Journal of Experimental Psychology: Applied
1995, Vol. 1, No. 3, 227–244

1076-898X/95/$3.00

Following Instructions: Effects of Principles and Examples

Richard Catrambone
Georgia Institute of Technology

Two experiments examined whether people's success using general instructions for operating a device (a word processor) would be improved if the instructions were supplemented with principles or examples. The principle explained relevant internal workings of the word processor. Its inclusion aided initial performance and later transfer, presumably by helping the learner to explain a difficult step. The example was predicted to reduce ambiguities and therefore, start-up time; its presence aided initial performance but appeared to have no effect on transfer. Besides demonstrating that general instructions can be improved through examples and principles, the studies suggest that a learner can access features of the initial instruction-based representation when faced with a novel task, even after executing a procedure a number of times.

People often have difficulty following instructions for devices such as word processors. One reason for this difficulty is that users are unsure how to apply the instructions to the particular case on which they are working. Instructions that are tailored for specific cases help users apply the instructions initially, but users later can have difficulty transferring those instructions to novel cases. Conversely, users learning from instructions written at a more general level are able to apply the instructions to a wide variety of tasks, but they have great difficulty using the instructions initially (Catrambone, 1990).

This trade-off in initial performance and later transfer as a function of type of instructions are troubling. It would be desirable to find a way to help learners use instructions successfully and quickly, as well as transfer them easily to novel tasks. Although some success has been found in revising the specific instructions to aid generalization (Catrambone, 1990), the approach taken here is to examine whether the initial use of general instructions can be improved while still maintaining good transfer to novel tasks.

Research that examines factors influencing how well people can use instructions has important practical implications for instructional design. Users often give up on product instructions that are hard to follow or understand initially and instead try the tasks on their own (Mack, Lewis, & Carroll, 1983). Not surprisingly, this can lead to considerable trouble operating the device or software and may ultimately affect future decisions to use products from the same company.

General Versus Specific Instructions

Catrambone (1990) found a trade-off between specific and general instructions for people learning deletion tasks on a word processor. Specific instructions are defined as instructions that provide the exact steps for a given procedure such as deleting a word. General instructions are defined as those that provide a procedure covering a wide number of cases within a task domain, such as how to delete *any* text. Participants using general deletion instructions took roughly twice as long as specific-instruction participants to perform the initial tasks. However, specific-instruction participants were less successful at transferring to novel tasks. General-instruction participants had little

This research was supported by the Office of Naval Research Grant N00014-91-J-1137. I thank Susan Bovair, Dan Fisk, Alex Kirlik, Franz Schmalhofer, and Neff Walker for comments on the manuscript of this article.

Correspondence concerning this article should be addressed to Richard Catrambone, School of Psychology, Georgia Institute of Technology, Atlanta, Georgia 30332. Electronic mail may be sent to rc7@prism.gatech.edu.

difficulty with novel tasks, possibly because they had been led to form a general procedure.

Two sets of production rules were written to represent the procedural knowledge presumably conveyed by the general and specific instructions. It was assumed that individual production rules are learned at the same rate, as long as they are matched for number of clauses. Thus, given that there were fewer productions needed to represent the general instructions, it was expected that it would take no longer to learn the general instructions than it would take to learn the specific instructions. Learning time was assessed by how long it took participants to accomplish an initial task. However, although the production rules formed from the general instructions were sufficient to accomplish the initial task, it appeared that participants did not learn these rules very easily when studying the instructions. Rather, they required multiple readings and attempts at the task. Conversely, participants seemed to quickly learn the productions conveyed by the specific instructions.

Why do learners have difficulty implementing general instructions? One possibility is that general instructions provide relatively little information to constrain potential user actions. Thus, the search space for appropriate actions can be large. The present study examines two ways of helping learners to implement general instructions, thereby improving their initial usability: the use of principles and examples. These two ways are examined because of prior work in the mental model and problem-solving literatures that demonstrates the benefits of this type of information for people learning procedures and learning how to solve problems. Although the potential benefits of adding principles and examples to general instructions are discussed in detail below, it is worth highlighting the major issues first.

With respect to examples, the problem-solving literature has shown that learners can readily map the steps from an example to a new problem that is very similar to the example. However, they frequently have great difficulty applying the example to novel cases. If an example is added to general instructions, it is unclear whether the strengths of general instructions and examples will be combined or whether the weaknesses will dominate. That is, will the example help learners perform initial tasks quickly while the general instructions support transfer to novel cases, or will performance reflect start-up difficulties associated with general instructions and transfer difficulties associated with examples?

With respect to principles, prior work in the mental model literature has suggested that a device model can help a learner infer steps in new situations (e.g., Kieras & Bovair, 1984). Thus, a principle can help a learner make appropriate decisions with respect to implementing general instructions for an initial task. However, given the good transfer produced by general instructions, it is unclear whether a principle would additionally impact transfer to later novel tasks.

Effects of Mental Models on Learning and Executing Procedures

There is evidence that people can learn instructions for operating a device more successfully if they have some understanding of how the device behaves internally or if they understand the goal that a set of steps achieves. For example, Kieras and Bovair (1984) demonstrated that learners who studied a mental model for a device (a *phaser bank*) learned procedures for operating the device more quickly than learners who did not previously learn the mental model. In addition, those who first learned the mental model were also better able to optimize procedures. Kieras and Bovair argued that the type of extra information that helps people learn procedures more effectively is information that helps them infer the steps of procedures.

Smith and Goodman (1984) demonstrated the usefulness of an explanatory schema in instructions. They compared a group of participants who followed a set of steps for assembling an electric circuit to a group who received a structurally oriented explanatory schema with the steps. This schema consisted of statements that provided a rationale or goal for carrying out sets of steps (e.g., "The next thing that you will have to do is to assemble the on–off switch."). Thus, the schema included the underlying model that constrained the steps needed to build the circuit. When assembling a new circuit, participants who had previously received the explanatory schema were more accurate, even though the required steps were not identical to the ones followed during training. The schema participants were presumably aided by the fact that the same underlying model was present for both circuits, and this model helped them comprehend the new steps.

Learning From Examples

A good deal of research has shown that learners prefer to learn from examples (e.g., Chi, Bassok, Lewis, Reimann, & Glaser, 1989; LeFevre & Dixon, 1986; Pirolli & Anderson, 1985). One reason examples are often preferred may be that they provide an instantiation of a procedure to guide behavior. However, one well-established difficulty is that learners often have trouble generalizing examples to novel problems (Catrambone & Holyoak, 1990; Cooper & Sweller, 1987; Reed, Dempster, & Ettinger, 1985; Ross, 1989). Thus, although examples aid performance on problems that are almost identical to the examples, they frequently provide little benefit and sometimes cause interference for novel problems.

Two Ways of Improving General Instructions

The major goal of the present study was to examine whether the initial usability of general instructions can be improved through the use of principles and examples. A secondary goal was to examine whether principles and examples affected performance on later similar tasks as well as novel tasks.

Use of a Principle

Smith and Goodman (1984) found that learners carried out steps more quickly when they were provided with information that gave a rationale for the steps. Other studies indicate that background knowledge or elaborations may help initial performance (Hale, 1993; Reder, Charney, & Morgan, 1986).

Kieras and Bovair (1984) suggested that information about how a device operates will aid user performance to the extent that it helps the learner infer the steps needed to operate the device. Thus, information about a device that does not support this type of inferencing will presumably not aid performance. This functional definition of *useful information* may at first seem circular, but in practice a researcher or instruction writer can show a priori how the information might help a learner infer or comprehend certain steps (Kieras, 1990). Other work has suggested that information that supports and guides a user's exploration of a device can aid long-term performance (e.g., Carroll, Mack, Lewis, Grischkowsky, & Robertson, 1985).

Use of Examples

The problem-solving literature has frequently demonstrated the difficulty learners have when using examples on new tasks that are not isomorphic to the examples (e.g., Reed et al., 1985). An issue investigated in Experiment 1 is whether an example that is not isomorphic to the initial task can aid performance if the example is used in conjunction with general instructions. The example could provide the learner with concrete advice on how to implement certain steps while the general instructions provide sufficient structure to enable the learner to determine how to adapt steps from the example to the current task. Conversely, it is possible that a nonisomorphic example will not aid initial performance or may even hinder it compared with a situation in which an example is not present.

With respect to transfer to later tasks, it is unclear whether the presence of an example will hurt later transfer despite the fact that general instructions alone appear to support transfer (Catrambone, 1990). Given that learners have been shown to ignore instructions in favor of examples (LeFevre & Dixon, 1986), it is possible that an example could be more salient than the general instructions and thus hinder transfer performance.

Overview of the Experiments Plus a Performance Assumption

The experiments examined participants as they learned to do deletion tasks (Experiment 1) and format tasks (Experiment 2) on a word processor. In both experiments, the primary issues explored were whether the presence of an example or principle would aid initial performance and whether their presence would affect later performance on similar and new tasks.

When a learner successfully executes a procedure described in a set of instructions, the resulting procedure representation may be more a function of having carried out the procedure and less a function of the details of how it was presented in the instructions. That is, it is assumed that the memory that results from carrying out the task is more salient than the memory for the instructions that were read. This assumption is consistent with the finding of Ross and Kennedy (1990). They discovered that once learners had used a prior case to solve a new problem, their subsequent problem-solving performance seemed to be guided by a representation formed from the problem-solving experience rather than details of the prior cases. It is also consistent with the work of Luchins (1942), who found that when learners discovered a solution approach that worked for a particular problem type, they would continue to use that approach for other problems even when a simpler or more efficient approach was possible.

The assumption of the primacy or salience of the *execution* memory predicts that once a learner has successfully completed a task, later occurrences of the same task are less likely to be influenced by the presence of a principle or example, even if initial performance was affected by them. The likelihood of execution memory playing a strong role in subsequent performance presumably will increase with the number of prior executions of the procedure on similar tasks. However, if learners later face a novel task for which the old procedure will not work, they might access the memory of a principle or example that could influence their performance.

One limit of the above assumption is that if a learner has great difficulty carrying out an initial task and makes many mistakes and false turns, the resulting procedure representation may be so poor that the learner has to refer to the instructions again in order to be able to carry out subsequent executions of the task (e.g., Mack et al., 1983). In this case, the nature of the instructions may continue to exert an influence on performance until the procedure is better learned.

Experiment 1

Experiment 1 used deletion tasks on a word processor as the domain for investigating the effects of principles and examples on initial performance and later transfer. For the word processor used in this experiment, which was chosen because it was expected to be unfamiliar to potential participants, deletion is accomplished by (a) putting the cursor at the beginning of the to-be-deleted text, (b) pressing a function key, (c) selecting the *Delete* command from a menu of block operations, (d) highlighting the to-be-deleted text, and (e) pressing *Enter*. When the deletion function is selected, the computer displays the prompt, "Delete what?" and waits for a target character to be specified. When the user specifies the target, the computer then highlights the text up to that character. The target character can be pressed several times if it occurs more than once in the to-be-deleted text.

Although the concepts of highlighting and doing string searches are familiar to experienced users, they are often bewildering to inexperienced users. Consider the general instructions for deletion in Table 1. Step 7 in the table instructs the user to type the end character "over and over until the text is highlighted" in order to highlight the to-be-deleted text. This step could be difficult for an inexperienced user to understand because it might seem odd to have to type a character while trying to delete something. The user might also be confused as to why the target character would need to be pressed more than once. Thus, when faced with the first deletion task of the experiment such as deleting a word that requires a single-press of the target character, a learner could have a fair amount of difficulty.

One way of reducing confusion is to provide a principle that explains what the computer does each time the target is pressed, thus giving the learner a rationale for pressing the key and for pressing it multiple times. This principle is given in Step 7 under the heading "Additional information for principle participants." This principle could help the learner comprehend Step 7 more easily and carry out the initial task more quickly than a learner who was not presented with the principle.

As mentioned earlier, once a learner has successfully executed the general procedure on the initial task, the resulting representation of the procedure in memory is assumed to be sufficient to allow this learner to carry out isomorphic tasks successfully. In this case, features of the learning materials, such as a statement of a principle, are predicted to no longer influence performance. Additional trials with the same task will presumably produce a speed-up in performance, but this should be true regardless of whether the learner was given the

Table 1
Deletion Instructions Used in Experiment 1: Deleting Text

Step 1:

DELETING TEXT

Use the arrow keys to move the cursor so that it is under the first character of the text you want to delete.

Step 2:
Press the F4 key.

Step 3:
Notice that the following "menu" appears at the bottom of the computer screen:
Move Copy Delete Recall Save.
Also notice that Move is highlighted.

Step 4:
Press the right-arrow key twice to highlight:
Delete.

Step 5:
Press the ENTER key.

Step 6:
Notice that at the bottom of the screen the computer displays the question:
Delete what?

Step 7:
Type the character at the end of the text you want to delete, typing it over and over until the text is highlighted.

Additional information for principle participants that followed the previous text:	Additional information for example participants that followed the previous text:
Each time you type the character, the computer "searches" in a forward direction starting from the point at which the cursor is located, until the computer finds the character. When the computer finds the character, it highlights all the text it searched through on the way to finding the character.	Matches initial task: For example, if the word you wished to delete was airplane then you would type the letter e. Does not match initial task: For example, if the word you wished to delete was telephone then you would type the letter e three times.

Step 8:
If you highlighted the wrong text and you want to start over, press the ESC key. After you press the ESC key, you will have to move the cursor back to the beginning of the text you wish to delete and then press the F4 key in order to restart the deletion procedure.

Step 9:
If you highlighted the correct text, then press the ENTER key.

Step 10:
Notice that the text that was highlighted disappears. END

Note. ESC = escape.

principle or an example because it is assumed that the strengthening process that produces speed-up (e.g., Anderson, 1983) will be operating on a representation that will be similar across groups.

When a learner faces later novel tasks, the representation formed from the training tasks may no longer work because it may have been tailored to the specific training tasks. Therefore, the principle might become relevant again as the learner tries to problem solve. For instance, when learners are faced with a novel task that requires them to press the target character multiple times, those learners who saw the principle explaining what the computer does when the target character is pressed might be quicker to determine what they need to do.

Another way of helping a learner perform initial tasks is with the use of an example. Once again, consider Step 7 in Table 1. Under the heading "Additional information for example participants" are examples of how the step for specifying the target might be implemented. The example labeled "Matches initial task" is expected to allow the learner to understand more rapidly that the last letter of the to-be-deleted word is in fact the letter that should be pressed. In addition, the example is isomorphic to the initial deletion task attempted by participants in Experiment 1; in both cases the target character needs to be pressed only once in order to highlight the to-be-deleted word.

The example in Table 1 labeled "Does not match initial task" shows a case in which the target character needs to be pressed multiple times in order to highlight the word. If learners have difficulty adapting the example to a nonisomorphic case, then this example will not help performance as much as the single-press example and may not even help performance relative to a no-example situation. On the other hand, if the general instructions help learners adapt the example, then performance on the initial task might be better than the no-example case and perhaps as good as the single-press example case.

Two scenarios are considered likely with respect to the effect of the example manipulation on performance on later novel tasks involving multiple presses of the target. One possibility is that participants who had seen the multiple-press example will be quicker on the first multiple-press task because it will match the example. A second possibility is that the combination of the general instructions and practice on the earlier tasks will mitigate any effects of an early example, perhaps by making the example less likely to be accessed during later tasks.

If the effect of the examples and principle is confined entirely to the initial task, and performance on later novel tasks is good for all groups, this would suggest that they are effective in aiding comprehension of general instructions. Presumably this would occur if learners, by applying the general instuctions to the initial task, form a sufficiently general representation for performing a variety of tasks. Alternatively, if there is also an effect of principle or example on later tasks, this would suggest that the procedure learned by doing the initial task did not support inferencing so well that other factors could not improve performance.

In summary, it was hypothesized that a principle would aid performance on the initial task and possibly on the first novel task. It was also hypothesized that an example matching the initial task would aid performance relative to the no-example condition and possibly the mismatching example condition. It was hypothesized that the mismatching (multiple-press) example might also aid performance on the first novel task because that task requires multiple presses of the target character.

An additional issue examined in Experiment 1 is whether any benefit in initial performance from a principle or example in the instructions is offset by time needed to read this additional information. The value of a principle or example on initial performance time would be lowered if learners required an offsetting amount of time to process this supplemental information. However, it is possible that the hypothesized improved comprehension from a principle or example might prevent a significant increase in reading time.

Method

Participants. Participants were 73 (41 men and 32 women) college-aged students at the Georgia Institute of Technology who received course credit. Participants' computer experience as indicated on a questionnaire was confined to using word processing software on a Macintosh. The interface for the Macintosh software is considerably different from the interface for the word processor used in this experiment.

Materials and procedure. Participants performed word processing tasks on an IBM PS/2 Model 80 computer. This will be referred to as the *task computer.* Participants were first shown several features of the task computer and the word processing program. The computer screen displayed an empty document. Participants were shown what the cursor was and were told how text automatically wraps around the screen. The *Shift* key and the arrow keys on the keyboard were pointed out, and participants were shown how to move the cursor around the screen with the arrow keys. Participants were told that their keystrokes would be collected and time-stamped by the computer. Participants were then asked to type a paragraph to allow them to get familiar with the keyboard.

After typing the practice paragraph, participants were shown how to access the instructions for doing various procedures. Participants read instructions from the computer screen of a second IBM PS/2 Model 80. This will be referred to as the *instruction computer.* The procedures were retrieving a document, deleting text, inserting text, and exiting a document. Each procedure (e.g., retrieving a document) was broken into a series of screens and labeled Screen 1, Screen 2, and so on. The last screen for a particular procedure always concluded with the word *End.* Table 1 presents the instructions for deleting text. Note that the headings for each screen were Screen 1, Screen 2, etc., not Step 1, Step 2, etc.

The deletion instructions were identical for all participants except for Step 7 (see Table 1). The principle seen by half of the participants provided an explanation of how the computer searches for target characters. For participants who received an example, it demonstrated either the deletion of a word requiring a single keypress, thus matching the first task, or the deletion of a word requiring multiple keypresses. For participants who received the principle and an example, the principle came first.

The instructions could be viewed only one screen at a time. The contents of a screen became visible when the participant held down the space bar. When the space bar was not depressed, an outline of the instructions for all the topics appeared on the display. The outline consisted of rows of dashes where each row corresponded to a screen. The row representing the first screen of information for a particular procedure (such as retrieving a document) consisted of the title of the procedure rather than dashes. This allowed participants to keep visual track of where they were in the instructions. In addition, one row in the outline was always at a higher intensity than the others. This row corresponded to the screen that would appear if the space bar was pressed.

Participants could go forward or backward through the instruction screens by pressing the Next Page key or the Previous Page key. Participants' movements through the instructions were saved to the computer's memory. After being shown how to read instructions, participants were presented the first document on which they were to work. The various changes that had to be made were indicated on the document. Text to be deleted was underlined in red ink. The name of the document was printed in the upper left-hand corner because the name was needed in order to retrieve the document. Subsequent documents were similarly annotated.

Training phase. During the training phase, participants deleted six words as well as doing insertion tasks and document retrieval and exiting. The three training documents required the following tasks in order: (a) retrieve the document, (b) delete a word, (c) insert a phrase, (d) delete a word, (e) insert a phrase, and (f) exit the document.

The phrase insertions were always seven words long. The word deletions during the training phase always required a single keypress of the target character (i.e., the last letter never occurred earlier in the word). The insertion and document retrieval and exiting tasks were included to make the training phase somewhat realistic.

Prior to doing a task (such as retrieving a document) for the first time, participants were required to read the complete instructions for that task. This was done in order to make sure participants saw all the steps for the procedure at least once and would be less tempted to guess about how to do a step later on. Participants were told that once they were done reading the instructions for a particular procedure, they could not look back at it while they attempted to do the task.

Once participants began a task, if they did not know what to do at a particular point or made a mistake from which they could not recover, they had to reread the instructions for that procedure and then redo the task. Once participants successfully completed a task, they did not have to read the instructions for that procedure again unless they later made a mistake from which they could not recover.

The time to do a deletion task was defined as the moment the function key was pressed until the *Enter* key was pressed, causing the appropriate text to disappear from the screen. This performance time included time spent redoing the task (and possibly rereading the instructions) if the participant made a mistake. Frequently, participants would reread the instructions after successfully completing a task but before beginning the next task. This reading time was included in the appropriate reading time cells in Table 3 but was not counted as part of the time to do a deletion task. For instance, if a participant reread the instructions after completing the first deletion task but before beginning the second deletion task, this reading time would be part of the category "Before second success" in Table 3 but would not be counted in the time to perform the second deletion task because the rereading occurred before the task was begun.

Test phase. During the test phase, participants performed a total of eight deletion tasks across four different documents. The first document required deleting two words where the final letter also occurred in the middle of the word, thus requiring the target letter to be pressed multiple times ("mysterious" and "presence"). The second document required deleting garbage text at the beginning of two words ("[swk]various" and "[hox]twill"). The third and fourth documents required deleting garbage text in the middle of two words in which the final letter of the garbage text also occurred earlier in the garbage text, thus requiring multiple presses of the target letter ("str[mpotp]ange," "dis[bfafuf]tress," "poss[znopntn]ible," and "cond[qwniw]uct").

Design. The between-subjects independent variables were presence of a principle and presence and type of example (no-example, single-press, and multiple-press). These variables were crossed forming six experimental groups. The number of participants in each group is indicated in Table 2. The within-subject variables were serial position and training versus test phase. The dependent variables were time to perform each deletion task and time spent reading the instructions. Participants performed a total of 14 deletion tasks.

Table 2
Time to Perform Deletion Tasks in Experiment 1 (in Seconds)

	Example					
	Single press[a]		Multiple press[b]		None	
Task	Principle ($n = 12$)	No principle ($n = 12$)	Principle ($n = 11$)	No principle ($n = 13$)	Principle ($n = 13$)	No principle ($n = 12$)
			Training phase			
Word						
1st	29.76	34.88	35.52	46.09	43.24	63.43
2nd	7.43	8.11	8.10	8.46	8.43	8.58
3rd	8.43	8.67	11.23	9.85	9.89	10.89
4th	8.24	6.99	9.03	7.85	7.17	7.47
5th	9.26	7.66	7.29	7.53	8.74	8.34
6th	6.44	5.56	6.49	5.77	7.65	6.14
Average (Words 2–6)	7.96	7.40	8.43	7.89	8.37	8.29
			Test phase			
Multiple Keypress						
1	9.17	14.85	9.08	9.66	11.07	15.45
2	7.63	9.99	8.30	8.34	9.38	9.53
Garbage						
Beginning 1	6.57	7.42	6.72	7.87	8.08	7.40
Beginning 2	6.93	8.72	8.17	7.73	8.10	8.81
Middle 1	16.12	14.36	10.80	11.98	14.30	14.21
Middle 2	10.03	10.42	10.38	9.95	14.76	11.81
Middle 3	9.38	10.28	11.67	9.83	11.10	11.18
Middle 4	7.94	10.86	10.03	10.81	11.05	9.30
Average (Words 2–8)	9.23	10.29	9.43	9.50	10.97	10.32

[a]Matches first task. [b]Does not match first task.

Results

Training tasks. Table 2 presents the time required by each group to do the various deletion tasks. Several analyses were performed on the data. First, in order to examine whether there was an effect of a principle or example on initial task performance, I carried out an analysis of variance (ANOVA) on the time participants required to do the first deletion task. This analysis found that both a principle, $F(1, 67) = 4.67, p = .03, MSE = 556.90$, and example, $F(2, 67) = 4.92, p = .01$, reduced the time to carry out the first task. The interaction of the two variables was not significant, $F(2, 67) = 0.64, p = .53$.

Pairwise comparisons (collapsed over principle) were performed in order to examine the performance differences from the example manipulation. The groups receiving the example that matched the first task performed that task significantly faster than the no-example group ($p = .0038$). There was no difference between the matching-example groups and the mismatching-example groups ($p > .20$) or between the mismatching-example and no-example groups ($p > .09$).

Performance on the rest of the training phase deletion tasks was analyzed using serial position (second word, third word, etc.) as the within-subjects variable and example and principle as the between-subjects variables. There was no main effect of principle, $F(1, 67) = 1.05, p = .31, MSE = 13.51$, or example, $F(2, 67) = 1.03, p = .36$, nor was their interaction significant, $F(2, 67) = 0.16, p = .85$. There was a significant effect of serial position, $F(4, 268) = 24.86, p < .0001, MSE = 4.51$, and the interaction of serial position with example, $F(8, 268) = 2.36, p = .02$. No other interactions were significant (both *p*s $> .25$).

Test tasks. An analysis of performance on the first test task, deleting a word requiring multiple presses of the target character, revealed an effect of principle in the expected direction, $F(1, 67) = 5.90, p = .02, MSE = 38.63$. There was at best only a marginal effect due to example, $F(2, 67) = 2.47, p = .09$. The interaction of principle and example was not significant, $F(2, 67) = 1.09, p = .34$.

Performance on the rest of the test tasks was analyzed using serial position as the within-subjects variable and example and principle as the between-subjects variables. There was no main effect of principle, $F(1, 67) = 0.09, p = .77, MSE = 37.96$, or example, $F(2, 67) = 1.69, p = .19$. There was a significant effect of serial position, $F(6, 402) = 19.30, p < .0001, MSE = 16.99$. None of the interactions was significant (all *p*s $> .17$).

Reading time. Table 3 presents the time participants spent reading the deletion instructions. Reading time was broken into four mutually exclusive categories: prior to the first deletion attempt, prior to the first deletion success, prior to the second deletion success, and total remaining time spent reading deletion instructions. Reading time was analyzed using category of reading as the within-subjects variable and example and principle as the between-subjects variables.

There was no main effect of principle, $F(1, 67) = 1.90, p = .17, MSE = 1186.30$, or example, $F(2, 67) = 0.98, p = .38$, nor was their interaction significant, $F(2, 67) = 0.81, p = .45$. There was a significant effect of reading category, $F(3, 201) = 113.48, p < .0001, MSE = 953.25$. No interactions with reading category were significant (all *p*s $> .27$).

Table 3
Time Spent Reading Deletion Instructions in Experiment 1 (in Seconds)

	Example					
	Single press[a]		Multiple press[b]		None	
Time of reading	Principle ($n = 12$)	No principle ($n = 12$)	Principle ($n = 11$)	No principle ($n = 13$)	Principle ($n = 13$)	No principle ($n = 12$)
Before first attempt	94.36	90.17	99.75	97.40	93.68	88.22
Before first success	13.06	21.19	24.97	32.02	19.90	47.33
Before second success	47.92	33.33	45.69	62.64	32.90	62.87
Rest of deletion tasks	0.00	8.46	2.02	0.17	2.06	0.00

[a]Matches first task. [b]Does not match first task.

Discussion

It was hypothesized that a principle and an example would aid initial task performance. These hypotheses were supported. This suggests that both factors aid learners in instantiating general instructions, although presumably through different mechanisms. The performance benefit from the principle is consistent with the claim that information that helps learners infer the steps they should take will aid them in carrying out an initial task using general instructions. The significant effect of example, coupled with the pairwise comparisons, shows that participants receiving an example that matched the first training task outperformed the no-example participants. There was also a trend for the mismatching example participants to outperform the no-example participants. A conservative conclusion from these results is that a matching example aids performance. A less strict conclusion is that an example that is not isomorphic to the initial task can also aid performance, at least relative to having no example, when coupled with general instructions.

The significant effect of serial position for the training phase indicates that participants generally were faster on later training tasks. The fact that the principle did not interact with serial position during the training phase is consistent with the assumption that learners were operating on a similar representation after completing the first training task.

The interaction of serial position and example is difficult to interpret. An examination of Table 2 shows that from the third word deletion task until the sixth, only the multiple-press example group consistently got faster. It is not clear why the other groups had a less consistent pattern.

It was hypothesized that a principle would aid performance on a later novel task because the procedure learned during training would need to be adapted and a principle was expected to make the adaptation easier to carry out. This hypothesis was supported. This suggests that when an alteration to the procedure was needed—multiple presses of the target character rather than a single press—participants could access their knowledge of the relevant principle in order to figure out the alteration more quickly.

There was a trend toward superior performance by the multiple-press example group on the first novel task. However, the weakness of this result suggests that the effect of an example may be unreliable on a task similar to it if a series of intervening tasks have occurred. This result is consistent with the findings of Pirolli (1991) who found in the domain of programming that examples appeared to aid the initial acquisition of a skill but play little role in later performance. Nevertheless, it is unclear why a principle should be accessed when needed for a later task while an example is less likely to be accessed. One possible explanation comes from a study by Marshall (1991) examining students solving arithmetic word problems. She found that students possessing abstract knowledge, as determined by interviews, preferred to use it instead of details from example and were also the most successful solving problems. Marshall suggested that these more successful learners may first encode an example and then build an abstract network around it. The use of general instructions in the present study may have facilitated this approach, thus making the example less accessible for later tasks but providing learners with a more robust procedural representation.

The effect of serial position in the test phase reflects the difficulty participants experienced when faced with the first one or two deletion tasks involving garbage text in the middle of words (see Garbage Middle 1 and Garbage Middle 2 in Table 2). These tasks involved for the first time placing the cursor inside a word rather than at the beginning in order to do the task. This feature was not inferable from the principle or illustrated in an example. Nevertheless, although this feature slowed performance, it obviously did not cause great difficulty for participants. The lack of an effect of principle or example for the tasks beyond the first test task is consistent with the training phase finding that after completing a novel task, learners tend to have similar procedural representations for that type of task despite differing instructions.

The reading time analysis indicates that there was not a significant overall time cost associated with the principle or examples being added to the general instructions. Whereas the group lacking both principles and examples had the shortest reading time prior to the first task and a group receiving both a principle and an example had the longest time, the difference was not dramatic and was less than half the size of the performance difference between these groups on the first deletion task. Although the inclusion of a principle and example increased the number of words in the instructions, this additional information may have helped participants comprehend Step 7 more rapidly. The additional time to read the principle and example may have been balanced by the reduction in time to comprehend the instructions for that step. It appears that factors such as when a participant was reading the instructions was a more important influence on reading time.

Experiment 2

The primary goal of Experiment 2 was to test whether the findings from Experiment 1 would extend to other tasks. Although the general design of Experiment 2 was very similar to that of Experiment 1, it differed in that there was not a matching versus mismatching example manipulation. Rather, the example manipulation was example (matching the first task) versus no example. As a side effect of this simplified design, there was not a condition in which the first novel task in the test phase matched the example.

The tasks used in Experiment 2 involved formatting a document: changing line spacing, margins, pitch, header spacing, and footer spacing. These tasks are structurally similar in that they are accomplished by selecting an item from an initial format menu and then going through lower level menus until a field appears in which a value for some format feature can be specified. The challenge to learners is to understand the notion of menu hierarchy and to exploit it when attempting tasks that involve new menus.

As in Experiment 1, it was expected that information about the device that could help learners infer or comprehend procedural steps would aid initial performance and possibly later transfer. The format instructions were identical for all participants except for Steps 6 and 7 (see Table 4). The principle seen by half of the participants was included in Step 6 and briefly explained the notion of a menu hierarchy. It was hypothesized that this explanation would help participants understand that the menus were connected and that if a particular menu did not seem to match the current goal, it was possible that a submenu connected to that menu might be appropriate. Participants not seeing this principle were hypothesized to be less likely to understand this relationship, at least early in learning. The result of understanding the menu structure could lead to an increased willingness to explore the menu structure in order to find the items that might be needed for certain tasks.

The example in Step 7 seen by half of the participants illustrated how to carry out the first task, which was changing line spacing. The effect of having an example showing the specific menus to be traversed for initial tasks should have been powerful. In this case, participants knew exactly which choices to make. The only variable would be the exact value to fill in the field; the example illustrated changing the spacing from double to triple, whereas the first task involved changing it from double to single. Changing the value of a variable in a problem that is isomorphic to a studied example is typically easy for college-aged learners (Catrambone, 1994; Reed et al., 1985). Thus, it was hypothesized that participants receiving this example should carry out the first task more quickly than those without the example.

The test phase tasks were also structurally similar to the training tasks but involved choosing an unfamiliar item from the top level format menu. Besides being unfamiliar, this item was semantically unrelated to the tasks that had to be performed. Catrambone (1990) found that a sizeable number of participants performing menu-based tasks required a hint when performing a new task that involved traversing unfamiliar menus. Although some participants spontaneously began to search through the menus in order to find items that seemed related to the target task, about one third had to be prompted to try this strategy. If the principle helps learners to infer steps, then participants receiving the principle might be more likely to develop the search strategy sooner.

If participants receiving the principle perform similarly to those not receiving it on the initial test phase task, this would suggest that the principle was relatively ineffective at this stage or that the practice received during the training phase was sufficient to help all participants learn the notion of connected menus and thus search the menu structure for new items. Conversely, if participants receiving the principle perform better on the initial test phase task compared with no-principle participants, this would suggest that the principle was valuable in helping participants realize they could search through the structure in order to find needed items. It would also suggest that the

Table 4
Formatting Instructions: Changing the Format of a Document

Step 1:

CHANGING THE FORMAT OF A DOCUMENT

Suppose you have a document on the screen and you want to change some formatting feature. Use the arrow keys to move the cursor so that it is under the first letter of the first word of the document.

Step 2:
Press the F7 key.

Step 3:
Notice that the following "main menu" appears at the bottom of the computer screen
Line—format   Page—format   Tabs
Also notice that Line—format is highlighted.

Step 4:
Any time you wish to make a format change, you have to first choose an item from the main menu that corresponds to the type of change you wish to make.

Step 5:
To choose an item from a menu, you do the following two things:
(1) Highlight the item you want. Use the right and left arrow key to do this (unless the item is already highlighted).
(2) When you have highlighted the item you want, then press the ENTER key in order to actually choose the item.

Step 6:
Once you have chosen a menu item, one of two things will happen. Either a new "sub-menu" will appear at the bottom of the screen or a "request" for information will appear at the bottom of the screen. If a sub-menu appears at the bottom of the screen, then choose the menu item you want using the arrow keys (if necessary) and then press the ENTER key. If a "request" appears at the bottom of the screen, press the ESC key to erase the current information that is to the right of the request and then type in the new information.

Additional information for principle participants that followed the previous text:

Most menus for making format changes have other menus connected to them. So when you choose a menu item, this will probably cause another menu to appear. This process of going from "higher level" menus to "lower level" menus will continue until you reach a menu that is as "low level" as possible. When you choose an item from one of these low level menus, the computer will then be prepared to accept new information about a particular formatting feature.

Step 7:
After you type in the new information (usually a number), press the ENTER key.

Additional information for example participants that followed the previous text:

Here is an example. Suppose the line spacing of your document was double spacing, and you wanted to change it to triple spacing. You would first select Line—format from the menu that reads   Line—format   Page—format   Tabs
and then you would choose "Spacing" from the menu that appears next.
When you choose "Spacing" you would then see the following information printed at the bottom of the screen   Enter line spacing (1, 2, or 3): 2
You would then press the ESC key, type a 3, and then press the ENTER key.

Table 4 (*continued*)

Step 8:
Notice that now the main menu reappears at the bottom of the screen.
Step 9:
Press the ESC key once in order to make this menu disappear.
Step 10:
You are now finished making the format change.
Note though that the document on the screen might still look the same after you make a format change; the format change would be seen if you printed out the document on paper.
One general piece of advice is that if you are making a format change and want to cancel the last step you did, press the ESC key. END

Note. ESC = escape.

practice during the training phase was not sufficient to help participants learn the notion of connected menus but rather simply helped reinforce a particular set of procedural steps.

Method

Participants. Participants were 60 (40 men and 20 women) college-aged students at the Georgia Institute of Technology who received course credit. As in Experiment 1, participants' computer experience was confined to using word processing software on a Macintosh.

Materials. Participants read instructions for the following procedures: retrieving a document, changing the format of a document, inserting text, and leaving a document. As in Experiment 1, each procedure was broken into a series of screens. Table 4 presents the instructions for making format changes.

The format instructions were identical for all participants except for Steps 6 and 7 (see Table 4). The principle briefly explained the notion of a menu hierarchy, and the example demonstrated how to change line spacing from double spaced to triple spaced.

Procedure. The general procedure was similar to that of Experiment 1. As in Experiment 1, participants performed the tasks on an IBM PS/2 Model 80 computer. After being told that their keystrokes would be collected and time-stamped by the computer, participants were shown the first document on which they were to work. They were explicitly told that the tasks they would be doing (other than inserting text and retrieving and exiting documents) were format tasks. The document's name and the various changes that were to be made were indicated on each document. Subsequent documents were similarly annotated.

Training phase. During the training phase, participants changed the line spacing, margins, and pitch three times each as well as doing insertion tasks and document retrieval and exiting. The three format tasks all involved choosing the Line—format menu option and then choosing the relevant items (spacing, margins, or pitch options) from the resulting submenu.

The three training documents each required the following tasks in order: (a) retrieve the document, (b) change the line spacing, (c) insert a phrase, (d) change the margins, (e) insert a phrase, (f) change the pitch, and (g) exit the document. The rules for reading and rereading instructions were the same as in Experiment 1. Because of software difficulties, reading times were not recorded.

Test phase. During the test phase participants performed two types of format tasks three times each. Each task was performed once in each of three documents. These new tasks involved selection of a new item (Page—format) from the main format menu. The first type of task involved changing the spacing of the header of a document, and the second involved changing the spacing of the footer of a document. The structure for these tasks was the same as those performed in the training phase, but the top-level menu choice was new, as were the resulting submenus.

Design. The between-subjects independent variables were presence of a principle and an example. These variables were crossed, forming four groups of 15 participants each. The within-subject variable was serial position and phase (training vs. test). The dependent variable was time to perform each format task. Participants performed a total of 15 format tasks.

Results

The time for each format task was determined as follows. In the training phase, the time span was measured from the time when the function key was pressed until the new format value was typed. The measure of time to do a task in the test phase differed depending on whether the task was the first one for a particular document. If the task was the first one, then the time was measured from the time when the function key was pressed until the format value was typed. If the format task was not the first one for a document, then the time span was measured from the time when the prior format task was finished until the value was typed for the current task. This different measure was necessary for test phase tasks because participants could retain the format menu during all the format tasks for a particular document because there were no intervening insertion tasks.

Training tasks. Table 5 presents the time required by each group to do the various format tasks. Several analyses were performed on the data that were analogous to the analyses in Experiment 1. First, in order to examine whether there was an effect of a principle or example on initial task performance, I carried out an ANOVA on the amount of time participants required to do the first occurrence of each type of format task. On the first line-spacing task, there was a significant effect of principle, $F(1, 56) = 4.08, p = .048, MSE = 584.37$, and example, $F(1, 56) = 5.58$, $p = .02$. The interaction of the two variables was not significant, $F(1, 56) = 1.47, p = .23$. There was no effect of principle, example, or their interaction for the initial margin or pitch tasks (all $p > .27$).

Performance on the rest of the training phase format tasks was analyzed with serial position as the within-subjects variable and example and prin-

Table 5
Time to Perform Format Tasks in Experiment 2 (in Seconds)

	Example		No example	
Task	Principle	No principle	Principle	No principle
Training phase				
1st line spacing	36.99	42.04	44.17	64.35
1st margins	37.56	36.62	40.33	46.30
1st pitch	28.36	28.69	27.48	27.91
2nd line spacing	24.31	25.56	21.82	20.47
2nd margins	17.68	19.64	18.68	18.77
2nd pitch	18.17	18.13	18.87	19.37
3rd line spacing	13.83	14.87	13.67	12.19
3rd margins	15.64	14.98	16.67	17.73
3rd pitch	17.35	15.78	16.66	16.58
Average (2nd & 3rd occurrence of each task)	17.83	18.16	17.73	17.52
Test phase				
Spacing				
1st header	22.19	30.22	26.99	33.83
1st footer	16.11	16.64	16.34	17.66
2nd header	13.78	14.91	10.76	14.31
2nd footer	17.61	17.70	17.43	17.22
3rd header	13.20	15.81	14.61	16.06
3rd footer	13.76	14.37	13.00	12.86
Average (2nd & 3rd occurrence of each task)	14.58	15.70	13.95	15.11

Note. $n = 15$ for each of the four samples.

ciple as the between-subjects variables. There was no main effect of principle, $F(1, 56) = 0.0, p = .98$, $MSE = 380.92$, example, $F(1, 56) = 0.03, p = .86$, or their interaction, $F(1, 56) = 0.02, p = .90$. There was a significant effect of serial position, $F(5, 280) = 10.99$, $p < .0001$, $MSE = 54.69$. No interactions with serial position were significant (all $ps > .36$).

Test tasks. An analysis of performance on the first test task, changing the header spacing, revealed an effect of principle, $F(1, 56) = 5.90, p = .02$, $MSE = 140.42$. There was not a significant effect from example, $F(1, 56) = 1.89, p = .17$, or the interaction of principle and example, $F(1, 56) = 0.04$, $p = .85$. An analysis of performance on the second test task, changing the footer spacing, found no effect of principle, example, or their interaction (all $ps > .53$).

Performance on the rest of the test tasks (the second and third occurrences of the header and footer spacing tasks) was analyzed with example and principle as the between-subjects variables and serial position as the within-subjects variable. There was no main effect of principle, $F(1, 56) = 0.86, p = .36$, $MSE = 89.58$, example, $F(1, 56) = 0.25, p = .62$, or their interaction, $F(1, 56) = 0.0$, $p = .98$. There was a significant effect of serial position, $F(3, 168) = 6.26, p = .0005, MSE = 34.51$. No interactions with serial position were significant (all $ps > .58$).

Discussion

As predicted, performance on the initial format task was aided by the presence of a principle and an example. The remainder of the training tasks were affected only by serial position. Taken together, these results suggest that the process of completing the first task left participants with similar representations for carrying out the rest of the training tasks. The speed-up over the rest of the tasks is a reasonable effect of practice.

When participants were faced with their first novel task in the test phase, those who studied the principle were again the fastest performers, as had been predicted. This is consistent with the hypothesis that the principle helped these participants to reason more effectively about how to carry out a novel aspect of the task, namely, finding the appropriate menu. Although a formal measure was not taken, observation of participants suggested that the difficulty faced with the first header-spacing task was finding the appropriate menu path (starting with Page—format). Principle participants tended to begin to explore the menu structure sooner than the no-principle participants. There was no effect of principle when participants performed the first footer-spacing task, probably because most of them noticed the relevant footer menu item when they found the header menu item.

Finally, performance on the remainder of the tasks showed an effect only of practice. Again this is consistent with the interpretation that during the process of successfully completing a task, learners are likely to form a representation that will allow them to perform the same tasks successfully in the future.

General Discussion

Prior work has suggested that instructions written at a level general enough to apply to a wide variety of tasks have the benefit of helping people apply the procedure to later unfamiliar tasks. However, their disadvantage is that they are hard to apply initially compared with instructions tailored to specific tasks (Catrambone, 1990). The current study was designed to examine whether the use of principles and examples can help people use general instructions more effectively.

It was conjectured that an example could aid initial performance because learners have been frequently shown to be quite good at applying an example to a simiiar problem (e.g., Catrambone, 1994; Reed et al., 1985). However, people typically have difficulty generalizing examples. Thus, it was unclear whether when general instructions were combined with examples, the result would be more a function of the strengths or weaknesses of these two types of information.

The benefits of a principle were hypothesized to manifest themselves in initial start-up as well as transfer. However, given that prior work had found that general instructions supported transfer, it was unclear whether a principle could improve performance further.

The results indicate that both a principle and an example aid the application of a general procedure to an initial task, and there appears to be no interaction between these factors. The presence of an example that matches the initial task appears to aid performance more than an example that does not match the initial task. In addition, the presence of an example does not appear to retard transfer to later, novel tasks. The fact that the combination of general instructions with an example leads to superior initial performance without hurting transfer suggests that the strengths of these two information sources dominate their weaknesses when they are combined. In addition, participants receiving a principle were aided in transfer to novel tasks, suggesting that even the relatively good transfer produced by general instructions can be enhanced by the presence of a principle. Finally, the reading time results from Experiment 1 indicate that adding principles and examples to general instructions does not have to lead to longer overall reading times.

The present study serves as a demonstration that general instructions can be improved through the addition of examples and principles. This demonstration was by no means a foregone conclusion given prior work, for instance, on the transfer difficulties that examples can cause. Nevertheless, detailed guidelines for creating general instructions, examples, or principles cannot be created based on the present work.

It is speculated though that as instructions move more toward generality, the importance of principles and examples in improving start-up performance is likely to increase. The attempt to create the materials for the study made it clear that it is crucial when writing documentation to consider the cases that one feels belong to the same task category. This decision will affect how general the resulting general instructions should be: For instance, is *highlighting* a task? Is *deleting a word* a task? Is *deletion* a task? Is *text manipulation* a task? The answers to these questions will remind writers to keep checking whether the instructions being created are at least formally sufficient (as measured perhaps by a production rule analysis) to accomplish all instances of the task that they can reasonably imagine.

Implications for Models of Procedural Learning

The results from this study have implications for the procedural representations learners form from instructions and from actually doing tasks. It appears that although manipulations to instructions can affect initial task performance, the fact that learners with and without a principle or example performed similarly on subsequent similar tasks suggests that a representation for the procedure, or a revision of the original representation, is formed once the initial task is carried out. This new representation controls performance on tasks that are similar to the initial task. This speculation is consistent with the present results, as well as the results of Ross and Kennedy (1990) who found that learners seem to form a generalization when they apply an example to a new problem and then use this generalization for later tasks.

The suggestion that learners rely on the representation of the procedure formed while doing the first task is also consistent with the notion of the formation of a situation model from an initial text base (Kintsch, 1986). The *text base* is the representation a learner forms while initially reading a text, whereas the *situation model* is a representation of the situation described in the text. With respect to the present study, the text base would be the representation formed from the instrutions, and the situation model might be the representation resulting from applying the instructions to the initial tasks. Kintsch (1986) argues that people use the situation model in order to perform tasks such as inferencing. The present results indicate though that when the situation model does not sufficiently support inferencing for a novel task, learners can still access the text base for help.

Implications for Design of Instructions

Instructions written at a general level have the benefit of being shorter because they do not need to cover many specific cases. This is an appealing feature to users because they usually dislike reading large amounts of instructional material (Mack et al., 1983). However, the danger is that the general instructions will be difficult to use initially because the user will not be given the exact steps for each situation. The results of the present study indicate that this difficulty can be reduced through the use of examples and through the use of principles that aid inferencing. In addition, general instructions augmented with principles could lead to better performance on later tasks because the principle might help the learner avoid forming incorrect associations between superficial details of specific instructions and the more basic rules for interacting with the device.

The fact that overall reading time was not significantly increased by the additional text required to provide the principle and example in Experiment 1 provides added incentive for the documentation approach outlined above. Obviously, if enough text is added in order to provide a principle or example, reading time must increase at some point. Perhaps parameters can be derived to guide one's determination of when the increase in reading time is likely to be greater than time saved in performing initial tasks.

Future Work

The experiments presented here were neither designed to test detailed hypotheses concerning how principles and examples aid the application of general instructions nor specific ways of constructing principles or examples. For instance, it may be the case that general instructions benefit from principles as a function of the ambiguity of the steps. Steps might be ambiguous because of the nature of the domain and the variety of tasks that the instructions cover. However, users may sometimes perform better with a well-written minimalist set of instructions if ambiguity can be controlled (Carroll, Smith-Kerker, Ford, & Mazur-Rimetz, 1987–1988). Conditions favoring minimal versus more detailed instructions can possibly be derived from models of procedural learning (e.g., Anderson, 1983; Kieras & Bovair, 1986).

Another issue for future work rests on Kieras and Bovair's (1984) suggestion that one-way mental models can be useful is when they provide a mechanism to help learners infer a procedure. A study that systematically manipulates the number of steps in a procedure and the number of those steps that can be inferred from a given principle or mental model would provide evidence for the type of situations in which principles would be most useful in helping users comprehend instructions.

The construction of instructions and the principles and examples that could be included in them would presumably benefit by constraints developed through the application and development of models of human cognition. The general instructions used in the present experiments were developed from a set of production rules used to model the knowledge needed to carry out tasks with a particular word processing program (Catrambone, 1990). However, the principle and examples were created from the researcher's intuitions and familiarity with the mental model and problem-solving literatures. Continuing work may focus on developing a systematic approach for the design of principles and examples and their integration into instructions.

References

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

Carroll, J. M., Mack, R. L., Lewis, C. H., Grischkowsky, N. L., & Robertson, S. R. (1985). Exploring exploring a word processor. *Human–Computer Interaction, 1,* 283–307.

Carroll, J. M., Smith-Kerker, P. L., Ford, J. R., & Mazur-Rimetz, S. A. (1987–1988). The minimal manual. *Human–Computer Interaction, 3,* 123–153.

Catrambone, R. (1990). Specific versus general procedures in instructions. *Human–Computer Interaction, 5,* 49–93.

Catrambone, R. (1994). Improving examples to improve transfer to novel problems. *Memory & Cognition, 22,* 606–615.

Catrambone, R., & Holyoak, K. J. (1990). Learning subgoals and methods for solving probability problems. *Memory & Cognition, 18,* 593–603.

Chi, M. T. H., Bassok, M., Lewis, R., Reimann, P., & Glaser, R. (1989). Self explanations: How students study and use examples in learning to solve problems. *Cognitive Science, 13,* 145–182.

Cooper, G., & Sweller, J. (1987). Effects of schema acquisition and rule induction on mathematical problem-solving transfer. *Journal of Educational Psychology, 79,* 347–362.

Hale, C. R. (1993). *The role of explanations in symptom–fault rule acquisition.* Paper presented at the 34th Annual Meeting of the Psychonomic Society, Washington, DC.

Kieras, D. E. (1990). The role of cognitive simulation models in the development of advanced training and testing systems. In N. Frederiksen, R. Glaser, A. Lesgold, & M. Shafto (Eds.), *Diagnostic monitoring of skill and knowledge acquisition* (pp. 51–74). Hillsdale, NJ: 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., & Bovair, S. (1986). The acquisition of procedures from text: A production–system analysis of transfer of training. *Journal of Memory and Language, 25,* 507–524.

Kintsch, W. (1986). Learning from text. *Cognition and Instruction, 3,* 87–108.

LeFevre, J., & Dixon, P. (1986). Do written instructions need examples? *Cognition and Instruction, 3,* 1–30.

Luchins, A. S. (1942). Mechanization in problem solving. *Psychological Monographs, 54* (Whole No. 248).

Mack, R. L., Lewis, C. H., & Carroll, J. M. (1983). Learning to use word processors: Problems and prospects. *ACM Transactions on Office Information Systems, 1,* 254–271.

Marshall, S. P. (1991). *Schemas in problem solving: An integrated model of learning, memory, and instruction* (Technical Report No. 91-02). San Diego, CA: San Diego State University Center for Research in Mathematics and Science Education.

Pirolli, P. L. (1991). Effects of examples and their explanations in a lesson on recursion: A production system analysis. *Cognition and Instruction, 8,* 207–259.

Pirolli, P. L., & Anderson, J. R. (1985). The role of learning from examples in the acquisition of recursive programming skill. *Canadian Journal of Psychology, 39,* 240–272.

Reder, L. M., Charney, D. H., & Morgan, K. I. (1986). The role of elaborations in learning a skill from an instructional text. *Memory & Cognition, 14,* 64–78.

Reed, S. K., Dempster, A., & Ettinger, M. (1985). Usefulness of analogous solutions for solving algebra word problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 106–125.

Ross, B. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 456–468.

Ross, B. H., & Kennedy, P. T. (1990). Generalizing from the use of earlier examples in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 42–55.

Smith, E., & Goodman, L. (1984). Understanding written instructions: The role of an explanatory schema. *Cognition and Instruction, 1,* 359–396.

Received March 17, 1994
Revision received February 28, 1995
Accepted May 2, 1995 ■