DOCUMENT RESUME

ED 253 858 CS 007 936

AUTHOR	Kieras, David E.
TITLE	The Role of Prior Knowledge in Operating Equipment from Written Instructions. Final Report.
INSTITUTION	Michigan Univ., Ann Arbor.
SPONS AGENCY	Office of Naval Research, Arlington, Va. Personnel and Training Research Programs Office.
REPORT NO	FR-85/ONR-19
PUB DATE	20 Feb 85
CONTRACT	N00014-84-K-0731
NOTE	36p.
PUB TYPE	Reports - Research/Technical (143)
EDRS PRICE	MF01/PC02 Plus Postage.
DESCRIPTORS	Adults; *Equipment Utilization; Knowledge Level; Memory; *Prior Learning; *Reading Comprehension; *Reading Research; *Transfer of Training; *Writing Skills
IDENTIFIERS	*Direction Following

ABSTRACT

A series of experiments investigated the role of prior knowledge in tasks involving the operation of equipment from written instructions. The experiments covered two situations. In the first, the prior knowledge was already possessed by the subjects before the experiments. The studies involved comprehension and memory of technical prose, expertise in descriptions of familiar and unfamiliar pieces of equipment, and expertise effects in following instructions that differed in organization. In the second situation, the prior knowledge was provided as part of the training involved in the experiments. These studies concerned the role of knowledge of how a system works and transfer of training from previously learned operating procedures to new procedures. The results support methodological, theoretical and practical conclusions. Methodologically, many traditional prose recall paradigms should be used with caution in the investigation of prior knowledge, and careful attention should be paid to the relationship between the knowledge being supplied to the subject and the exact tasks that the subject is expected to perform. Theoretically, the results support what is perhaps becoming the consensus model of cognitive architecture, namely the ACT class of theories described in Anderson's most recent textbook, the "Architecture of Cognition" (1983). Practically, the results provide a good foundation for future applied research on the arrangement, sequence and content of instructional materials. (FL)

ED253858

THE ROLE OF PRIOR KNOWLEDGE IN OPERATING EQUIPMENT FROM WRITTEN INSTRUCTIONS

FINAL REPORT

David E. Kieras

University of Michigan

Report No. 19 (FR-85/ONR-19)

February 20, 1985

This research was supported by the Personnel and Training Research Programs, Office of Naval Research, under Contract Number N00014-84-K-0731, Contract Authority Identification Number NR 667-473. Reproduction in whole or in part is permitted for any purpose of the United States Government.

Approved for Public Release; Distribution Unlimited.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER FR-85/ONR-19	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) The role of prior knowledge in operating equipment from written instructions		5. TYPE OF REPORT & PERIOD COVERED Final Report 1 Sept. 1981-28 Feb.1985
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) David E. Kieras		8. CONTRACT OR GRANT NUMBER(s) N00014-84-K-0731
9. PERFORMING ORGANIZATION NAME AND ADDRESS University of Michigan Ann Arbor, MI		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61153N; RR 042-06 RR 042-06-02; NR 667-473
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, VA 22217		12. REPORT DATE February 20, 1985
		13. NUMBER OF PAGES 42
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for Public Release; Distribution Unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Instructions, expertise, comprehension, training.

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This final report summarizes a set of results on the role of prior knowledge in how people operate electronic equipment from written instructions. These results cover two situations: In the first, the prior knowledge is possessed by subjects prior to the experiment. These studies involved comprehension and memory of technical text, expertise in descriptions of familiar and unfamiliar pieces of equipment, and expertise

DD FORM 1473 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE S/N 0102-014-6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered) effects in following instructions that differ in organization. In the second situation, the prior knowledge was provided as part of the training involved in the experiments. These studies concerned the role of knowledge of how a system works, and transfer of training from previously-learned operating procedures to new procedures. Simulation models were constructed and compared in detail to the data, yielding significant theoretical conclusions about the mechanisms involved in the effective use of prior knowledge. The work has considerable practical significance for the design and evaluation of materials and documentation for equipment training and maintenance.

ABSTRACT

This final report summarizes a set of results on the role of prior knowledge in how people operate electronic equipment from written instructions. These results cover two situations: In the first, the prior knowledge is possessed by subjects prior to the experiment. These studies involved comprehension and memory of technical text, expertise in descriptions of familiar and unfamiliar pieces of equipment, and expertise effects in following instructions that differ in organization. In the second situation, the prior knowledge was provided as part of the training involved in the experiments. These studies concerned the role of knowledge of how a system works, and transfer of training from previously-learned operating procedures to new procedures. Simulation models were constructed and compared in detail to the data, yielding significant theoretical conclusions about the mechanisms involved in the effective use of prior knowledge. The work has considerable practical significance for the design and evaluation of materials and documentation for equipment training and maintenance.

The Role of Prior Knowledge in Operating Equipment from Written Instructions

David E. Kieras
University of Michigan

INTRODUCTION

The goal of this project was the study of the role of prior knowledge in tasks involving the operation of equipment from written instructions. The work has been done in two categories: In the first, prior knowledge was not manipulated; rather the prior knowledge under investigation was the knowledge possessed by the person prior to being in the experimental situation. This project began by investigating the role of prior knowledge in the comprehension of simple technical prose, since a large body of research has been done on how people learn information from written prose, and the previous ONR project focussed on comprehension of descriptive technical text. This work was followed by a set of studies on what experts know about electronic equipment compared to non-experts, along with a comparison of their strategies for following instructions for operating a piece of electronic equipment.

In the second category of work, prior knowledge was manipulated in experiments in which subjects were given information about equipment in a fixed sequence, with the question being how later learning was influenced by the previously acquired knowledge. One major topic in this category is the role of how it works knowledge, which is the knowledge that a person might have about the internal structure and functioning of a piece of equipment. It is unclear whether having such knowledge is beneficial, and if so, under what conditions it is. A series of experimental studies and a computer simulation modeling effort provided some answers. Since much equipment operation is learned under rote learning conditions, a second topic was the effects of prior knowledge of procedures for operating a piece of equipment. This work arrived at a precise theoretical description of knowledge of procedures and transfer of training between procedures.

PRE-EXPERIMENTAL PRIOR KNOWLEDGE

Prior Knowledge and Learning from Technical Prose

Background. We all have the strong intuition that it should be easier to learn from material if one is already familiar with the general subject matter area. However, there have been very few studies showing that this was true in a convincing way, and in fact, some reports have suggested the opposite. This unclear state of empirical affairs presents a serious theoretical problem, because all of the current theories of comprehension assume that new information from the material is understood in terms of prior knowledge, and then is integrated with this prior knowledge. Thus, the availability of relevant prior knowledge should have profound effects on the process of learning information from text. The unclear state of the empirical literature is probably due to the fact that it is very difficult to conduct a well-controlled study in which the amount of true prior knowledge possessed by subjects is the variable of interest. Manipulating the amount of prior knowledge within an experiment is not a good approach, because there is ample reason to believe from the experimental literature that such experimentally acquired knowledge does not have the same properties as the true prior knowledge that subjects have before coming to an experiment.

Approach. The approach was based on the logic of quasi-experimental design, in which the sampling of subjects is used to vary the amount of prior knowledge, and statistical methods, mainly multiple regression, were used to control nuisance variables such as general reading ability, word frequency and so forth. The materials were several passages that varied widely in their basic content familiarity both within passages and between passages. This insured that individual subjects would be almost certain to vary in the amount of prior knowledge that they had of the material. Each subject was tested to determine which passage facts they already knew.

Since subjects can change their reading behavior substantially as a function of the reading task, three different reading tasks were used, and measures were collected not only of the amount of information recalled, but also of the time spent reading or studying the material. In one task, subjects were allowed to study the material for as long as they chose, on a self-paced sentence-by-sentence basis, with the knowledge that they would be tested for recall later. In a second reading task, the subjects knew that they would be tested for later recall, but were allowed to study each sentence for a fixed amount of time. In the third test, subjects did not know that they would be tested for later recall, but instead read the passages a sentence at a time in order to identify the topic of the passage. All subjects were then tested for recall.

Results. As expected, but not previously demonstrated, there were clear effects of the amount of prior knowledge. However, the role of prior knowledge depended substantially on the reading task. If subjects knew they would have to recall the material, they spent more time studying unfamiliar portions of the material than they did on familiar portions. But in both cases, they recalled the information at the same level, regardless of familiarity. But if subjects were unaware of the requirement for later recall, or were limited in the amount of time they could study, then unfamiliar material was recalled less than familiar.

The quantitative size of the effects on study time could be accounted for by a simulation model that was based on the principle that representations of previously known information do not have to be constructed from "scratch", but rather, the previously present representation could simply be "tagged" as appearing in a particular passage context. This means that known facts do not have to be subjected to memorization processes to the same extent as unknown facts, meaning that study time would be a function of the amount of unknown information.

The simulation model was based on standard concepts in comprehension theory, being based on an augmented transition network parser that constructs semantic network representations. After each sentence was analyzed, the simulation would compare the content of the sentence with the contents of long term memory, and identify those portions of knowledge structure that were already present in long term memory. Then, only the new structure would have to be added to memory.

Significance. This work provided a demonstration of a theoretically important effect which had not appeared in the literature. It also showed that the effects of prior knowledge could be accounted for using standard theoretical concepts developed in comprehension research. One important substantive result is that readers who are experienced students have mnemonic strategies that are powerful enough to deal with extremely unfamiliar material. Another result is that a relatively simple interpretation of how prior knowledge is used in comprehension is viable, as well as more complex, and currently popular, notions based on schema theory (e.g., Rumelhart, 1980).

However, perhaps more importantly, this work made it clear that how prior knowledge is used in comprehension can be highly dependent on the task required of the reader. Standard recall experiments could give highly misleading results, because subjects can approach the task in a way that can eliminate differences between familiar and unfamiliar information.

The original approach planned in this project was to use standard comprehension paradigms to investigate how passages about equipment were processed in terms of the reader's prior knowledge. However, these results strongly suggested that the tasks used in such experiments would yield results that were either weak, or tied so directly to these specific tasks that they would not be directly relevant to the actual tasks involved in operating equipment. For this reason, the remaining experiments in the project always required the subject to engage in a relatively realistic task involving operating a piece of equipment.

Publications. The empirical work was described in detail in Technical Report Number 11 (Johnson and Kieras, 1982), and the simulation model was described in detail in Kieras (1983). The empirical work and the simulation modeling of the effects appeared in the archival publication Johnson and Kieras (1983). The methodology used to compare the simulation to the data, both for this study, and others discussed in this report, was described in Kieras (1984a). Further discussion of the theoretical mechanisms of task effects will appear in Kieras (in preparation).

Expert Knowledge of Equipment

Background. While there has been considerable research on the nature of expertise, there has been very little study of expertise with regard to equipment, especially actual electronic equipment. Compared to other domains, the domain of equipment has some important psychological properties. First of all, rather than dealing with abstract concepts, as in many fields of expertise that have been studied, expertise with equipment deals with physically concrete objects. Second, equipment is often very complex in terms of the different possible levels of analysis and types of information that can be involved. For example, knowledge of equipment ranges from the typical colors with which the equipment is painted, all the way to the physical principles involved in how the equipment operates. Third, equipment is something that a person interacts with, and which can have its own internal states and rules of behavior. Thus, unlike some forms of expertise which deal only with a person's skill, expertise with equipment involves not only the skill of operating the equipment, but also how the equipment will behave in response to what the operator does with it. Thus knowledge about equipment can have some important features that distinguish it from other knowledge domains.

Notice that the domain of electronic devices is very suited for the exploration of schema theory (see Rumelhart, 1980; Rumelhart & Ortony, 1977). A schema is an organized body of knowledge that represents a stereotyped, or a frequently ocurring, pattern of events that is used to organize perceptual and memory processes. For example, device schemas apparently conform to very concrete specific physical features of the device and their relationships. Devices themselves have a very strong hierarchical structure, since the entire device is made up of sub-devices. Furthermore, each sub-device normally has schematic features as well. For example, many electronic devices include some sort of audio amplifier connected to some sort of transducer, such as a loud speaker. Such devices almost always have a cluster of features that correspond to this common sub-device. For example, there will be a volume control, quite often a tone control, and a speaker which is normally a circular object behind a perforated grill. While details of the placement and appearance of these features will vary, quite often the two controls will be located adjacent to each other, normally clustered with other controls, and the speaker will usually be facing toward the user. The volume and tone controls are very rarely separated from each other, and the speaker almost never appears on the top surface of a device. Because of the strength and concreteness of these patterns, it would probably be much easier to construct a schema theory in this domain than in the traditional domains where the concept has been used, such as in the understanding of stories.

Approach. Both experts and non-experts were used in these studies. The experts were individuals who had years of experience in electronics, many of them being former military electronics technicians. The non-experts were ordinary undergraduate and graduate students with no special background. The task was to provide oral descriptions of a piece of equipment which was placed in front of the subject. The subject was videotaped, and the descriptions transcribed and analyzed for content. The equipment consisted of several devices, ranging from everyday items such as a tape recorder, to very specialized equipment. In a follow-up study, ordinary student subjects were asked to produce descriptions from memory of several everyday pieces of electronic equipment, such as a television set. The subjects were asked to produce information in several categories which were chosen on the basis of the previous experiment results. These responses were analyzed for content.

The value of the schema concept was explored with small-scale simulation models for how a device would be recognized in terms of schemas.

Results. The results of the first study showed that for both experts and non-experts, knowledge of electronic devices classifies naturally into the categories of: (1) the function of the device; (2) the operating procedures; (3) how the device works internally; (4) how the device behaves externally; and (5) the power source of the device. The knowledge appeared to be organized in terms of a hierarchy, in that the categories of function, operation, and how-it-works were applied recursively to not just the device as a whole, but to each of its controls and other external features.

There was strong evidence that the knowledge was organized in terms of schemas, in that there were many cases in which the subjects manifested having definite expectations about device features, and knowledge of general conventional patterns of features. The devices were recognized and categorized almost immediately by key patterns of features. This is exactly the mechanism that schema organization would entail. In some cases, involving unusual devices, an incorrect schema can apparently be triggered by a subset of features, leading to serious misperceptions of the device as a whole. For example, one expert was confused by an unusual device that was actually a form of signal generator, but which had its output connectors on the left-hand side of the front panel, a position customarily used for input connectors.

A surprising result was the prominence of procedural information as opposed to how-it-works information, which even expert subjects tended not to produce. Another surprise was the high frequency of mention of the power source of the device. Finally, there were many obvious features of the devices that subjects often failed to mention, such as the fact that an electric clock has hands on a dial numbered from 1 to 12.

The second study clarified these questions by prompting subjects for descriptions from memory of everyday devices. They were asked to describe the function of each device, the features such as dials and indicators which they used to recognize the device, and, in addition, the features one would expect to see on the device, how the device was operated, and how the device works inside. There was high agreement between subjects about the function, recognition features, expected features, and operating procedures. The surprising prominence of power source information reappeared in these results, which suggests that power sources, being a common feature of many different electronic devices, plays an important role in this knowledge domain.

The how-it-works knowledge, was sketchy and inconsistent, as in the first experiment, although subjects did produce a considerably greater quantity than they did in the first experiment. Note that unlike the function and features of a device, which are frequently experienced and concrete, the how-it-works knowledge is generally abstract or "invisible," being hidden inside the device, and normally not involved in the routine use of the device. For complex devices such as televison sets, the how-it-works knowledge was less consistent and sketchier than for the simpler devices. Thus, the lack of how-it-works descriptions in the first study was artifactual, at least to some extent, because for the simplest devices there was a reasonable amount of such information produced. Perhaps the face-to-face interaction and task demands of describing orally a presented device biased subjects against producing much of this knowledge in the first study.

The hypothesis that device knowledge is organized in terms of schemas was further confirmed by the fact that the recognition features and expected features were distinct sets, even though there was some overlap. In terms of schema theory, certain features would trigger the activation of a schema, which would then make the expected features available.

These results also contain many of the effects attributed to schemas, such as expectations on the part of the perceiver, and confusion when confronted by a stimulus that does not quite match a schema. However, it is hard to define how this characterization could be made more precise on the basis of empirical data; the research would quickly become a matter of simply surveying the population stereotypes for various pieces of equipment. While this might be useful, it would not advance the theoretical concepts very far.

Some simulation modelling of schema processes was done. A set of production rules were written that would recursively instantiate the schema for a device, by first recognizing low level schemas for individual clusters of features on the device. These sub-schemas would then be slot-fillers for higher level schemas. After working its way up two or three levels, the schema for the entire device could then be instantiated. Working from a different tack, other production rules could recognize particular patterns of distinguishing features of a device, and directly instantiate the corresponding schema. For example, a box with an antenna and a tuning dial is almost certainly a radio. While this model was promising, extending it did not have an obvious direction. It is hard to say what any empirical consequences of such a model might be, other than the obvious ones already described. However, it is important to note that much of schema theory is not very well developed at the level of rigorous simulation models. Thus, further work along these lines might help firm up the theory.

Significance. A major lesson of these results is that the knowledge that people have about electronic devices is incredibly rich and detailed, encompassing a very broad range of kinds of information. Perhaps there has been some tendency for psychologists and educators to assume that equipment was relatively simple, in the sense that there was relatively little one needed to know about it in order to use it. While in some sense this may be true, it is also clear that people know a very large amount of information of several different types about equipment.

There are many facets of knowledge about equipment that could be explored with further research. For example, from the results it is clear that familiar items of equipment often have stereotypical layouts of the controls, indicators, and so forth. One might wonder if this is true of all classes of equipment, or just those that are very common. For example, does the extremely specialized equipment used in the military also follow stereotypical patterns in its control layout? Another related question is whether sterotypical external layouts of equipment are in fact related to the internal structure of the equipment. For example, one intuition is that the external features of a device that are most closely related to the purpose of the equipment are often large and centrally located. An example is that measuring instruments usually have a large meter centrally placed on the front panel. A signal generator usually has a large dial centrally positioned. How will the user react to a signal generator that for design reasons has a small dial located in an off-center position? Note that modern test equipment, in which digital circuitry and displays are used heavily, often seem to depart seriously from previous customs on the placement of controls and indicators. For example, the front panel layout of service oscilloscopes has been fairly standard, because to a great extent because this layout corresponded to the optimum arrangement of the traditional large and bulky vacuum tube circuitry. Newer, more compact circuitry can be arranged in many different ways, thus possibly leading to front panel arrangements that are no longer familiar to most users.

Publications. These results are described in Technical Report No. 12 (Kieras, 1982a), and were discussed in a presentation at the Joint Services Workshop on Artificial Intelligence Applications to Maintenance, whose proceedings were published (see Kieras, 1984b).

Expertise Effects in Following Instructions

Background. A very common task is one in which the user of a piece of equipment follows step-by-step instructions for operating it. Sometimes the goal of the user is to learn the procedure, but often the user is in a strictly one-shot situation, meaning that the user is not trying to learn the procedure, only follow it once. Theoretically, the user must obtain specific pieces of procedural knowledge from the individual instruction steps, and then immediately execute the procedural knowledge. While this should be a very simple task, it is clear from everyday experience that instructions may not be written well enough to be followed easily. Furthermore, the experience of the user should play some role in the ability to follow instructions correctly. This study was designed to answer three questions about how people function in a task involving operating a piece of equipment from written instructions.

The first question concerned the instruction format. Two formats were examined. In both instruction formats, the overall task of the subject was to get the device into a specified state; this task was stated at the beginning of the instructions. The first format was step-by-step instructions, in which each step concerned the setting of a individual control on the equipment. The subject had to read each step in the presented order, and was expected to carry it out immediately. The other instruction format was a hierarchical menu. Following the initial task statement, the subject was presented with a menu of choices, each of which consisted of a natural "chunk" in the operation of the device. For example, if the task was to get a radio tuned to a specific station, the first menu would contain the choices of getting the radio powered up, getting the radio tuned to the correct station, and making the final adjustments to the radio controls. At each level of the hierarchy, the subject could either attempt to execute the task using the information available at that level, or could choose to get more detail by selecting one of the choices. This would produce another menu. At the very bottom of the hierarchy was the same step-by-step instructions as in the other condition. Thus, the contrast was between following a linear sequence of steps to operate the controls, or having the same steps arranged at the bottom of a hierarchy that allowed the subject to read only to the level of detail that he or she desired in order to complete the task. Intuition would hold that this highly organized form of instructions would be superior to the linear instructions. Furthermore, Smith and Goodman (1982), found that a similar form of hierarchical organization did improve performance.

The second question was whether there would be substantial differences between experts and non-experts in following instructions. It was expected that experts would be faster overall. However, expertise should also be relative to the specific device being operated. This was examined by including a wide range of electronic devices, ranging from everyday items, such as an ordinary portable radio, to ones familiar only to experts, such as a dual-trace triggered oscilliscope, to devices that would be unfamiliar even to electronics experts, such as a laboratory physiological stimulator, or a unique device.

The third question concerns the nature of the prior knowledge that subjects have about devices. In the work described above, it was concluded that knowledge of devices is organized as schemas, which reflect stereotypical arrangements of events. A straightforward extension of this idea is that knowledge of operating procedures for familiar devices should also have a stereotypical pattern. If people's knowledge of how to operate devices has stereotypical properties, this should be reflected in the pattern of menu choices that people make in the hierarchical instructions condition, or a chunking effect in execution time of the linear instructions. Also, there were many cases in the hierarchical instructions condition in which the device was operated completely from memory; that is, subjects read only the portion of the instructions that stated the overall task. In this case there should be schema-based stereotypicality in the sequence of operations that people performed from memory.

<u>Approach</u>. Several devices were used, and there were two types of subjects: experts, who had extensive electronic

experience, and non-experts who did not. The instruction type was a between subjects manipulation, but each subject operated each one of the several devices. A laboratory computer was used to present the instructions, and measure the time each segment of the instructions was viewed. The subjects' behavior was videotaped in order to permit detailed scoring of the subjects' activities.

Before beginning the experiment, the subjects answered a questionaire in which their familiarity with specific items of equipment was assessed. The major dependent variable was the total time required to complete the task. In the linear instruction condition, the individual time spent on each step was an additional dependent variable. In the menu format condition, the major variables were the time spent on each frame of the instructions, and the choices made in moving down the menu hierarchy. Finally, the sequence of operating the controls was also assessed, although this only has particular value in the menu hierarchy condition, in cases where very few or no frames of instructions were read. Otherwise, the controls operated are almost completely determined by what the presented instruction step actually says to do.

Results. Contrary to intuition, the hierarchical menu format was not superior overall to the linear step-by-step instruction format. The menu format was superior only if the subject was familiar with the type of device, and was sometimes substantially inferior. For example, while the menu condition took less than half the time of the step-by-step condition for operating the radio, in other cases, in which the subject was not familiar with the device, the menu condition could be 39% to 71% slower in total task time than the step-by-step condition.

The basic reason for the menu condition being inferior in some cases appeared to be that subjects would often mistakingly attempt to operate an unfamiliar device on the basis of very little instructions; this "go it alone" approach could sometimes be disasterous. For example, expert subjects often attempted to operate the physiological stimulator only on the basis of the main task statement. The task required connecting an indicator light to the stimulator, but there were several possible connectors where the light could be plugged in. Many expert subjects plugged the indicator lamp into the wrong jack, and then spent a long time trying to accomplish the task. Thus, the intuition that the menu hierarchy is better than the linear sequence must be strongly qualified by whether the user is familiar enough with the equipment to take advantage of the hierarchy, and also whether the subject is likely to attempt to operate the equipment without help from the instructions when the device is unfamiliar.

Experts were faster overall than the non-experts and were able to operate equipment with fewer instructions in the menu conditions. In terms of total time, experts were about one third faster than non-experts. However, there were strong effects of specific experience with the device as well as general expertise effects. Another major expertise effect was that experts were much better at executing complicated physical activities, such as plugging in a cord, and also at complex physical activities that are familiar to experts, such as zeroing a meter. Informally, it appeared that non-expert subjects often would spend a lot of time fumbling with cords and connectors, while experts seem to know exactly how to perform these activities smoothly and precisely. This result is particularly interesting, because it shows that expertise with electronic equipment has other components in addition to cognitive factors.

The results showed that prior knowledge played a specific role in following instructions in the step-by-step condition. Examination of the times to complete individual steps showed that a step that was always read in the menu condition took appreciably longer to execute than one that was never read. Assuming that the menu choices reflect the familiarity of the procedure chunks, the amount of time taken to complete a step is thus a function of its predictability on the basis of prior knowledge.

Unfortunately the information in the pattern of menu choices was very limited. The hope was that the subjects' choice of branches of the menu would reflect their knowledge of stereotypical portions of operating the equipment. This would only be true if the device was familiar to the subject. But, in this case, the subject would need to read very little of the instruction hierarchy. Thus, the subject being familiar with the equipment meant that the subject made very few choices that would reveal that familiarity. Future research along these lines should take this into account. There were, however, some interesting effects in the menu choices. Almost everybody, including the non-experts, knew how to get even the expert equipment powered up, even if it was unfamiliar.

Another interesting effect is that non-experts could learn how to perform repeated activities in the course of following the instructions. For example, on one device, two indicator lights had to be plugged in one after the other. While half of the non-experts read the detailed instructions for plugging in the first light, very few read the instructions for doing the same to the second light. Similar effects showed up in more complicated situations. The implication is that subjects are not simply executing the instructions as they are read, and then simply forgetting the instruction content as they proceed to the next instruction. Rather, they seem to be able to immediately generalize the content of one set of instructions and apply it immediately to a similar situation. This suggests that subjects' ability to induce and generalize procedures is very rapid and powerful; a similar conclusion was reached in entirely different tasks (Kieras and Bovair 1985, Kieras, 1984c).

An important aspect of the role of prior knowledge was revealed by the sequence of actions performed by subjects in the menu condition, considering those cases in which the subject operated the equipment successfully using only the main task statement. It was expected that these sequences of actions would show fairly stereotypical patterns. However, the detailed analysis of the sequences yielded the conclusion that there is in fact very little stereotypical content in the activity. The initial stages of operating some of the devices were fairly patterned; for example, all of the subjects plugged in the radio before operating any other controls. Thus, there is some tendency for the power-up operations to be done prior to other steps. However, there were very few other patterns.

This lack of stereotyped patterns led to the conclusion that subjects operate familiar equipment from memory not by executing "canned" procedures, but by problem-solving within the constraints imposed by the nature of the device. For example, many equipment controls have loose sequential constraints on their operation, but these constraints do not predetermine a particular sequence of operations. Thus, the subjects made many idiosyncratic passes over the controls, and the overall state of the device gradually converges to the desired one. In many cases, the number of control operations performed is considerably more than is technically required.

The best characterization of how a piece of equipment is operated from memory seems to be that people determine what constraints need to be satisfied, and then operate the controls in a manner that meets the constraints and accomplishes the task, but does not necessarily follow any fixed order. Thus, the major prediction of schema theory with regard to how equipment is operated, namely stereotyped sequences of actions, does not appear to hold.

Notice that this is a task situation in which subjects were not specifically trained to operate a piece of equipment, but rather were operating the equipment based upon their general prior knowledge. A distinction should be made between what people do when they have a highly automated skill at operating a particular piece of equipment, a result of intensive training and practice, and the ability to operate equipment in a more general setting, in which each piece of equipment is familiar, but not highly practiced. Under certain conditions, the strategies used by experts may be less effective with unfamiliar equipment than the performance of non-experts who are following strict step-by-step instructions.

Even though this problem-solving approach is rather sub-optimal from a strictly technical point of view, it is in fact very robust. That is, the subjects could apply the same approach to any device within a class with which they were familiar. For example, almost any electronics expert would operate almost any type of volt-ohm-milliammeter. They would simply recognize which constraints have to be satisfied before the desired measurement could be obtained, and would work with the controls until these constraints were satisfied. It is this robustness of expert knowledge which is particularly valuable in operating with equipment. In this experiment the experts could operate some of the completely novel devices without any instuctions, and do so quite often without any serious mistakes or inefficiencies.

Thus, expertise at operating a variety of equipment does not consist of having a set of canned procedures for operating different devices, but rather of having a set of powerful problem-solving heuristics which can be applied to devices that might be unfamiliar, but which may not be very efficient when applied to familiar devices.

Significance. These results are important to the design of equipment maintenance documentation, which is usually used in just the manner explored in this study. The intuition that the hierarchical instructions are clearly better than a strict linear sequence is false; the experience of the user is critical. With regard to training individuals to become expert users of a broad variety of equipment, it should be recognized that teaching strict procedures is probably not the appropriate course. Clearly, if the individual is being trained to operate one piece of equipment under stressful conditions, training specific operating procedures to the point where they become automatically executed is clearly optimum. However, if the individual is being trained to do maintenance work that might involve a large variety of equipment, an understanding of the general constraints involved in successful functioning of the equipment would probably be more productive than attempting to teach specific operating sequences for each individual piece of equipment.

Publications. These results are described in Technical Report No. 14 (Kieras, Tibbitts, & Bovair, 1984), and are also cited in Kieras (1984b).

EXPERIMENTALLY-ACQUIRED PRIOR KNOWLEDGE

Prior Knowledge of How a System Works

Background. There has been a long-standing disagreement over whether users of equipment should be fully informed about how the equipment works inside. Should the user simply be told how to get the job done with the equipment, or should the user be told how the equipment works? For example, the training materials for word processors normally come with extensive discussion of how to accomplish various tasks with the word processor, but normally have little or no discussion or description about how the system itself works.

On the other hand, there are classical results in experimental psychology that suggest very strongly that understanding how a system works would make it "meaningful," and this would greatly improve a person's ability to learn the procedures and to remember them later.

There have been many attempts in experimental psychology to demonstrate just such a beneficial effect of how-it-works knowledge in the context of systems like text editors, but these attempts have almost uniformly failed to demonstrate the desired effects (e.g., Alexander, 1982; Foss, Smith, & Rosson, 1982). The work done under this project not only demonstrated these effects in the context of a simple control panel device, but also shed considerable light on the conditions under which these effects would appear, and provided a theoretically-based simulation model that explains the effects.

Approach. Subjects were asked to learn how to operate a simple control panel device that had a few switches, push buttons, and indicator lights. All of the studies involved comparing two conditions: In the rote condition, subjects learned how to operate the device strictly by rote, without any knowledge of the internal functioning or structure of the device, or without any explanation of the behavior of the device. In the model condition, before attempting to learn how to operate the device, the subjects learned a "mental model," in the form of a block diagram of the internal structure of the device, and learned how the internal components of the device were related to each other and to the controls. Performance in the learning task was then compared.

In the first study, subjects were explicitly trained to operate the device in several situations. These situations corresponded not only to normal operating conditions, but also to situations in which some internal component to the device was malfunctioning, and an alternative procedure had to be executed in order to compensate for the malfunction. This was intended to simulate the situations involved in working with real equipment.

After subjects had learned all the procedures, they were then given a series of retention tests. The major variables were the time taken to study the how-it-works explanation, the time taken to complete the training phase, the execution speed and accuracy during the retention tests, and certain qualitative features of retention, such as whether the recalled procedure was an improvement upon those that had been taught. That is, the instructed procedures had been designed so that some of them were quite inefficient. If the subject knew how to correctly interpret the indicator lights, the procedure could be considerably shortened.

In the second and third studies, the subjects inferred how to operate the device, rather than being explicitly trained to operate it. In this procedure, the subjects were given the device, told what the desired goal state was, and then were free to try to operate the controls to arrive at that goal state. Although operators of equipment are not normally put in this situation, this procedure proved to be a simple and effective way of determining and examining the effects of having the how-it-works knowledge.

Results. In the first experiment, in which the procedures were explicitly trained, the group who had studied how the system works learned the procedures faster, retained them better, and executed them faster, even after one week. A typical effect size was a 20% improvement. Of special interest is that the model group made the procedures more efficient far more often than the rote group. This means that not only was the how-it-works knowledge producing a general improvement in performance, but also a qualitative improvement in subjects' ability to deal intelligently with the device. The time taken to learn the mental model was roughly the same as the savings in training time, but notice that the model materials were not optimized. Thus, with no additional training time penalty, the model subjects were able to deal with the device much better.

A simple explanation for these results is that knowing how the system worked made it more "meaningful." However, this explanation is not detailed enough. A more precise hypothesis is that the how-it-works knowledge allowed the subjects to infer the procedures, which would give the subject two independent means of executing a procedure correctly. That is, the direct rote memory for the procedure failed, the subject could reconstruct the procedure based on inference from knowledge of how the device worked. In some cases, this inferred procedure could in fact be more efficient than the instructed one.

The second and third experiments confirmed this inference hypothesis. The subjects were asked to infer the procedures rather than learn them from explicit training. The second experiment simply compared a group trained on the model with a rote group in the procedure inference paradigm. The results were quite simple; the model group could infer the optimum procedures on the first try, whereas the rote group took several tries to arrive at the same procedures by sophisticated, but limited, trial and error approaches. Think-out-loud protocols showed quite clearly that the model group was basing their inferences on knowledge of how the system worked, whereas the rote group was basing their inferences on the superficial details of how the device looked and behaved, such as the fact that there were subtle mnemonic relationships between the labels on various controls and indicators. Another interesting result from the think-out-loud protocols is that the rote group subjects tended to view the device as unreliable and capricious in its behavior, and even suspected duplicity on the part of the experimenter. These affect-laden reactions suggest that much of what we think of as "computer anxiety" could in fact be due to the cognitive problem of not being able to explain or predict how a system is behaving.

The third experiment was intended to determine what aspects of knowledge of how the device worked was important. The previous experiments had supplied information about how the system worked in terms of a fantasy explanation based on the "Star Trek" television series. Namely, the control panel was described as being the control panel for a "phaser bank" aboard the "starship Enterprise". While this fantasy certainly motivated and interested subjects, there is an obvious concern about whether the effects produced were due to general properties of this fantasy. Also, the material included not only a description of the internal components of the system and their relations to each other, but also some discussion of the fictitious principles of physics involved. These principles may have provided some organizing structure, and thus might have produced the effects.

The third study was designed to demonstrate more clearly the nature of the critical information in the mental model, based on the idea that a good mental model supports inference of the procedures. The two factors compared were whether the how-it-works material had the fantasy and fictitious principle content or not, and whether the material provided the <u>system topology</u>, which is information about how the components and controls are connected to each other. Such materials included a block diagram and discussion about the actual controls on the control panel and how they were connected to the actual internal components of the system. The materials in the no-topology information with fantasy-principle content also presented a block diagram, but this diagram did not include any of the actual controls of the system; rather it corresponded to an idealized general description of how systems of this sort worked, rather than the specific system that the subject was dealing with.

The condition corresponding to no fantasy-principle content and no system topology information was essentially the same as the rote condition of the earlier studies. The condition with the fantasy-principle content and the topology information was the same as the previous model conditions. The critical comparison is whether the fantasy and principle content provides any performance facilitation, or whether the topology information is the critical content.

The results were very clear. The fantasy-principle content provided no facilitation at all; the critical information was the system topology information. The fantasy-principle condition with no topology information provided a general discussion of the principles of how systems of this type worked, but this did not allow subjects to infer the actual procedures needed to operate the device. Rather, it was critical to know how the controls related to each other and to the components.

Thus the general conclusion can be stated in terms of a criterion for when how-it-works knowledge will be of value to the user: How-it-works knowledge will be of value only if it is specific enough to allow the user to infer the exact operating procedures. Thus, the earlier attempts to demonstrate positive benefits of understanding how a word processor works probably did not provide knowledge that was specific enough. On the other hand, only some of the information about how a system works should be important; trying to understand technical detail that is not needed in order to be able to infer the operating procedures is simply a waste of time.

This leads to a second criterion for when how-it-works knowledge should be provided to the user of a piece of equipment: How-it-works knowledge should only be provided if it is actually necessary or advantageous to the user to be able to infer the procedures rather than learn them by rote. Notice that the ordinary telephone system is so easy to learn by rote that it is doubtful whether being able to infer how to operate it from knowledge of the switching mechanisms would be of any value. Similar arguments can be made for everyday systems like the automobile.

In the case of word processors, many of the commands that are involved in operating a word processor are either determined arbitrarily by the person who wrote the software, or are obvious to the user in terms of the text editing task itself. For example, no explanation is necessary for why pressing the up-arrow cursor key causes the cursor to move up. Likewise, no amount of explanation of the principles of computing or word processor design will explain why "EUN" is the keystroke sequence that will exit the editor; this was simply an arbitrary decision on the part of the designer.

Depending on the specifics of the design of the system, there may indeed be aspects of the how-it-works knowledge that is important for the user to know. However, this knowledge should be very specific, and severely limited in technical detail. It is a problem for future research to determine whether this knowledge has characteristics that would allow it to be determined on an a priori basis.

To explore these hypotheses theoretically, a simulation model was constructed for how procedures could be inferred from knowledge of the system topology. The simulation model has a declarative and a procedural component. The declarative component is a propositional representation of the block diagram, or topological description, of the system. The content of the block diagram, and the explanation that accompanied it, was basically the power flow connections through the system. That is, the diagram started at the main power source and went through various internal components and switches until it arrived at the final component of the system which required the power. The procedural component is a set of about 50 production rules which operate on the declarative representation. These production rules generate a plan for operating the device, and then execute the plan. If the plan fails to produce the desired result, the rules attempt to determine the problem in the system that caused the failure, and then attempt to devise a new plan.

The plan is constructed by finding a path through the block diagram of the system that routes power from the source to the desired point. The plan consists of a list of the control settings that will establish the route. In order to devise the plan, the production rules essentially simulated the internal state changes of the device. This corresponds to a popular notion of the role of mental models as allowing the person to simulate internally the states of the external world. Thus, the rules modelled the flow of power through the system, and made simple inferences about the conditions of individual components based upon the states of the indicator lights.

Although rather simple, this simulation model represents a potentially broad and important class of mental models, namely, any system in which some commodity, such as energy or information, is routed from one point to the next, through discrete and all-or-none components. Furthermore, models of this class have potential significance other than as simulation models for cognitive processing. They bear a strong resemblance to certain problems now being attacked in artificial intelligence, in which it is desired to troubleshoot or analyze the operation of a system based on a description of its internal structure.

The processing in the simulation model was compared in considerable detail to the response times of individual actions by subjects in the topology information conditions of third experiment. The basic question in this comparison was whether the simulation and human subjects performed their inference processes at the same points in the sequence of actions performed, given the cases where subjects performed the same sequence of actions as the simulation. If so, the relative amount of time required for the inferences should be accounted for by the model.

It was found that a reasonable portion of the variance in the response times could be accounted for by the simulation model, supporting the explanation for the role of how-it-works knowledge. The simulation was based on the principle of inferring procedures from a logically minimum required amount of knowledge of how the system worked, namely, its topology, along with a few simple and general principles for how power flows from one point to the next in such a system. A matter for further research is whether it is possible to formally characterize which portions of a device description are logically required for this form of inference.

Significance. This work bears very directly upon a basic issue in the preparation of training materials and of documentation for equipment. There has been some controversy for some time about the role of training in basic electronics theory in the training maintenance personnel. For example, Bond and Towne (1979) report that a common experience is that standard training in electronics theory is of little or no value in troubleshooting even complex equipment.

Traditional electronics training deals with very general principles, which although important, may only rarely explain the behavior of a piece of equipment at the level of analysis required for troubleshooting and repair. For example, much electronic repair of complex systems is done by identifying a defective module and replacing it. The logic of identifying the defective module usually involves reasoning based on tracing the power or signal flow through a set of interconnected modules. With complex systems, this reasoning may in fact be quite subtle, but it simply does not involve basic electronics theory, such as Ohm's law, or the details of transistor functioning. Rather, the logic of identifying a defective module is likely to be specific to the behavior of the modules and the topology of the system being repaired; only when one is troubleshooting at the individual discrete component level will the more basic electronics theory become important.

Continuation of this research, and an attempt to apply it more directly to training situations, potentially can result in much more efficient and effective training approaches. Also, potentially more effective equipment documentation could be prepared by ensuring that the critical how-it-works information required is prominent, and not obscured by unimportant information. It is not known at this time whether documentation in fact provides this critical system topology information in an easily used way; this would be a topic for further research.

Publications. The details of the experimental work is described in Technical Reports Nos. 13 (Kieras & Bovair, 1983) and 15 (Kieras, 1984c), and the simulation model and its comparison to data are described in detail in Technical Report No. 15 (Kieras, 1984c). A condensed presentation of the experimental results appeared in Kieras and Bovair (1984). A paper based on Kieras (1984c), concerning the simulation model and the comparison to data, has been submitted to Cognitive Science, and word on acceptance is expected within a few months.

This work has been described in several conference and colloquium talks. The work was presented at the 1983 Psychonomics Society meetings, in a colloquium series on applied cognitive psychology at the University of Michigan in December 1983, and in a colloquium at Bell Labs in February of 1984. Many reprint requests have also been received and responded to.

Prior Knowledge of Procedures in Rote Learning

Background. In recent years there has been developing a theory of the nature and acquisition of cognitive skill (Anderson, 1982). According to this theory, people have both declarative and procedural knowledge. Declarative knowledge is knowledge of facts, whereas procedural knowledge is knowledge of how to do things. Thus, in the context of operating equipment, the job of the learner is to acquire the knowledge of how to actually operate the equipment, which is procedural knowledge. Knowledge of how the equipment works, as discussed above, is declarative knowledge.

The theory goes beyond this simple distinction, however, to propose specific representations for both declarative and procedural knowledge. In line with established cognitive theory, declarative knowledge is represented as a semantic network. Procedural knowledge is represented as a set of production rules. A production rule is in the form:

IF (condition) THEN (action).

A production rule consists of a condition and an action; if the condition is satisfied, then the rule is "fired," and the action is performed. A set of production rules consists simply of a large set of such rules, with no built-in constraints upon the order in which the rules may be executed. Rather the order in which the rules fire is specified by the conditions and actions. The conditions can test for both external events, such as outside stimuli, or internal conditions such as the state of the semantic network representation or the contents of a working memory. The actions can both modify the external situation by means of overt responses, or can modify the state of memory. Thus, a piece of procedural knowledge consists of a set of production rules whose conditions and actions cause the rules to be fired in the correct order to produce the proper sequence of overt actions. Anderson's work has focussed on developing principles of learning that describe how an initial set of production rules can become more compact and efficient as learning proceeds. This theory has been able to explain many of the important and classical results in learning, such as the exact mathematical shape of the learning curve.

In the theory of cognitive skill, the initial set of production rules for a particular skill is assumed to be derived from declarative knowledge that is the original input to the system. That is, when first learning a skill, the learner would acquire a body of declarative knowledge that provides the specifications for the skill to be learned. These specifications would be interpreted by some general problem solving process, itself represented as production rules, and as a by-product of the activity of these general rules, specific rules for the particular skill will be formed.

Quite often in learning to operate equipment, the learner gets the initial specifications for the skill in the form of written step-by-step instructions. In terms of the theory, what the learner must do is to derive a correct set of production rules from the content of these instructions. Presumably, the written instructions would be comprehended by mechanisms similar to those already proposed in current theories of reading comprehension (e.g. Kieras, 1982b, 1983). These mechanisms would result in the learner having a declarative representation of the procedure available in memory immediately after reading the instructions. A general set of instruction-following processes, a pre-existing set of production rules, would then interpret this representation and carry out the correct procedure. Again, as a by-product of the activity of these general procedures, the specific production rules for the particular procedure would then be formed.

Anderson's work has focused almost completely upon the processes that occur once the correct production rules have already been formed. This work focused on the process by which the written instructions were translated into production rules. Thus, this work complements Anderson's, in that it focuses on the very initial stages of learning a skill from written material.

Approach. The rote condition of the first experiment on how-it-works knowledge described above yielded data on learning a series of procedures by rote. The time required to learn the individual procedures varied over a very large range, and appeared to be a function of the order in which the procedures were learned. After translating the procedures into production rules, it appeared that the excursions in training time could be accounted for by a simple transfer of training hypothesis. This hypothesis held that in learning a procedure, production rules that had been learned in a previous procedure could be transferred into the representation for the new procedure if the rule was either identical, or very similar to, the rule required for the new procedure. Thus, the time required to learn a procedure would be mostly a function of the number of new production rules required by the procedure. This simple rule could account for most of the variance in the rote learning data from the how-it-works knowledge experiment described above.

The work described in this section was then undertaken to provide a more comprehensive test of this simple transfer of training theory. The subjects learned procedures for operating a simple control panel device, which was the same one used in the studies of how-it-works knowledge described above. However, in this work, subjects were not provided any information about how the system worked. Rather, they learned how to operate the device strictly by rote. The training was done by explicitly listing the individual steps in the procedures, and having the subjects study these instructions, followed by attempting to reproduce the procedure from memory. This process was repeated until the subjects had learned the procedure.

A simulation model was constructed to rigorously simulate the transfer process. The production rule representations for each of the 10 procedures used were then put through the transfer simulation in various training orders. A set of three training orders was then chosen and used in the experiment. A set of step-by-step instructions was devised for each procedure that had the property that each individual sentence stating an instruction step corresponded very well to the contents of one production rule.

In the experiment, subjects first read through the step by step instructions, and then attempted to execute the procedure from memory. If they made a mistake, they were cycled back through the instructions. They repeated this alternation between reading the instructions and trying to execute the procedure until they successfully executed the procedure three times in a row. Then they went on to the next procedure in the specific training order. The variables of interest were the total time taken to learn a procedure, the time spent reading each individual instruction step, the accuracy of execution of each individual step in the procedure during training, and the speed and accuracy of retention in each step in the procedures in a final test for memory in the procedures.

The transfer simulation model was used to make rigorous <u>a priori</u> predictions of the number of new and transferred production rules in each procedure as a function of the training order. It was expected that these predictions could account for a substantial portion of the variance in training times, as was suggested by the preliminary analysis of the data discussed above. By examining the relationship between the reading times of individual instruction steps with the accuracy of execution of the corresponding step in the procedure, it should be possible to essentially track the acquisition of individual production rules. In this way, information could be obtained on the very initial stages of acquiring a procedure from written instructions.

Results. As expected, the analysis of the procedures in terms of the transfer hypothesis was able to account for a considerable proportion of the variance in training times. The number of new production rules required by a procedure was the single most important of the possible predictor variables considered, and alone could account for 69% of the variance in training times, and was a better predictor of training time on a single procedure than a subject's own mean training time.

A detailed regression analysis revealed that there were other effects involved in procedural learning as well. Of special interest is an apparent "overload" effect, in which a certain complex procedure was the first to be learned, and considerably more training time was required than would be predicted on the basis of the amount of new production rule information involved. Thus, while training time on the whole is very closely related to the number of new production rules, there are other aspects of training order which can be very important. These aspects can be clearly identified by applying the production rule analysis.

A matter for further research is exploring the nature of some of the additional effects, and testing the generality of the transfer theory, and whether the production rule analysis is as powerful in a variety of different task domains as it is with a simple control panel. Notice that Polson and Kieras (1985) have applied a similar analysis to the learning of a word processing system, and found similar predictive power of the production rule analysis. Thus, it appears that the production rule analysis is very general, but further research is needed to explore its limits.

A detailed analysis of the reading times for the individual instruction steps showed that people essentially cease to spend time on instruction steps once the corresponding steps in the procedure has been mastered. What is particularly surprising, however, is that this effect only appears for instruction steps that correspond to new production rules as defined by the transfer theory. Procedure steps that correspond to previously learned production rules do not show such a sharp decline in reading time. In other words, upon the very first reading of a procedure, subjects can distinguish between those instruction steps that correspond to production rules they already know, and those that correspond to rules that have to be learned.

instruction steps that are already known are read for very little time, right from the outset, whereas steps for new rules are read and studied until they are mastered, after which the reading time drops down to the same as that for steps already known.

What this pattern of reading time effects suggests is that the correct execution of a procedure depends on when a correct declarative representation of the procedure has been formed, and not upon when a production rule representation for it has been formed. That is, the pattern of reading time results appears to resemble what would be expected from powerful comprehension processes that can compare and manipulate declarative representations. It does not appear to be easily explained by the learning rules proposed by Anderson (1982) that are defined in terms of operations upon procedural knowledge. This means that the initial stages of learning from written text have more to do with comprehension processes than originally believed.

A preliminary analysis of the retention data suggests that the production rule analysis may also be very powerful in explaining the details of retention of procedures. Many of the errors made in recalling procedures could be accounted for by interference between two of the production rules in the procedures. Such rules had very similar conditions, differing in literally only one bit of information, but different actions, one being the correct action, and the other producing an incorrect action. About 95% of the errors in recall were due to the similar incorrect rule being fired instead of the correct one.

However, the severity of the interference of the incorrect rule was strongly related to classical variables from interference theory. For example, the amount of practice with the two rules, and their ordering during training, were very important. This shows that the traditional degree of learning and proactive versus retroactive interference considerations are at work in procedure retention. More importantly, the exact details of how these classical variables show up in recall can apparently be easily characterized in terms of the production rule analysis. A matter for future research is to clarify these effects further, and construct a rigorous model that predicts where these retention interference effects will occur.

Significance. This work is unique and unprecedented in that it is the first time that important quantitative features of the learning process could be accounted for with such power and precision by a completely a priori analysis. As shown by the Polson and Kieras work, this analysis also appears to be very general, but additional research is needed to confirm this claim. If this theory of learning and transfer can be successfully extended, it provides a very powerful analytic tool for investigating and improving the efficiency of training materials and training sequences.

As applied in a related project being conducted by Kieras and Polson under sponsorship of the IBM corporation, the analysis can be used to evaluate proposed designs for user interfaces of computer systems (see Kieras & Polson, in press; Polson & Kieras, 1985). That is, a high-quality user interface is one in which there is relatively little in the way of procedural knowledge that has to be learned in order to operate the system, and in which there will be strong positive transfer of production rules from one procedure to the next, corresponding to a "consistent" user interface. Thus, the practical significance of this work could be very large; further research will tell whether this potential is real.

On the purely theoretical front, this work has played an important role in clarifying the nature of procedural knowledge, and how it is acquired from written material. Together with other work being conducted under ONR sponsorship, such as Anderson's, there should soon be a comprehensive body of theory directly related to training issues at a level of precision and practical value that was simply not available before.

Publications. This work has been described in Technical Report No. 16 (Kieras & Bovair, 1985), and in a paper at the 1984 Cognitive Science Society Meetings. Related work from the Kieras and Polson project has been described in Kieras and Polson (in press), and in conference presentations by Polson and Kieras (1985). A journal article based on Technical Report No. 16 will be submitted within the next few months.

SUMMARY

The work in this project has made important contributions to the understanding of the role of prior knowledge in operating devices from written instructions. In terms of experimental methodology, some useful conclusions can be stated. First, readers have specialized strategies for dealing with unfamiliar material, which means that many traditional prose recall paradigms should be used with caution in the investigation of prior knowledge. The direction of this project had to be changed, because as originally proposed, it would have relied heavily on standard comprehension paradigms and thus was vulnerable to producing misleading conclusions about prior knowledge. Second, in experiments investigating mental models, or other forms of prior knowledge, careful attention should be paid to the relationship between the knowledge being supplied to the subject, and the exact tasks that the subject is expected to perform. Previous research in this area has not considered this relationship in enough detail, leading to many failed experiments and confusion over the role of prior knowledge, especially in the how-it-works domain.

On the theoretical front, this work has continued to demonstrate the power and effectiveness of rigorous theoretical analysis of the sort that can be represented in a simulation model. The effects of prior knowledge in prose memory situations, how-it-works knowledge utilization, and transfer of training, can all be explained by simulation models in a way that is theoretically precise, and in many cases quantitative and empirically powerful. On the whole, the results support what is perhaps becoming the consensus model of cognitive architecture, namely the ACT class of theories described in Anderson's most recent textbook, The Architecture of Cognition (1983).

The practical significance of these results is substantial, but will require further research to fully realize. With regard to instructional materials, both for immediate execution and long term learning, there are several important conclusions regarding the arrangement, sequence, and content of the material. If the conclusions from this research are confirmed by further research, it will be possible to make very precise decisions about what should be included in both training materials and operating instructions for equipment. Good choices could be made about what level of detail of how-it-works knowledge should be included in training materials and equipment documentation. The exact sequence and content of procedural instructions can be chosen with great precision. The overall arrangement and content of instructions for immediate execution can be chosen with regard to the expertise and knowledge of the users of the instructions. Thus, these results of this project clearly provide a good foundation for future applied research.

REFERENCES

Alexander, J. H. (1982). Computer text-editing: The development of a cognitive skill. Doctoral Dissertation, University of Colorado Department of Psychology.

Anderson, J. R. (1982). Acquisition of cognitive skill. Psychological Review, 89, 369-406.

Anderson, J. R. (1983). The architecture of cognition. Cambridge: Harvard University Press.

Bond, N. C. & Towne, D. M. (1979). Troubleshooting complex equipment in the military services: Research and prospects (Technical Report No. 92). University of Southern California, Behavioral Technology Laboratories.

Foss, D. J., Smith, P. L., & Rosson, M. B. (1982, November). The novice at the terminal: Variables affecting understanding and performance. Paper presented at the Psychonomic Society Meeting, Minneapolis.

Johnson, W., & Kieras, D. E. (1982). The role of prior knowledge in the comprehension of simple technical prose (Technical Report No. 11, UARZ/DP/TR-82/ONR-11). University of Arizona, Department of Psychology.

Johnson, W., & Kieras, D. (1983). Representation-saving effects of prior knowledge in memory for simple technical prose. Memory and Cognition, 11, 456-466.

Kieras, D. E. (1982a). What people know about electronic devices: A descriptive study (Technical Report No. 12 UARZ/DP/TR-82/ONR-12). University of Arizona, Department of Psychology.

Kieras, D. E. (1982b). A model of reader strategy for abstracting main ideas from simple technical prose. Text, 2, 47-82.

Kieras, D. E. (1983). A simulation model for the comprehension of technical prose. In G. H. Bower (Ed.), The Psychology of Learning and Motivation, 17. New York, NY: Academic Press.

Kieras, D. E. (1984a). A method for comparing a simulation model to reading time data. In D. Kieras & M. Just (Eds.), New methods in reading comprehension research, Hillsdale, N. J.: Erlbaum.

Kieras, D. E. (1984b). The psychology of technical devices and technical discourse. In Artificial Intelligence in Maintenance: Proceedings of the joint services workshop, Air Force Human Resources Laboratory, Brooks Air Force Base, Texas 78235, Report No. AFHRL-TR-84-25.

Kieras, D. E. (1984c). 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. (in preparation). Theoretical mechanisms of task and strategy effects in comprehension. To appear in B. Britton (Ed.) Executive control processes in reading.

Kieras, D. E. & Bovair, S. (1983). The role of a mental model in learning to operate a device (Technical Report No. 13 UARZ/DP/TR-83/ONR-13). University of Arizona, Department of Psychology.

Kieras, D. E. & Bovair, S. (1984). The role of a mental model in learning to operate a device. Cognitive Science, 8 191-219.

Kieras, D. E., Tibbitts, M., & Bovair, S. (1984). How experts and non-experts operate electronic equipment from instructions. (Technical Report No. 14, UARZ/DP/TR-84/ONR-14). University of Arizona, Department of Psychology.

Kieras, D., & Bovair, S. (1985). The Acquisition of Procedures from Text: A Production-System Analysis of Transfer of Training (Technical Report No. 16, TR-85/ONR-16). University of Michigan, Department of Humanities.

Kieras, D. E., & Polson, P. G. (In press). An approach to the formal analysis of user complexity. International Journal of Man-Machine Studies.

Polson, P. G., & Kieras, D. E. (1985, April). A quantitative model of the learning and performance of text editing knowledge. Presented at the CHI 85 conference on Computer-Human Interaction, San Francisco.

Rumelhart, D. E. (1980). Schemata: The building blocks of cognition. In R. Spiro, B. Bruce, & W. Brewer (Eds.), Theoretical issues in reading comprehension. Hillsdale, N.J.: Erlbaum.

Smith, E. E. & Goodman, L. (1982). Understanding instructions: The role of explanatory material (Technical Report No. 5088). Bolt Beranek and Newman, Inc.

APPENDIX

Reports, Publications, Presentations
Contract No. N00014-84-K-0731, NR 667-473

Technical Reports

Johnson, W., & Kieras, D. E. (1982). The role of prior knowledge in the comprehension of simple technical prose (Technical Report No. 11, UARZ/DP/TR-82/ONR-11). University of Arizona, Department of Psychology.

Kieras, D. E. (1982). What people know about electronic devices: A descriptive study (Technical Report No. 12 UARZ/DP/TR-82/ONR-12). University of Arizona, Department of Psychology.

Kieras, D. E. & Bovair, S. (1983). The role of a mental model in learning to operate a device (Technical Report No. 13 UARZ/DP/TR-83/ONR-13). University of Arizona, Department of Psychology.

Kieras, D. E., Tibbitts, M., & Bovair, S. (1984). How experts and non-experts operate electronic equipment from instructions. (Technical Report No. 14, UARZ/DP/TR-84/ONR-14). University of Arizona, Department of Psychology.

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., & Bovair, S. (1985). The Acquisition of Procedures from Text: A Production-System Analysis of Transfer of Training (Technical Report No. 16, TR-85/ONR-16). University of Michigan, Department of Humanities.

Archival Publications

Johnson, W., & Kieras, D. (1983). Representation-saving effects of prior knowledge in memory for simple technical prose. Memory and Cognition, 11, 456-466.

Kieras, D. E. (1983). A simulation model for the comprehension of technical prose. In G. H. Bower (Ed.), The Psychology of Learning and Motivation, 17. New York, NY: Academic Press.

Kieras, D. E. (1984). A method for comparing a simulation model to reading time data. In D. Kieras & M. Just (Eds.), New methods in reading comprehension research, Hillsdale, N. J.: Erlbaum.

Kieras, D. E. & Bovair, S. (1984). The role of a mental model in learning to operate a device. Cognitive Science, 8 191-219.

Kieras, D. E. (1984). The psychology of technical devices and technical discourse. In Artificial Intelligence in Maintenance: Proceedings of the joint services workshop, Air Force Human Resources Laboratory, Brooks Air Force Base, Texas 78235, Report No. AFHRL-TR-84-25.

Kieras, D. E. (in press). The how, when, and why of doing a cognitive simulation. Behavior Research Methods, Instrumentation, and Computers.

Kieras, D. E. (submitted). A simulation model of a mental model: Inferring how to operate a device from knowledge of how it works. Submitted to Cognitive Science.

Kieras, D. E., & Bovair, S. (in preparation). The acquisition of procedures from text: A production system analysis of transfer of training. To be submitted.

Kieras, D. E. (in preparation). Theoretical mechanisms of task and strategy effects in comprehension. To appear in B. Britton (Ed.) Executive control processes in reading.

Presentations

Kieras, D. E. (1981, June). Prior knowledge and device operation. Presented in Meeting on instructions organized by Henry M. Halff, Navy Personnel Research and Development Center, San Diego.

Kieras, D. E., & Johnson, W. (1981, November). Prior knowledge in the comprehension of simple technical prose. Presented at the Psychonomic Society Meetings, Philadephia.

Kieras, D. E., & Johnson, W. (1982, April). Prior knowledge and comprehension. In H. Ellis & L. Bourne (Co-chairs), Workshop in Human Learning, Memory, and Cognition, at the Rocky Mountain Psychological Association Meetings, Albuquerque.

Kieras, D. (1982, June). Technical prose and technical knowledge. Colloquium briefing presented at the Naval Personnel Research and Development Center, San Diego.

Kieras, D. (1983, October). The psychology of technical devices and technical discourse. Invited presentation at the Joint Services Workshop on Artificial Intelligence in Maintenance, Automatic Testing, Maintenance Aiding and Training. Boulder, Colorado.

Kieras, D., & Bovair, S. (1983, November). Mental models and learning how to operate a device. Presented at the Psychonomic Society Meetings, San Diego.

Kieras, D. E. (1983, December). Mental models and interacting with devices. Colloquium address presented in the Colloquium Series on Applied Cognitive Science, University of Michigan, Ann Arbor.

Kieras, D. E. (1984, March). The role of mental models in operating devices. Invited Colloquium presented at Bell Laboratories, Murray Hill.

Kieras, D. E., & Bovair, S. (1984, June). The acquisition of procedures from text. Presented at the Sixth Annual Conference of the Cognitive Science Society, Boulder, Colorado.

Kieras, D. E. (1984, November). The how, when, and why of doing a cognitive simulation. In A. Lesgold (Chair), Tutorial on Cognitive Simulation, held at the Meetings of the Society for Computers in Psychology, San Antonio.