anderson, J. R. (1978). Arguments concerning representations for mental imagery. Psychological Review, 85, 249-277.
- Anderson, J. R. (1983). The architecture of cognition. Cambridge: Harvard University Press.
- Kieras, D. E. (1981). Knowledge representations in cognitive psychology. In L. Cobb & R. M. Thrall (Eds.), Mathematical Frontiers of the Social and Policy Sciences, AAAS Selected Symposium 54, Boulder, Colorado: Westview Press.
- Kieras, D. E. (1984). A simulation model for procedural inference from a mental model for a simple device (Technical Report No. 15, UARZ/DP/TR-84/ONR-15). University of Arizona, Department of Psychology.
- Kieras, D. E. (1985). The why, when, and how of cognitive simulation: A tutorial. Behavior Research Methods, Instruments, & Computers, 17 (2), 279-285.
- Kieras, D. E. & Bovair, S. (1984). The role of a mental model in learning to operate a device. Cognitive Science, 8, 255-273.
- Kieras, D. E., & Bovair, S. (1986) The acquisition of procedures from text: A production-system analysis of transfer of training. Journal of Memory and Language, 25, 507-524.
- Kieras, D. E., & Polson, P. G. (1985). An approach to the formal analysis of user complexity. International Journal of Man-Machine Studies, 22, 365-394.
- Polson, P. G., & Kieras, D. E. (1985). A quantitative model of the learning and performance of text editing knowledge. CHI'85 Conference Proceedings, 207-212.

1987/02/03 Distribution List
 (Michigan/Kieras) NR 667-547

Dr. Beth Adelson
 Department of Computer Science
 Tufts University
 Medford, MA 02155

AFOSR,
 Life Sciences Directorate
 Bolling Air Force Base
 Washington, DC 20332

Dr. Robert Ahlers
 Code N711
 Human Factors Laboratory
 Naval Training Systems Center
 Orlando, FL 32813

Dr. Ed Aiken
 Navy Personnel R&D Center
 San Diego, CA 92152-6800

Dr. John Allen
 Department of Psychology
 George Mason University
 4400 University Drive
 Fairfax, VA 22030

Dr. William E. Alley
 AFHRL/MOT
 Brooks AFB, TX 78235

Dr. John R. Anderson
 Department of Psychology
 Carnegie-Mellon University
 Pittsburgh, PA 15213

Technical Director, ARI
 5001 Eisenhower Avenue
 Alexandria, VA 22333

Dr. Patricia Baggett
 University of Colorado
 Department of Psychology
 Box 345
 Boulder, CO 80309

Dr. Eva L. Baker
 UCLA Center for the Study
 of Evaluation
 145 Moore Hall
 University of California
 Los Angeles, CA 90024

Dr. James D. Baker
 Director of Automation
 Allen Corporation of America
 401 Wythe Street
 Alexandria, VA 22314

Dr. Meryl S. Baker
 Navy Personnel R&D Center
 San Diego, CA 92152-6800

Dr. James Ballas
 Georgetown University
 Department of Psychology
 Washington, DC 20057

Dr. Harold Bamford
 National Science Foundation
 1800 G Street, N.W.
 Washington, DC 20550

prof. dott. Bruno G. Bara
 Unità di ricerca di
 intelligenza artificiale
 Università di Milano
 20122 Milano - via F. Sforza 23
 ITALY

Dr. Isaac Bejar
 Educational Testing Service
 Princeton, NJ 08450

Leo Beltracchi
 United States Nuclear
 Regulatory Commission
 Washington DC 20555

Dr. John Black
 Teachers College
 Columbia University
 525 West 121st Street
 New York, NY 10027

Dr. Arthur S. Blalwes
 Code N711
 Naval Training Systems Center
 Orlando, FL 32813

Dr. R. Darrell Bock
 University of Chicago
 NORC
 6030 South Ellis
 Chicago, IL 60637

Dr. Deborah A. Boehm-Davis
 Department of Psychology
 George Mason University
 4400 University Drive
 Fairfax, VA 22030

Dr. Sue Bogner
 Army Research Institute
 ATTN: PERI-SF
 5001 Eisenhower Avenue
 Alexandria, VA 22333-5600

Dr. Gordon H. Bower
 Department of Psychology
 Stanford University
 Stanford, CA 94306

Dr. Richard Braby
 NTSC Code 10
 Orlando, FL 32751

Dr. Robert Breaux
 Code N-095R
 Naval Training Systems Center
 Orlando, FL 32813

Commanding Officer
 CAPT Lorin W. Brown
 NROTC Unit
 Illinois Institute of
 Technology
 3300 S. Federal Street
 Chicago, IL 60616-3793

Dr. John S. Brown
 XEROX Palo Alto Research
 Center
 3333 Coyote Road
 Palo Alto, CA 94304

Dr. John Bruer
 The James S. McDonnell
 Foundation
 Univ. Club Tower, Suite 1610
 1034 South Brentwood Blvd.
 St. Louis, MO 63117

Dr. Bruce Buchanan
 Computer Science Department
 Stanford University
 Stanford, CA 94305

Mr. Donald C. Burgy
 General Physics Corp.
 10650 Hickory Ridge Rd.
 Columbia, MD 21044

Maj. Hugh Burns
 AFHRL/IDE
 Lowry AFB, CO 80230-5000

Dr. Patricia A. Butler
 OERI
 555 New Jersey Ave., NW
 Washington, DC 20208

Joanne Capper
 Center for Research into
 Practice
 1718 Connecticut Ave., N.W.
 Washington, DC 20009

Dr. Pat Carpenter
 Carnegie-Mellon University
 Department of Psychology
 Pittsburgh, PA 15213

Dr. John M. Carroll
 IBM Watson Research Center
 User Interface Institute
 P.O. Box 218
 Yorktown Heights, NY 10594

Dr. Robert Carroll
 OP 01B7
 Washington, DC 20370

LCDR Robert Carter
 Office of the Chief
 of Naval Operations
 OP-01B
 Pentagon
 Washington, DC 20350-2000

Dr. Fred Chang
Strategic Technology Division
Pacific Bell
2600 Camino Ramon
Rm. 3S-453
San Ramon, CA 94583

Dr. Davida Charney
English Department
Penn State University
University Park, PA 16802

Dr. Eugene Charniak
Brown University
Computer Science Department
Providence, RI 02912

Dr. L. J. Chmura
Computer Science and Systems
Code: 7590
Information Technology Division
Naval Research Laboratory
Washington, DC 20375

Dr. Yee-Yeen Chu
Perceptronics, Inc.
2111 Erwin Street
Woodland Hills, CA 91367-3713

Dr. William Clancey
Stanford University
Knowledge Systems Laboratory
701 Welch Road, Bldg. C
Palo Alto, CA 94304

Dr. Charles Clifton
Tobin Hall
Department of Psychology
University of
Massachusetts
Amherst, MA 01003

Dr. Stanley Collyer
Office of Naval Technology
Code 222
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. Lynn A. Cooper
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

LT Judy Crookshanks
Chief of Naval Operations
OP-112G5
Washington, DC 20370-2000

Phil Cunniff
Commanding Officer, Code 7522
Naval Undersea Warfare
Engineering
Keyport, WA 98345

CAPT P. Michael Curran
Office of the CNO
Director, Naval Medicine
Pentagon, Room 4D471, OP-939
Washington, DC 20350-2000

Dr. Cary Czychon
Intelligent Instructional
Systems
Texas Instruments AI Lab
P.O. Box 660245
Dallas, TX 75266

Brian Dallman
3400 TTM/TTGXS
Lowry AFB, CO 80230-5000

Dr. Natalie Dehn
Department of Computer and
Information Science
University of Oregon
Eugene, OR 97403

Goery Delacote
Directeur de L'informatique
Scientifique et Technique
CNRS
15, Quai Anatole France
75700 Paris FRANCE

Dr. Thomas E. DeZern
Project Engineer, AI
General Dynamics
PO Box 748
Fort Worth, TX 76101

Dr. Andrea di Sessa
University of California
School of Education
Tolman Hall
Berkeley, CA 94720

Dr. Stephanie Doan
Code 6021
Naval Air Development Center
Warminster, PA 18974-5000

Dr. Emanuel Donchin
University of Illinois
Department of Psychology
Champaign, IL 61820

Defense Technical
Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC
(12 Copies)

Dr. Jean-Pierre Dupuy
Ecole Polytechnique
Crea 1 Rue Descartes
Paris, FRANCE 75005

Mr. Ralph Dusek
ARD Corporation
5457 Twins Knolls Road
Suite 400
Columbia, MD 21045

Edward E. Eddowes
CNATRA N301
Naval Air Station
Corpus Christi, TX 78419

Dr. William Epstein
University of Wisconsin
W. J. Brogden Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

Dr. Edward Esty
Department of Education, OERI
Room 717D
1200 19th St., NW
Washington, DC 20208

Dr. Beatrice J. Farr
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Marshall J. Farr
Farr-Sight Co.
2520 North Vernon Street
Arlington, VA 22207

Dr. Pat Federico
Code 511
NPRDC
San Diego, CA 92152-6800

Dr. Paul Feltovich
Southern Illinois University
School of Medicine
Medical Education Department
P.O. Box 3926
Springfield, IL 62708

Mr. Wallace Feurzeig
Educational Technology
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02238

Dr. Craig J. Fields
ARPA
1400 Wilson Blvd.
Arlington, VA 22209

Dr. Gerhard Fischer
University of Colorado
Department of Computer Science
Boulder, CO 80309

J. D. Fletcher
9931 Corsica Street
Vienna VA 22180

Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

Dr. Norman Frederiksen
Educational Testing Service
Princeton, NJ 08541

Dr. Michael Friendly
Psychology Department
York University
Toronto ONT
CANADA M3J 1P3

Dr. Michael Genesereth
Stanford University
Computer Science Department
Stanford, CA 94305

Dr. Herbert Ginsburg
Teachers College
Columbia University
525 West 121st Street
New York, NY 10027

Lee Gladwin
Route 3 -- Box 225
Winchester, VA 22601

Dr. Robert Glaser
Learning Research
& Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Arthur M. Glenberg
University of Wisconsin
W. J. Brogden Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

Dr. Marvin D. Glock
13 Stone Hall
Cornell University
Ithaca, NY 14853

Dr. Sam Glucksberg
Department of Psychology
Princeton University
Princeton, NJ 08540

Dr. Daniel Gopher
Industrial Engineering
& Management
TECHNION
Haifa 32000
ISRAEL

Dr. Sherrille Gott
AFHRL/MODJ
Brooks AFB, TX 78235

Dr. James G. Greeno
University of California
Berkeley, CA 94720

Dr. Dik Gregory
Behavioral Sciences Division
Admiralty Research
Establishment
Teddington
Middlesex, ENGLAND

Prof. Edward Haertel
School of Education
Stanford University
Stanford, CA 94305

Dr. Henry M. Halff
Halff Resources, Inc.
4918 33rd Road, North
Arlington, VA 22207

Dr. Ronald K. Hambleton
Prof. of Education & Psychology
University of Massachusetts
at Amherst
Hills House
Amherst, MA 01003

Dr. Wayne Harvey
Center for Learning Technology
Educational Development Center
55 Chapel Street
Newton, MA 02160

Dr. Barbara Hayes-Roth
Department of Computer Science
Stanford University
Stanford, CA 95305

Dr. Frederick Hayes-Roth
Teknowledge
525 University Ave.
Palo Alto, CA 94301

Dr. Joan I. Heller
505 Haddon Road
Oakland, CA 94606

Dr. Shelly Heller
Department of Electrical Engi-
neering & Computer Science
George Washington University
Washington, DC 20052

Dr. Per Helmersen
University of Oslo
Department of Psychology
Box 1094
Oslo 3, NORWAY

Dr. John Holland
University of Michigan
2313 East Engineering
Ann Arbor, MI 48109

Dr. Thomas Holzman
Lockheed Georgia
Dept. 64-31
Zone 278
Marietta, GA 30063

Ms. Julia S. Hough
Lawrence Erlbaum Associates
6012 Greene Street
Philadelphia, PA 19144

Dr. James Howard
Dept. of Psychology
Human Performance Laboratory
Catholic University of
America
Washington, DC 20064

Dr. Barbara Hutson
Virginia Tech
Graduate Center
2990 Telestar Ct.
Falls Church, VA 22042

Dr. Alice Isen
Department of Psychology
University of Maryland
Catonsville, MD 21228

Dr. R. J. K. Jacob
Computer Science and Systems
Code: 7590
Information Technology Division
Naval Research Laboratory
Washington, DC 20375

Neil Jacobstein
Manager, Research and
Advanced Development
Teknowledge, Inc.
525 University Ave.
Palo Alto, CA 94301-1982

COL Dennis W. Jarvi
Commander
AFHRL
Brooks AFB, TX 78235-5601

Dr. Robin Jeffries
Hewlett-Packard Laboratories
P.O. Box 10490
Palo Alto, CA 94303-0971

CDR Tom Jones
ONR Code 125
800 N. Quincy Street
Arlington, VA 22217-5000

Mr. Daniel B. Jones
U.S. Nuclear Regulatory
Commission
Division of Human Factors
Safety
Washington, DC 20555

Dr. Douglas H. Jones
Thatcher Jones Associates
P.O. Box 6640
10 Trafalgar Court
Lawrenceville, NJ 08648

Dr. Jane Jorgensen
University of Oslo
Institute of Psychology
Box 1094, Blindern
Oslo, NORWAY

Dr. Ruth Kanfer
University of Minnesota
Department of Psychology
Elliott Hall
75 E. River Road
Minneapolis, MN 55455

Dr. Milton S. Katz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Frank Keil
Department of Psychology
Cornell University
Ithaca, NY 14850

Dr. Wendy Kellogg
IBM T. J. Watson Research Ctr.
P.O. Box 218
Yorktown Heights, NY 10598

Dr. Dennis Kibler
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. Peter Kincaid
Training Analysis
& Evaluation Group
Department of the Navy
Orlando, FL 32813

Dr. Walter Kintsch
Department of Psychology
University of Colorado
Campus Box 345
Boulder, CO 80302

Dr. Paula Kirk
Oakridge, TN 37831-0117

Dr. David Klahr
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Janet L. Kolodner
Georgia Institute of Technology
School of Information
& Computer Science
Atlanta, GA 30332

Dr. David H. Krantz
2 Washington Square Village
Apt. # 15J
New York, NY 10012

Dr. Benjamin Kuipers
University of Texas at Austin
Department of Computer Sciences
T.S. Painter Hall 3.28
Austin, TX 78712

Dr. David R. Lambert
Naval Ocean Systems Center
Code 441T
271 Catalina Boulevard
San Diego, CA 92152-6800

Dr. Pat Langley
University of California
Department of Information
& Computer Science
Irvine, CA 92717

Dr. Marcy Lansman
University of North Carolina,
The L. L. Thurstone Lab.
Davie Hall 013A
Chapel Hill, NC 27514

Dr. R. W. Lawler
ARI 6 S 10
5001 Eisenhower Avenue
Alexandria, VA 22333-5600

Dr. Alan M. Lesgold
Learning Research and
Development Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Jim Levin
Department of
Educational Psychology
210 Education Building
1310 South Sixth Street
Champaign, IL 61820-6990

Dr. John Levine
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Michael Levine
Educational Psychology
210 Education Bldg.
University of Illinois
Champaign, IL 61801

Dr. Clayton Lewis
University of Colorado
Department of Computer Science
Campus Box 430
Boulder, CO 80309

Matt Lewis
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Don Lyon
P. O. Box 44
Higley, AZ 85236

Vern Malec
NPRDC, Code P-306
San Diego, CA 92152-6800

Dr. Jane Malin
Mail Code SR 111
NASA Johnson Space Center
Houston, TX 77058

Dr. William L. Maloy
Chief of Naval Education
and Training
Naval Air Station
Pensacola, FL 32508

Dr. Elaine Marsh
Naval Research Laboratory
Code 7510
4555 Overlook Avenue, Southwest
Washington, DC 20375-5000

Dr. Sandra P. Marshall
Dept. of Psychology
San Diego State University
San Diego, CA 92182

Dr. Richard E. Mayer
Department of Psychology
University of California
Santa Barbara, CA 93106

Dr. Gail McKoon
CAS/Psychology
Northwestern University
1859 Sheridan Road
Evanston, IL 60201

Dr. Joe McLachlan
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. James S. McMichael
Navy Personnel Research
and Development Center
Code 05
San Diego, CA 92152

Dr. Barbara Means
Human Resources
Research Organization
1100 South Washington
Alexandria, VA 22314

Dr. Douglas L. Medin
Department of Psychology
University of Illinois
603 E. Daniel Street
Champaign, IL 61820

Dr. Jose Mestre
Department of Physics
Hasbrouck Laboratory
University of Massachusetts
Amherst, MA 01003

Dr. Al Meyrowitz
Office of Naval Research
Code 1133
800 N. Quincy
Arlington, VA 22217-5000

Dr. Ryszard S. Michalski
University of Illinois
Department of Computer Science
1304 West Springfield Avenue
Urbana, IL 61801

Prof. D. Michie
The Turing Institute
36 North Hanover Street
Glasgow G1 2AD, Scotland
UNITED KINGDOM

Dr. George A. Miller
Department of Psychology
Green Hall
Princeton University
Princeton, NJ 08540

Dr. Lance Miller
IBM-FSD Headquarters
6600 Rockledge Drive
Bethesda, MD 20817

Dr. Andrew R. Molnar
Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

Dr. William Montague
NPRDC Code 13
San Diego, CA 92152-6800

Mr. Melvin D. Montemerlo
NASA Headquarters
Washington, DC 20546

Dr. Nancy Morris
Search Technology, Inc.
5550-A Peachtree Parkway
Technology Park/Summit
Norcross, GA 30092

Dr. Randy Mumaw
Program Manager
Training Research Division
HumRRO
1100 S. Washington
Alexandria, VA 22314

Dr. Allen Munro
Behavioral Technology
Laboratories - USC
1845 S. Elena Ave., 4th Floor
Redondo Beach, CA 90277

Dr. David Navon
Institute for Cognitive Science
University of California
La Jolla, CA 92093

Mr. William S. Neale
HQ ATC/TTA
Randolph AFB, TX 78150

Dr. T. Niblett
The Turing Institute
36 North Hanover Street
Glasgow G1 2AD, Scotland
UNITED KINGDOM

Dr. A. F. Norcio
Computer Science and Systems
Code: 7590
Information Technology Division
Naval Research Laboratory
Washington, DC 20375

Commanding Officer,
Naval Research Laboratory
Code 2627
Washington, DC 20390

Dr. Harold F. O'Neil, Jr.
School of Education - WPH 801
Department of Educational
Psychology & Technology
University of Southern
California
Los Angeles, CA 90089-0031

Dr. Michael Oberlin
Naval Training Systems Center
Code 711
Orlando, FL 32813-7100

Dr. James V. Olsen
Director,
Waterford Testing Center
1681 West 820 North
Provo, UT 84601

Office of Naval Research,
Code 1133
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142BI
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142
800 N. Quincy St.
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142PS
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 1142CS
800 N. Quincy Street
Arlington, VA 22217-5000
(6 Copies)

Special Assistant for Marine
Corps Matters,
ONR Code 00MC
800 N. Quincy St.
Arlington, VA 22217-5000

Dr. Judith Orasanu
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

CDR R. T. Parlette
Chief of Naval Operations
OP-112G
Washington, DC 20370-2000

Dr. James Paulson
Department of Psychology
Portland State University
P.O. Box 751
Portland, OR 97207

Dr. Douglas Pearce
DCIEM
Box 2000
Downsview, Ontario
CANADA

Dr. Virginia E. Pendergrass
Code 711
Naval Training Systems Center
Orlando, FL 32813-7100

Military Assistant for Training
and
Personnel Technology,
OUSD (R & E)
Room 3D129, The Pentagon
Washington, DC 20301-3080

LCDR Frank C. Petho, MSC, USN
CNATRA Code N36, Bldg. 1
NAS
Corpus Christi, TX 78419

Dr. Steven Pinker
Department of Psychology
E10-018
M.I.T.
Cambridge, MA 02139

Dr. Tjeerd Plomp
Twente University of Technology
Department of Education
P.O. Box 217
7500 AE ENSCHEDE
THE NETHERLANDS

Dr. Peter Polson
University of Colorado
Department of Psychology
Boulder, CO 80309

Dr. Steven E. Poltrock
MCC,
Human Interface Program
3500 West Balcones Center Dr.
Austin, TX 78759

Dr. Mary C. Potter
Department of Psychology
MIT (E-10-032)
Cambridge, MA 02139

Dr. Joseph Psotka
ATTN: PERI-1C
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333

Dr. James A. Reggia
University of Maryland
School of Medicine
Department of Neurology
22 South Greene Street
Baltimore, MD 21201

Dr. Wesley Regian
AFHRL/MOD
Brooks AFB, TX 78235

Dr. Gil Ricard
Mail Stop C04-14
Grumman Aerospace Corp.
Bethpage, NY 11714

Mark Richer
1041 Lake Street
San Francisco, CA 94118

William Rizzo
Code 712
Naval Training Systems Center
Orlando, FL 32813

Dr. Linda G. Roberts
Science, Education, and
Transportation Program
Office of Technology Assessment
Congress of the United States
Washington, DC 20510

Dr. Ernst Z. Rothkopf
AT&T Bell Laboratories
Room 2D-456
600 Mountain Avenue
Murray Hill, NJ 07974

Dr. William B. Rouse
Search Technology, Inc.
5550-A Peachtree Parkway
Technology Park/Summit
Norcross, GA 30092

Dr. Roger Schank
Yale University
Computer Science Department
P.O. Box 2158
New Haven, CT 06520

Dr. Janet Schofield
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Karen A. Schriver
Department of English
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Hans-Willi Schroiff
Institut fuer Psychologie
der RWTH Aachen
Jaegerstrasse zwischen 17 u. 19
5100 Aachen
WEST GERMANY

Dr. Judith Segal
OERI
555 New Jersey Ave., NW
Washington, DC 20208

Dr. Robert J. Seidel
US Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333

Dr. Colleen M. Seifert
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UKSD
La Jolla, CA 92093

Dr. Ramsay W. Selden
Assessment Center
CCSSO
Suite 379
400 N. Capitol, NW
Washington, DC 20001

Dr. Daniel Sewell
Search Technology, Inc.
5550-A Peachtree Parkway
Technology Park/Summit
Norcross, GA 30092

Dr. Michael G. Shafto
ONR Code 1142CS
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. Sylvia A. S. Shafto
Department of
Computer Science
Towson State University
Towson, MD 21204

Dr. Ben Schneiderman
Dept. of Computer Science
University of Maryland
College Park, MD 20742

Dr. Ted Shortliffe
Computer Science Department
Stanford University
Stanford, CA 94305

Dr. Valerie Shute
AFHRL/MOE
Brooks AFB, TX 78235

Mr. Raymond C. Sidorsky
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Robert S. Siegler
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

LTCOL Robert Simpson
Defense Advanced Research
Projects Administration
1400 Wilson Blvd.
Arlington, VA 22209

Dr. Zita M Simutis
Instructional Technology
Systems Area
ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Derek Sleeman
Dept. of Computing Science
King's College
Old Aberdeen
AB9 2UB
UNITED KINGDOM

Dr. Gail Slemmon
Logicon
P.O. Box 85158
San Diego, CA 92138

Dr. Linda B. Smith
Department of Psychology
Indiana University
Bloomington, IN 47405

Dr. Alfred F. Smode
Senior Scientist
Code 07A
Orlando, FL 32813

Dr. Richard E. Snow
Department of Psychology
Stanford University
Stanford, CA 94306

Dr. Elliot Soloway
Yale University
Computer Science Department
P.O. Box 2158
New Haven, CT 06520

Dr. Richard Sorensen
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Paul Speckman
University of Missouri
Department of Statistics
Columbia, MO 65201

Dr. Kathryn T. Spoehtl
Brown University
Department of Psychology
Providence, RI 02912

Dr. Marian Stearns
SRI International
333 Ravenswood Ave.
Room B-S124
Menlo Park, CA 94025

Dr. Frederick Steinheiser
CIA-ORD
612 Ames
Washington, DC 20505

Dr. Albert Stevens
Bolt Beranek & Newman, Inc.
10 Moulton St.
Cambridge, MA 02238

Dr. David Stone
KAJ Software, Inc.
3420 East Shea Blvd.
Suite 161
Phoenix, AZ 85028

Dr. John Tangney
AFOSR/NL
Bolling AFB, DC 20332

Dr. Kikumi Tatsuoka
CERL
252 Engineering Research
Laboratory
Urbana, IL 61801

Dr. Martin M. Taylor
DCIEM
Box 2000
Downsview, Ontario
CANADA

Dr. Perry W. Thorndyke
FMC Corporation
Central Engineering Labs
1185 Coleman Avenue, Box 580
Santa Clara, CA 95052

Major Jack Thorpe
DARPA
1400 Wilson Blvd.
Arlington, VA 22209

Dr. Sharon Tkacz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Martin A. Tolcott
3001 Veazey Terr., N.W.
Apt. 1617
Washington, DC 20008

Dr. Douglas Towne
Behavioral Technology Labs
1845 S. Elena Ave.
Redondo Beach, CA 90277

Dr. Kurt Van Lehn
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Jerry Vogt
Navy Personnel R&D Center
Code 51
San Diego, CA 92152-6800

Dr. Ming-Mei Wang
Lindquist Center
for Measurement
University of Iowa
Iowa City, IA 52242

Roger Weissinger-Baylon
Department of Administrative
Sciences
Naval Postgraduate School
Monterey, CA 93940

Dr. Donald Weitzman
MITRE
1820 Dolley Madison Blvd.
MacLean, VA 22102

Dr. Keith T. Wescourt
FMC Corporation
Central Engineering Labs
1185 Coleman Ave., Box 580
Santa Clara, CA 95052

Dr. Douglas Wetzel
Code 12
Navy Personnel R&D Center
San Diego, CA 92152-6800

LCDR Cory deGroot Whitehead
Chief of Naval Operations
OP-112G1
Washington, DC 20370-2000

Dr. Heather Wild
Naval Air Development
Center
Code 6021
Warminster, PA 18974-5000

Dr. William Clancey
Stanford University
Knowledge Systems Laboratory
701 Welch Road, Bldg. C
Palo Alto, CA 94304

Dr. Michael Williams
IntelliCorp
1975 El Camino Real West
Mountain View, CA 94040-2216

A. E. Winterbauer
Research Associate
Electronics Division
Denver Research Institute
University Park
Denver, CO 80208-0454

Dr. Robert A. Wisher
U.S. Army Institute for the
Behavioral and Social
Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Frank Withrow
U.S. Office of Education
400 Maryland Ave. SW
Washington, DC 20202

Mr. John H. Wolfe
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Dan Wolz
AFHRL/MOE
Brooks AFB, TX 78235

Dr. George Wong
Biostatistics Laboratory
Memorial Sloan-Kettering
Cancer Center
1275 York Avenue
New York, NY 10021

Dr. Wallace Wulfeck, III
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Joe Yasatuke
AFHRL/LRT
Lowry AFB, CO 80230

Mr. Carl York
System Development Foundation
181 Lytton Avenue
Suite 210
Palo Alto, CA 94301

Dr. Joseph L. Young
Memory & Cognitive
Processes
National Science Foundation
Washington, DC 20550

Dr. Steven Zornetzer
Office of Naval Research
Code 114
800 N. Quincy St.
Arlington, VA 22217-5000

END

4-87

D T I C