

DOCUMENT RESUME

ED 257 036

CS 007 955

AUTHOR Kieras, David E.; Bovair, Susan.
TITLE The Acquisition of Procedures from Text; A Production-System Analysis of Transfer of Training. Technical Report No. 16.
INSTITUTION Michigan Univ., Ann Arbor.
SPONS AGENCY Office of Naval Research, Arlington, Va. Personnel and Training Research Programs Office.
REPORT NO TR-85/ONR-16
PUB DATE 29 Jan 85
CONTRACT N00014-84-K-0731
NOTE 33p.
PUB TYPE Reports - Research/Technical (143)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS Cognitive Measurement; *Cognitive Processes; *Learning Processes; *Learning Strategies; Measurement Techniques; *Reading Comprehension; *Transfer of Training
IDENTIFIERS *Text Learning

ABSTRACT

In the context of the currently developing theory of procedural knowledge and cognitive skill, the task of learning procedures from written instructions must involve the formation of production rules from the information available in the text. Since this process has not been explored in great depth, a study examined the relation of the production rule representation to transfer of training. Subjects learned procedures for operating a simple piece of equipment by reading step-by-step instructions. By measuring the reading time on individual steps and the accuracy of execution of the procedure, the acquisition of individual production rules was tracked. Since the procedures are related, some transfer of training was possible from procedures learned earlier. Results show that production rules, as a way to represent procedural knowledge, can provide a detailed account of the learning process, and that the transfer process is more similar to comprehension processes than to conventional practice mechanisms or Anderson's (1982) learning principles. (DF)

 * Reproductions supplied by EDRS are the best that can be made *
 * from the original document. *

X This document has been reproduced as
received from the person or organization
originating it.
Minor changes have been made to improve
reproduction quality.

• Points of view or opinions stated in this docu-
ment do not necessarily represent official NIE
position of policy.

ED 257 036

THE ACQUISITION OF PROCEDURES FROM TEXT: A PRODUCTION-SYSTEM ANALYSIS OF TRANSFER OF TRAINING

David E. Kieras & Susan Bovair

University of Michigan

Technical Report No. 16 (TR-85/ONR-16)

January 29, 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.

5007975



UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER TR 85/ONR-16	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) The Acquisition of Procedures from Text: A Production-system Analysis of Transfer of Training		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) David E. Kieras & Susan Bovair		8. CONTRACT OR GRANT NUMBER(s) N00014-84-K-0731
9. PERFORMING ORGANIZATION NAME AND ADDRESS University of Michigan Ann Arbor, MI 48109		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 29 Jan., 1985
		13. NUMBER OF PAGES 37
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) Learning, Transfer, Procedural knowledge,		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The current theory of cognitive skill describes knowledge of procedures in terms of a production rule representation which is constructed on the basis of an initial declarative (propositional) representation. In these terms, learning a procedure from written instructions consists of converting the propositional content of the written material into production rules. This process was studied in a transfer of training,		

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

experiment. Subjects learned from step-by-step instructions a series of related procedures for operating a simple device, with the major manipulation being the order of learning the procedures. Very strong transfer effects were obtained, which could be predicted very well by a simple model of transfer. Individual production rules can be transferred, or re-used in the representation of a procedure if they appeared in a previously learned procedure, meaning that learning time is mostly a function of the number of completely new production rules that must be acquired. Examination of the time required to read individual instruction steps suggests, however, that this transfer mechanism involves processes acting on declarative propositional representations of the production rules. This means that the transfer process is more similar to comprehension processes rather than conventional practice mechanisms, or Anderson's (1982) learning principles.

BEST COPY AVAILABLE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

ABSTRACT

The current theory of cognitive skill describes knowledge of procedures in terms of a production rule representation which is constructed on the basis of an initial declarative (propositional) representation. In these terms, learning a procedure from written instructions consists of converting the propositional content of the written material into production rules. This process was studied in a transfer of training experiment. Subjects learned from step-by-step instructions a series of related procedures for operating a simple device, with the major manipulation being the order of learning the procedures. Very strong transfer effects were obtained, which could be predicted very well by a simple model of transfer. Individual production rules can be transferred, or re-used in the representation of a procedure if they appeared in a previously learned procedure, meaning that learning time is mostly a function of the number of completely new production rules that must be acquired. Examination of the time required to read individual instruction steps, suggests, however, that this transfer mechanism involves processes acting on declarative propositional representations of the production rules. This means that the transfer process is more similar to comprehension processes rather than conventional practice mechanisms, or Anderson's (1982) learning principles.

BEST COPY AVAILABLE

1

The Acquisition of Procedures from Text:
A Production-System Analysis of Transfer of Training
David E. Kieras & Susan Bovair

Quite often people must learn procedures from written instructions. In the context of the currently developing theory of procedural knowledge and cognitive skill (Anderson, 1982), this task must involve the formation of production rules from the information available in text. This process has not been systematically explored; the results reported here provide an initial characterization. Two general conclusions follow from this work. The first is that a production rule representation can provide a very precise characterization of the relative difficulty of learning a set of related procedures. The second is that apparently there are powerful comprehension-like processes that operate very early in learning on declarative representations of production rules. This supplements Anderson's (1982) description of the acquisition of skill, in that many of the important processes involved in learning a procedure can take place before a procedural representation has been formed.

The approach was to have subjects learn procedures for operating a simple piece of equipment by reading step-by-step instructions. By measuring the reading time on individual steps, and the accuracy of execution of the procedure, it is possible to track the acquisition of individual production rules. Since the procedures are related, some transfer of training is possible from procedures learned earlier. The key result is that this transfer is predicted very well from the similarities between the production system representations for the procedures.

DESCRIPTION OF THE TASK

The subjects learned series of procedures for how to operate a device, consisting of a simple control panel. The goal of operating the device was to get a certain indicator light to flash. Note that this was a rote learning situation; the internal organization of the device was not taught to the subjects. Each procedure consisted of several steps, as illustrated in the step-by-step instructions in Tables 1 and 2. Table 1 is the procedure for a "normal" situation, in which the device is operating properly. Table 2 is the procedure for a "malfunction" situation, in which some internal component of the device was not operating. Depending on the nature of the malfunction, the device either could be made to work by an alternate procedure, or could not. The final step in each procedure was to signal success or failure in getting the device to work.

Table 1
Example of a Normal Procedure

If the command is to do the MA procedure, then do the following:

Step 1. Turn the SP switch to ON.

Step 2. Set the ES selector to MA.

Step 3. Press the FM button, and then release it.

Step 4. If the PF indicator flashes,
then notice that the operation is successful.

Step 5. When the PF indicator stops flashing, set the ES selector
to N.

Step 6. Turn the SP switch to OFF.

Step 7. If the operation was successful,
then type "S" for success.

Step 8. Procedure is finished.

BEST COPY AVAILABLE

Table 2
Example of a Malfunction Procedure

If the command is to do the MA procedure, then do the following:

- Step 1. Turn the SP switch to ON.
- Step 2. Set the ES selector to MA.
- Step 3. Press the FM button, and then release it.
- Step 4. If the PF indicator does not flash,
then notice that there is a malfunction.
- Step 5. If the EB indicator is on, and the MA indicator is off,
then notice that the malfunction might be compensated
for.
- Step 6. Set the ES selector to SA.
- Step 7. Press the FS button, and then release it.
- Step 8. If the PF indicator does not flash,
then notice that the malfunction can not be compensated
for.
- Step 9. Set the ES selector to N.
- Step 10. Turn the SP switch to OFF.
- Step 11. If the malfunction could not be compensated for,
then type "N" for not compensated.
- Step 12. Procedure is finished.

BEST COPY AVAILABLE



The Control-Panel Device

The device used in this experiment was the same as that used in Kieras & Bovair (1984), in which the major manipulation was whether subjects were taught a mental model for the internal organization and structure of the device. In this experiment, subjects only learned the device by rote. The mental model is included here only to explain the behavior of the device, and the rationale for the choice of procedures to be taught to the subjects.

The device is a slope-front box with a simple front panel, shown in Figure 1, consisting of four controls, and four indicator lights. A laboratory computer detects the positions of the controls and turns the indicator lights on and off. The four controls consist of a toggle switch (SP), a three-position selector (ESS), and two push-buttons (FM and FS). The four indicator lights are labeled SPI, EBI, MAI, and PFI. The labels are based on the mental model used in Kieras and Bovair (1984).

The behavior of the device can, most easily be described in terms of the diagram shown in Figure 2, which was used in the mental model experiments. Power flows through the device from left to right, controlled by the switches; and also affected by the state of the imaginary internal components, shown as boxes in the diagram. If a component is not functioning correctly, then power cannot flow through it and the device malfunctions. There are four components, EB, MA, SA, and PB.

The SP switch is the on/off switch device and the SP indicator is the pilot light. The other indicators are status lights for the associated components. Thus, if the power switch is on, the EB indicator will be on if the EB component is good, and off if the EB component is bad. Note that there is no indicator associated with the SA component.

Power flows to the PB component when the ESS selector is set to either MA or SA, and the corresponding button is pushed (FM for ESS-MA, FS for ESS-SA). When the PB component receives power, the PF indicator flashes four times and then stops until the button is released and pushed again. Whether these combinations of control settings will work depends on the status of the components in the obvious way. For example, if the MA is bad, then the MA indicator will not be lit, and the ESS-MA, FM combination will not flash the PF indicator. If the EB or the PB is bad, then the PF indicator cannot be flashed because power cannot reach it, no matter how the controls are set.




Figure 1. The control panel device.




Figure 2. The simulated internal structure of the device, not known to subjects.

Although there are sixteen possible states that the device can be in, only six of these can be distinguished in terms of the behavior of the device. For example, if the EB component is bad, then the status of the other components is irrelevant because the pattern of indicator lights will be the same in all cases. These distinguishable states are shown in Table 3, which shows for each state of the device what indicator lights will come on, and which settings will make the PF indicator flash. The states are labeled by the defective component label prefixed with X; for example, XMA means that the MA component is defective. While knowledge of the internal structure of the device, and component status for each malfunction, makes it easy to understand the behavior of the device, it is important to note that this information is not provided to the subjects in the experiment.

Table 3
Possible States of Control Panel Device.

Label	Sub-device status	Status of EBI, MAI	PFI flash on MA procedure	PFI flash on SA procedure
NORMAL	All good	Both on	yes	yes
XEB	EB bad, others any	Both off	no	no
XMA	MA only bad, others good	EBI on, MAI off	no	yes
XSA	SA only bad, others good	Both on	yes	no
XEB-XMA	PB bad, EB and MA good, SA any	Both on	no	no
XEB-XSA	SA bad, EB good, PB or SA or both bad	EBI on, MAI off	no	no

Operating Procedures

There are a total of twelve procedures that could be used to operate the device in its six possible states, six where the MA setting of the ESS is tried first, and six in which the SA is first. In the experiment, the subject was commanded to do either the MA procedure or the SA procedure, where these commands referred to which ESS setting was to be tried first. Of the twelve possible procedures, the PF indicator can be made to flash in six (NORMAL, XMA, and XSA). It was decided that the definition of a malfunction should be that the first settings tried would not work, and so SA-XMA and MA-XSA were not included, leaving a total of ten procedures; two normal and eight malfunction procedures. The procedure steps are listed in Table 4. Tables 1 and 2 give examples of the step-by-step instructions for a normal and a malfunction procedure.

Table 4
Procedures Used to Operate Control Panel Device.

MA procedures					
MA-NORMAL	MA-XEB	MA-XPB	MA-XMA	MA-XMA-XSA	
(1) SP on	(1) SP on	(1) SP on	(1) SP on	(1) SP on	(1) SP on
(2) ESS-MA	(2) ESS-MA	(2) ESS-MA	(2) ESS-MA	(2) ESS-MA	(2) ESS-MA
(3) FM push	(3) FM push	(3) FM push	(3) FM push	(3) FM push	(3) FM push
(4) ESS-N	(4) ESS-N	(4) ESS-N	(4) ESS-SA	(4) ESS-SA	(4) ESS-SA
(5) SP off	(5) SP off	(5) SP off	(5) FS push	(5) FS push	(5) FS push
(6) Tap "S"	(6) Tap "N"	(6) Tap "N"	(6) ESS-N	(6) ESS-N	(6) ESS-N
			(7) SP off	(7) SP off	(7) SP off
			(8) Tap "S"	(8) Tap "N"	(8) Tap "N"
SA procedures					
SA-NORMAL	SA-XEB	SA-XMA-XSA	SA-XSA	SA-XPB	
(1) SP on	(1) SP on	(1) SP on	(1) SP on	(1) SP on	(1) SP on
(2) ESS-SA	(2) ESS-SA	(2) ESS-SA	(2) ESS-SA	(2) ESS-SA	(2) ESS-SA
(3) FS push	(3) FS push	(3) FS push	(3) FS push	(3) FS push	(3) FS push
(4) ESS-N	(4) ESS-N	(4) ESS-N	(4) ESS-MA	(4) ESS-MA	(4) ESS-MA
(5) SP off	(5) SP off	(5) SP off	(5) FM push	(5) FM push	(5) FM push
(6) Tap "S"	(6) Tap "N"	(6) Tap "N"	(6) ESS-N	(6) ESS-N	(6) ESS-N
			(7) SP off	(7) SP off	(7) SP off
			(8) Tap "S"	(8) Tap "N"	(8) Tap "N"

The eight malfunction procedures can be divided into two types. The first is those in which the first ESS setting tried does not work, and the alternate setting might work, depending on the malfunction state. These were termed possibly compensatable malfunctions. In the second type, the alternate setting will not work, and so need not be tried. These were termed non-compensatable malfunctions. For example, the XEB state is a non-compensatable malfunction for either the MA or the SA command, and the MA-XMA and SA-XSA states are possibly compensatable malfunctions. This distinction was presented to the subjects as part of the overall instructions, to rationalize the details of the procedures.

THEORETICAL ANALYSIS

Transfer Effects

In earlier work with this device (see Kieras and Bovair, 1983) it was noticed that the time required to learn the procedures under rote conditions varied over a very wide range. Training times for the rote-learning subjects are shown in Figure 5, which shows the training time for each procedure in the order that they were learned. Note that the order of procedures was fixed, rather than randomized, as would traditionally be done. Note that rather than being a smooth descending learning curve, there are large peaks for the times of the third, fifth, and ninth procedures. The number of steps in each procedure does not explain this pattern, because while the number of steps does vary for different procedures, the difference is not very large, and is frequently in the wrong direction. For example, procedure 5 has 8 steps, which is more than procedures 3 and 4 with 7 each, but procedure 6 has 9 steps. Rather, the pattern could be explained by the fact that the first procedure contains all new information, the second (the other normal procedure) contains only a little new information, the third (the first malfunction procedure) contains some new information, the fourth (the second malfunction procedure) very little, and the fifth (the first possibly compensatable malfunction) quite a lot.

BEST COPY AVAILABLE




Figure 3. Predicted and observed training times for the rote learning condition from Kieras and Bovair (1983).

These data then suggest that the amount of new information in a procedure is a plausible candidate for a predictor of training time, but it is not well-defined. Transforming the instructions into production rules could provide a precise characterization of what is to be learned, thus it could be determined which rules could be transferred, resulting in a quantitative measure of the amount of new information, namely, the number of new production rules that must be learned.

Production Rule Representation

Table 5 provides an example production rule set for the procedure in Table 1. The syntax of these rules is very simple. Each rule is in the form:

(Label IF (condition) THEN (action)).

The production system's working memory contains the GOALS to be accomplished, and NOTES, which consist of non-goal items concerning current processes, the environment, or specifications of the tasks to be accomplished. See Kieras and Polson (in press) for a full description of the production system notation, along with a description of the user-device interaction simulation that was used to test the production rules for accuracy.

A set of production rules was written for each procedure used in the Figure 3 experiment, and tested in the user/device interaction simulation to check for accuracy and completeness. Writing the production rules was done using a computer text editor, and it became obvious that once the first set of rules was generated, then subsequent sets could be generated easily by copying the first set, doing a few substitutions, and adding a few rules when necessary. By analogy, the transfer process could consist of recognizing which new rules are identical to previously learned rules, which are extremely similar to existing rules, and which are totally new. The subject could then spend most of the training time acquiring the new rules, and merely "tagging" existing rules as applying to the new situation.

Table 5
Example of Production Rules

```
(MA-N-START
IF (AND (TEST-GOAL DO MA PROCEDURE)
        (NOT (TEST-GOAL DO ??? STEP)))
THEN ((ADD-GOAL DO SP-ON STEP)) )
```

```
(MA-N-SP-ON
IF (AND (TEST-GOAL DO MA PROCEDURE)
        (TEST-GOAL DO SP-ON STEP))
THEN ((OPERATE-CONTROL *SP ON)
      (WAIT-FOR-DEVICE)
      (DELETE-GOAL DO SP-ON STEP)
      (ADD-GOAL DO ES-SELECT STEP)) )
```

```
(MA-N-ES-SELECT
IF (AND (TEST-GOAL DO MA PROCEDURE)
        (TEST-GOAL DO ES-SELECT STEP))
THEN ((OPERATE-CONTROL *ESS MA)
      (WAIT-FOR-DEVICE)
      (DELETE-GOAL DO ES-SELECT STEP)
      (ADD-GOAL DO FM-PUSH STEP)) )
```

```
(MA-N-FM-PUSH
IF (AND (TEST-GOAL DO MA PROCEDURE)
        (TEST-GOAL DO FM-PUSH STEP))
THEN ((OPERATE-CONTROL *FM PUSH)
      (WAIT-FOR-DEVICE)
      (OPERATE-CONTROL *FM RELEASED)
      (DELETE-GOAL DO FM-PUSH STEP)
      (ADD-GOAL DO PFI-CHECK STEP)) )
```

```
(MA-N-PFI-CHECK
IF (AND (TEST-GOAL DO MA PROCEDURE)
        (TEST-GOAL DO PFI-CHECK STEP)
        (LOOK *PFI FLASHING))
THEN ((ADD-NOTE OPERATION SUCCESSFUL)
      (DELETE-GOAL DO PFI-CHECK STEP)
      (ADD-GOAL DO ES-N STEP)) )
```

BEST COPY AVAILABLE

Table 5 (Continued)

```

(MA-N-ES-N
IF (AND (TEST-GOAL DO MA PROCEDURE)
        (TEST-GOAL DO ES-N STEP)
        (LOOK *PFI OFF))
THEN ((OPERATE-CONTROL *ESS N)
      (WAIT-FOR-DEVICE)
      (DELETE-GOAL DO ES-N STEP)
      (ADD-GOAL DO SP-OFF STEP)) )

```

```

(MA-N-SP-OFF
IF (AND (TEST-GOAL DO MA PROCEDURE)
        (TEST-GOAL DO SP-OFF STEP))
THEN ((OPERATE-CONTROL *SP OFF)
      (WAIT-FOR-DEVICE)
      (DELETE-GOAL DO SP-OFF STEP)
      (ADD-GOAL DO TAP STEP)) )

```

```

(MA-N-TAP
IF (AND (TEST-GOAL DO MA PROCEDURE)
        (TEST-GOAL DO TAP STEP)
        (TEST-NOTE OPERATION SUCCESSFUL))
THEN ((DELETE-NOTE OPERATION SUCCESSFUL)
      (ADD-NOTE TYPE S-FOR SUCCESS)
      (DELETE-GOAL DO TAP STEP)
      (ADD-GOAL DO FINISH STEP)) )

```

```

(MA-N-FINISHED
IF (AND (TEST-GOAL DO MA PROCEDURE)
        (TEST-GOAL DO FINISH STEP)
        (TEST-NOTE TYPE S-FOR SUCCESS))
THEN ((DELETE-NOTE TYPE S-FOR SUCCESS)
      (DELETE-GOAL DO FINISH STEP)
      (DELETE-GOAL DO MA PROCEDURE)) )

```

BEST COPY AVAILABLE

Two basic transfer rules were defined: identity (from copying), and generalization (a form of substitution). Production rules can be considered identical if they have the same conditions and the same actions. The original definition of the generalization transfer rule was: if rules have the same actions, and only one point of difference in their conditions, then the rules could be generalized. This was done by replacing the differing point with a "wild-card" that matches any value. Thus, if the only point of difference between two rules was that one had the condition clause (TEST-GOAL DO MA PROCEDURE), and the other had (TEST-GOAL DO SA PROCEDURE), then this clause could be replaced by (TEST-GOAL DO ??? PROCEDURE), where "???" is a wild-card that will match any item in that position. This version of the generalization transfer rule was later modified as described below.

When these transfer rules were applied to the production rules for the procedure training order shown in Figure 3, the number of new rules that needed to be added for each procedure was determined. The assumption is that the only rules that require substantial effort to learn are the completely new ones; the identical and generalizable rules should be very easy to learn, since all or almost all of their content is already known. Thus, the number of new rules in a procedure should be closely related to the difficulty of learning the procedure. In these data, the number of new rules in a procedure accounts for 79% of the variance among the mean training times for the 10 procedures, supporting the value of the production system analysis of transfer in the learning of procedures. However, this result was based on only ten data points, and so is no more than suggestive.

EXPERIMENT

By using three different training orders, this study was designed to get a more comprehensive set of data on the relation of the production rule representation to transfer of training. The three different training orders were chosen by analyzing the production rule sets for each procedure using a transfer process simulation program, described below, and selecting training orders that produced substantial variation in the number of new rules in each procedure, and also the number of new rules in each serial position in the training order.

Overview

Each subject learned a series of 10 procedures in a fixed order. There were three different orders, chosen as described below, with a separate group of subjects for each order.

BEST COPY AVAILABLE

To learn each procedure, the subjects first read a set of step-by-step instructions for the procedure, such as those in Tables 1 and 2, and then attempted to execute the procedure on the device. If they made an error, they were immediately informed, and then began to read the instructions again. They were required to execute the procedure correctly three times in a row before proceeding to the next procedure. The data recorded were the reading time on each step of the instructions, the accuracy of each step while executing the procedure, and the speed and accuracy of a final retention test, which will not be discussed in this report.

Method

Transfer Simulation. A simulation program was written in LISP that could perform the transfer processes automatically. The program (COMBINE-HARVEST) can be given a set of production rules for a procedure which it then examines for possible transfer with the set of rules already known. It generates a new set of rules, and also reports on the number of rules considered identical to existing rules, the number that could be generalized with existing rules, and the number of new rules added. The output from COMBINE-HARVEST was tested in the user-device interaction simulation to check that a proper rule set was generated, that is, it followed the correct procedure in any given situation.

The program's generalization criteria were modified slightly from the original definition; some of these modifications were more restrictive than the original, while some extended the definition. The new generalization rule for transfer specifically excluded certain types of clauses from the generalization process. These are: clauses that sequence the firing of rules (e.g., goals of the form DO STEP X), clauses that look for a particular configuration of indicator lights on the device, and clauses that operate controls on the device. These changes in the generalization criteria mean that, in practice, only operations on notes, and goals can be generalized. The new generalization rule was extended so that it could generalize more than one clause in the condition, and could also generalize the equivalent clauses in the action part of the production rule.

The program this would generate the number of new rules required for each procedure. Different orders of procedures produced very different patterns of the number of new rules, both in terms of the serial position in the sequence, and when comparing particular procedures in different orders. Depending on the order in which it was processed, different training times would be predicted for the same procedure.

Training Order Conditions. Three different training orders were selected for the experiment that would maximize the predicted training time differences. Either a procedure would have different predicted training times because there were a different number of new rules to be acquired, or if the number of new rules was the same, then the procedure would be in a different position in terms of the order of learning. It was not possible to maximize these differences for all the procedures, but it was done for as many as possible. These different orders also produced different numbers of rules accepted by the transfer process as identical or generalized, because it seemed likely that recognizing identity could be a faster process than generalization. The training orders were therefore chosen to produce different values for the number of identical rules, and the number of generalized rules. A final constraint on the training orders were that they should be, in some sense, meaningful orders.

The selected training orders are shown in Table 6. In training order condition 1, the order is: all the MA procedures first, then the SA procedures. Within this division, normal procedures are first, not-compensatable malfunctions second, and possibly compensatable malfunctions last. Training order condition 2 is based on the idea that once the longest procedures are learned, the shorter procedures should be learned comparatively easily. Thus, the order for training order condition 2 is: possibly compensatable malfunctions first, non-compensatable malfunctions second, and normal procedures last, and within these groups MA procedures are presented before SA. The order of training order condition 3 is on the principle of underlying causes, even though subjects have no information on these causes. Thus the order is normal procedures first, then XEB malfunctions, then XPB malfunctions, then the XMA and the XSA pair, and finally the two XMA-XSA malfunctions. Within these pairings SA procedures came before MA procedures. This order is quite different from the other two in that orders 1 and 2 are based on the procedures themselves, actions carried out by the subject, while training, order condition 3 is based on the behavior of the device.

Instruction Materials. A set of step-by-step instructions were prepared for each procedure; examples appear in Tables 1 and 2. These were prepared so that each sentence in the instructions appeared to correspond to a single production rule, one for each step or action (internal or external) involved in the procedure, and care was taken that these steps corresponded to the production rule sets themselves, as illustrated by the correspondence between Tables 1 and 5 for the corresponding steps in the different procedures.

BEST COPY AVAILABLE

Table 6

Number of New Production Rules
for Each Training Order Condition

Training Order Conditions								
		1			2			3
Serial Position	Name	New rules added	Name	New rules added	Name	New rules added		
1.	MA-NORMAL	9	MA-XMA	13	SA-NORMAL	9		
2.	MA-XEB	5	MA-XMA-XSA	4	MA-NORMAL	2		
3.	MA-XPB	1	SA-XSA	5	SA-XEB	5		
4.	MA-XMA-XSA	4	SA-XPB	0	MA-XEB	0		
5.	MA-XMA	2	MA-XEB	1	SA-XPB	4		
6.	SA-NORMAL	2	MA-XPB	1	MA-XPB	1		
7.	SA-XEB	0	SA-XEB	0	MA-XMA	5		
8.	SA-XMA-XSA	1	SA-XMA-XSA	1	SA-XSA	0		
9.	SA-XSA	3	MA-NORMAL	2	SA-XMA-XSA	1		
10.	SA-XPB	0	SA-NORMAL	0	MA-XMA-XSA	0		

BEST COPY AVAILABLE

Apparatus. The device consisted of an actual physical control panel connected to a laboratory computer, which monitored the settings of the switches and push buttons and controlled the indicator lights accordingly. All instructions and commands to the subjects were presented on a standard video terminal positioned next to the device. A computer-assisted instruction facility was used to present all of the procedure training and the retention tests. The subject was seated in a small room at a table with the terminal and the control panel, and was observed by means of a video camera and monitor.

Subjects. Subjects were recruited through campus advertisements and were paid \$5 for their participation. Subjects were randomly assigned to each of the three training order conditions. A total of 70 subjects participated in the experiment. Ten subjects' data was discarded, leaving a total of 60 subjects, with 20 subjects in each condition. Of the 10 subjects whose data was discarded, two final subjects were discarded because their data was not needed; three subjects did not finish the training part of the experiment, one subject was discarded because of a fire alarm during the experiment, and the first four subjects were not used because the experimental software required changes.

Design. Training order condition was a between-subjects factor, with each subject randomly assigned to one of the three training order conditions, subject to the constraint that approximately equal numbers were maintained in the three conditions. Each subject learned all 10 procedures in all three conditions. Subjects were also assigned by gender, so that there would be an equal number of males and females in each condition.

Instructions and Procedure. The first part of the instructions familiarized the subjects with the layout and labels on the device. Subjects were then told that they would be trained in several procedures for operating the device. They were told that the goal of operating the device was to make the PF indicator flash. Part of their training would include procedures to be performed if the device malfunctioned. They were told that for some malfunctions the PF indicator would not flash at first, but it might be possible to change the control settings so that it would flash. This was called compensating for a malfunction, and it was pointed out that some malfunctions could not be compensated for. The subjects were instructed that whenever they were asked to turn the device to the initial state that they should set the SP switch off, the ESS selector to N, and not push any buttons.

The training procedure consisted of alternating reading and trying phases. In the reading phase, the subject read the procedure a single step at a time, in a self-paced reading paradigm. Then in the trying phase, the subject attempted to

execute the procedure correctly. After the attempt, the subject would return to the reading phase. This process was repeated until the subject had completed three correct attempts in a row. Then the subject would commence learning the next procedure.

In the reading phase, the subject would tap the space bar to read each step on the terminal screen, which appeared as one sentence, as illustrated in Tables 1 and 2. The previous step was erased from the screen. Subjects were instructed to study each step for as long as they felt necessary. The lab computer recorded how long the subject left each step on the screen, defined as the reading time. When the subject had read all the steps in the procedure, a command, such as "Do the MA procedure," would appear on the screen and the subject would then try to perform the procedure from memory. If the subject made a mistake while attempting the procedure, the lab computer immediately sounded a buzzer, as a signal to stop trying. Then the subject was returned to the beginning of the reading phase. If the subject performed all steps correctly, the computer sounded a bell tone, and either returned to the beginning of the reading phase or went on to the next procedure if the criterion had been achieved. Throughout the procedure, the subject was prompted by displays on the terminal screen, such as a message that they had made an error and were being returned to the reading phase.

Since some pilot subjects tended to ignore the indicators during training, the instructions included a notice that although it might seem unnecessary to pay attention to the indicator lights during training, during the testing phase at the end of the experiment, it would be necessary to rely on the pattern of indicator lights to choose the correct procedure.

After being trained to criterion in all 10 procedures subjects were instructed that they could take a short rest or break before starting the test. They were told that they would see each of the 10 procedures three times each in the test in a random order. No feedback was given during testing.

RESULTS

Training Time

The total training time for a procedure is defined as starting when a subject begins the first reading of the first sentence of the instruction steps, until completing the last step of the last attempted execution.

The first analysis was simply to verify the presence of gross effects of the the training order on training time. An analysis of variance was performed on the total training time for each procedure in each training order condition; the means are

shown in Table 7. There were main effects of training order condition and procedure, and an interaction between training order condition and procedure ($p < .05$). While female subjects were an average of ten seconds faster than males on the training, this difference was not significant, and there are no significant interactions of gender.

Multiple regression analyses were performed in order to test the predictions from the theoretical analysis. The dependent variable was the total training time (TRTIME) giving 600 data points, one for each subject on each procedure in each condition. The predictor variables were those derived from the production rule model, the number of new productions (NEW), the number of generalized production rules (GEN), and the number of identical or old production rules (OLD). Other predictor variables included the subject's mean training time for all procedures (SMEAN) to handle the within-subject design (see Pedhazur, 1982), the main effect of training order (ORDER), and two dummy variables (COND1 and COND2) to test for a main effect of condition, with condition 3 as the baseline. Since the first procedure trained usually required a disproportionately long time, a dummy variable, FIRST, indicated whether the procedure was the first to be trained. Two interaction variables, C1FIRST and C2FIRST, were defined to represent the interaction of condition and first procedure trained.

The results of the regression analyses on the training time are shown in Table 8. The table shows the coefficients in the final equation that includes all variables that entered the stepwise analysis. The F-ratios are the "F-to-remove", and so provide a test of significance of the coefficients in the final equation. Finally, the standardized regression coefficients allow comparisons of the importance of each variable independently of scale differences. About 76% of the variance in total training time was accounted for by the final equation.

Figure 4 shows the predicted and observed mean times and the final regression equation. The most important predictor variable was the number of new rules in each procedure (NEW), which alone could account for 69% of the variance, and uniquely accounts for about 47% of the variance. The partial and standardized regression coefficients for NEW are substantially larger than those for identical (OLD) rules and generalizable rules (GEN), which are very similar.

BEST COPY AVAILABLE

Table 7
 Mean Training Times (secs) for Each Procedure
 for the Three Training Order Conditions

Procedure	Training Order		
	Condition 1	Condition 2	Condition 3
MA-NORMAL	212.496	81.125	111.883
SA-NORMAL	89.863	92.814	221.958
MA-XEB	142.058	111.907	98.829
MA-XPB	109.430	96.679	108.727
MA-XMA	117.012	464.089	165.727
MA-XMA,XSA	161.478	190.291	139.679
SA-XEB	79.677	84.980	160.697
SA-XMA,XSA	86.568	99.644	95.250
SA-XSA	111.109	176.169	151.411
SA-XPB	117.109	136.013	191.817
Mean	122.727	153.371	144.598

Table 8
 Regression Analysis
 on Total Training Time ($N = 600$, $R^2 = .7623$)

Variable	Final Coeff.	Final Std.Coeff.	F
CONSTANT	-132.39		
MEAN	1.00	.410	389.78
NEW	19.38	.662	153.54
OLD	11.82	.499	88.44
GEN	11.07	.291	51.09
2RFRST	165.10	.324	125.04
RFRST	47.10	.155	16.04
ORDER	-3.93	-.124	18.32
COND2	-16.51	-.085	14.86




Figure 4. Mean predicted and mean observed training times for each training order condition.

In addition, there were other effects, notably some learning-to-learn effects (FIRST and ORDER), and an apparent "overload" effect, C2FIRST, in which the first procedure in the second training order condition took an extremely long amount of time to learn beyond that predicted by the number of new production rules. This procedure was MA-XMA, shown in Table 4, which involved trying the MA setting first, then the SA setting. The other two orders involved only relatively simple normal operation steps, which may have appeared obvious and natural. Thus, subjects in condition 2 were confronted with a first procedure in which the first few steps apparently have no effects. This sort of conceptual difficulty is clearly a matter for further research.

Despite these other effects, however, the production system variables provided by the transfer model explain the training times very well; in fact, the number of new rules alone accounts for 69% of the variance, and is a better predictor of training time on a single procedure than the subjects' individual means! Thus, by analyzing the procedures in terms of production rules, and the relations between them, it is possible to account for the difficulty of the learning the procedures with great precision.

Reading Time

The time required to read each sentence of the instructions was averaged over procedures, but classified by training trial (e.g., first reading, second reading, and so forth), and by the transfer status of the corresponding production rules (Old, Generalizable, New). Figure 5 shows these means. The key point is simple. There was a substantial difference in the reading times for instruction steps depending on the transfer status of the corresponding production rule. The reading times for generalizable and old rules were almost identical, but reading times for new rules were much longer for the first few readings. A key result is that this difference appears on the first reading, meaning that subjects can immediately distinguish whether a sentence describing a step corresponds to a new rule or to a known one, and can immediately govern their reading and study time accordingly. The difference between reading times on the first trial between New and Generalized is strongly significant ($z = 3.51, p < .01$).




Figure 5. Mean reading times for instruction sentences as a function of reading trial and the transfer status of the corresponding rule.

BEST COPY AVAILABLE

A second question about the reading time is how they relate to the acquisition of individual production rules. Figure 6 shows the reading times for individual sentences plotted in terms of relative trial to mastery. The trial of mastery of a sentence was defined as the reading trial after which the subject executed the corresponding step in the procedure correctly for all trials thereafter. The figure shows the mean reading times for sentences classified by whether the corresponding production rule was new, generalized, or identical.

These data were subjected to a fairly complex regression analysis, summarized in Table 9, in order to determine the significance of the apparent effects. In terms of nuisance variables, the reading time depends on the subject's mean and the number of WORDS in the sentence, and there is a simple main effect of relative trial number (RELTRL), corresponding to the overall downward trend. There is an apparent practice effect, because sentences whose steps are mastered later, as shown by larger values of MASTRL, are read for less time. The key results are: NEW sentences are read longer than Identical or Generalizable, which are almost the same; sentences before mastery, BEFMAS, are read longer than after; and the effect is mostly due to the New sentences, as shown by the interaction variable BMNEW. Thus, consistent with Figure 5, sentences that state new rules are studied until the corresponding rules are mastered, whereupon they are studied for much less time.

Table 9
Regression Analysis on
Individual Sentence Reading Times
(N = 21,449, R² = .40)

Variable	Final Coeff.	Final Std. Coeff.	F
CONSTANT	-.557		
WORDS	.069	.188	1223.60
SMEAN	.844	.316	3508.06
MASTRL	-.267	-.153	655.46
NEW	.723	.188	798.56
RELTRL	-.182	-.132	224.06
BEFMAS	.739	.210	599.37
BMNEW	1.247	.208	831.06

BEST COPY AVAILABLE




Figure 6. Mean reading times for instruction sentences as a function of relative trial of mastery and transfer status.

BEST COPY AVAILABLE

CONCLUSIONS

A basic conclusion is that production rules, as a way to represent procedural knowledge, can provide a detailed account of important learning processes. This supports the approach suggested by Kieras and Polson (in press) who suggest that the production-rule theory of skill acquisition is useful for practical applications. That there are other phenomena involved, such as the "overload" described above, is clarified by the production system analysis as well.

These results present a puzzle for the theory of skill acquisition as formulated by Anderson (1982). The transfer process defined here has many similarities to some of Anderson's compilation and tuning processes. However, his processes are defined in terms of operations on procedural representations. These are constructed as a by-product of the activity of general interpretive procedures that are driven by an initial declarative encoding. However, in these results, rules are being compared, modified, and constructed very rapidly, and apparently before they exist in a procedural form. As Figure 5 shows, a generalization process can apparently occur on the first reading, and is almost as fast as recognizing an identical rule. Although there is no rigorous basis at this time for saying so, it seems that these aspects of the results are not reasonably subsumed under Anderson's compilation and tuning processes.

Instead, perhaps the work of relating new and old rules is done by processes similar to those proposed for macroprocessing in comprehension (e.g., Kieras, 1982), which can compare, modify, and construct complex propositional representations while reading is going on. Thus, subjects translate the instruction sentence into a declarative representation of a complete production rule, which can then be related to other such representations. If an identical representation already exists, it can be re-used, as appears to be possible in other types of text (see Johnson & Kieras, 1983). Similar to Anderson's proposals, this declarative representation would be interpreted by a general procedure for following instructions, and the procedural form of the rules would eventually be formed by the compilation and tuning processes. However, correct execution of the procedure would begin when the declarative rule set has been successfully constructed, and the time required to do so would depend on how much use could be made of previously learned rule representations. Thus, when procedures are acquired from text, comprehension-like processes can play a major role early in learning, leaving the compilation and tuning processes to govern learning once the initial declarative form of the rules is in place.

REFERENCES

- Anderson, J. R. (1982). Acquisition of cognitive skill. Psychological Review, 89, 369-406.
- 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. (1982). A model of reader strategy for abstracting main ideas from simple technical prose. Text, 2, 47-82.
- 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.
- Kieras, D. E., & Polson, P. G. (In press). An approach to the formal analysis of user complexity. International Journal of Man-Machine Studies.
- Pedhazur, E. J. (1982). Multiple regression in behavioral research (2nd ed.). New York: Holt, Rinehart, & Winston.