MINIMALISM RECONSIDERED: SHOULD WE DESIGN DOCUMENTATION FOR EXPLORATORY LEARNING?

THOMAS R. WILLIAMS
DAVID K. FARKAS

INTRODUCTION

During the past decade, few efforts aimed at the articulation of new computer documentation models have inspired as much interest within the documentation community as has John Carroll's formulation of the tenets of "minimalism." In our experience, however, minimalism is often incorrectly understood by practitioners simply as a prescription for brevity. Minimalism is, in fact, more complex than this. In "The Minimal Manual" (1987-88) Carroll and his coauthors suggest that minimalism comprises four characteristics: (1) brevity (or perhaps conciseness), (2) focus on real tasks, (3) support of error recognition and recovery efforts, and (4) adoption of an instructional philosophy termed "guided exploration."

Of these four, three are largely non-controversial. Both within and outside of the area of computer documentation, almost everyone applauds brevity and finds fault with unnecessary length. Software publishers, in particular, enthusiastically support an approach to documentation that will save them printing and distribution costs. Also, the minimalist position on error recovery is a position most of us would endorse as good, old-fashioned, common sense. When we give people directions, we regularly include guidance such as "If you reach the freeway entrance, you have gone too far." The goal of letting the user accomplish real work is also sound and generally accepted. Motivation is a powerful driver in the learning process. Certainly it makes sense to facilitate the learner's desire to achieve a practical and meaningful goal while using the program.

The minimalist notion of guided exploration, however, is the most questionable precept of minimalism, and it is the one we wish to examine here. We believe, for reasons we will detail, that guided exploration poses a dilemma to the documentation writer attempting to reconcile minimalist goals with the minimalist means articulated to achieve those goals. In short, we believe that the stated minimalist goal of enabling the learner to accomplish real work while learning a program is often thwarted by the act of compelling that learner to induce, through trial and error, the correct procedures needed to accomplish that work.

GUIDED EXPLORATION AS AN INSTRUCTIONAL THEORY

The notion of "guided exploration" owes its origins to the concept of "discovery learning," popular among learning theorists in the late '60s and early '70s. An approach to learning whose most visible, and perhaps most articulate, advocate was Jerome Bruner, discovery learning broadly embraced the view that principles, rules, and problem-solving techniques ought to be taught with a minimum of teacher intervention and a maximum of trial-and-error learning on the part of the student. In general, the discovery learning movement was a reaction against the "teaching of facts." Bruner summed up the discovery learning philosophy as follows:

> We teach a subject, not to produce little living libraries from that subject, but rather to get a student to think...for himself, to take part in the process of knowledge getting. Knowing is a process, not a product. (1966, p 72)

Discovery learning is fundamentally inductive. General principles are distilled by the learner from exposure to specific, and sometimes narrow, instances of those principles. Advocates of discovery learning argued that it improved the memorability of the material learned and that it fostered motivation because discovery itself is inherently motivating.

Briefly, "guided exploration," minimalism's version of discovery learning, embodies the belief that learners learn best when allowed to explore a program and to discover for themselves how it works. To encourage such discovery, Carroll advocates the use of open-ended exercises, sometimes preceded by the invitation to "try them and see." His manual also includes "On Your Own" exercises in which the user presumably can practice skills acquired through previous exposure to incompletely articulated procedures: "Procedural details were deliberately specified incompletely to encourage learners to become more exploratory, and therefore, we hoped, more involved in the learning activity...." (Carroll et al., 1987-88, p. 129). The assumption underlying this approach is that guided exploration capitalizes on the strategies that most users would spontaneously adopt if they were left to learn on their own.

It is certainly the case that people apply many different strategies when they attempt to learn computer software and that some of those strategies can legitimately be characterized as exploration. While some learners do, of course, work through print or online tutorials step by step or, consulting a user's guide, read each conceptual overview and then conscientiously read and follow each step of the associated procedure, many others disdain documentation entirely and, more often than not, attempt to learn computer software through unsupported problem-solving activities. There is, in fact, broad anecdotal (and some empirical) evidence that many users hate to read manuals and attempt to learn software without the aid of documentation even when documentation is available (Horton, p. 253; Brockmann, p. 31). The important question we wish to address, however, is this: Is it worthwhile to create forms of print or online documentation designed specifically to <u>compel</u> exploratory learning?

CHALLENGES TO GUIDED EXPLORATION

We would like to pose three challenges to the guided exploration model that we believe need to be addressed before documenters rush to the conclusion that the standard documentation set should include, or should be replaced by, documentation embracing the guided exploration philosophy. We offer these challenges in the hope that they will inspire enlightened discussion as well as future empirical research aimed at the resolution of issues important to all of us in the documentation community. The challenges are these: We believe (1) that in all but the most simple applications guided exploration may be both inefficient and ineffective, (2) that guided exploration is excessively authoritarian in nature—that it unnecessarily dictates to learners the depth of learning that must be achieved, and (3) that it focuses on the learner's acquisition of declarative knowledge at the expense of his or her acquisition of procedural knowledge.

Efficiency and Effectiveness

The notion of learning by discovery in domains other than computer documentation has not been without its critics, and their assessments of the liabilities inherent in discovery learning seem as germane to computer documentation as they do to fourth-grade social studies textbooks. Among the most compelling criticisms has been the claim that discovery learning is inefficient that there is little reason to compel a student to suffer through a lengthy trial-and-error process to solve a trivial problem that a teacher (or manual) could explain in a few moments (VanderZanden & Pace, 1980). If learning is inefficient, of course, it is not likely that it will facilitate the minimalist goal of helping the learner to get his or her real work done. Not all problems presented to a computer user, of course, are trivial, but surely simple command sequences and simple procedures seem likely candidates for characterization as trivial problems whose solutions might more efficiently be provided explicitly by the documentation rather than inferred by the documentation user.

But what, then, about nontrivial procedures? Would it be worthwhile to define a "difficulty threshold" and to adopt guided exploration as the instructional technology of choice for those procedures that exceed that threshold? Again, this strategy presents a dilemma to the documentation designer. We would argue that if the procedures that must be learned are difficult to induce, the probable consequences are (1) an increased likelihood that the procedures will be induced incorrectly (a point we will address in more detail subsequently), and (2) a significant reduction in the learner's motivation.

As Carroll notes, "The most important factor in learning is learner motivation, but this is also the factor least amenable to extrinsic control via design. If learners want to undertake a particular activity, then letting them attempt to do so is perhaps the best design step we can take" (Carroll et al.,1987-88, p. 126).

We would agree with Carroll up to a point. Letting users attempt to accomplish real goals as opposed to sitting them in front of a computer and telling them to do what to them appears to be meaningless exercises is certainly likely to positively influence their motivation. However, motivation is even more greatly influenced by the learner's expectations that his or her efforts will result in success, and those expectations, in turn, are greatly influenced by his or her perceptions of how easily a required task is likely to be accomplished (Weiner et al., 1976). It is doubtful to us that compelling a learner to induce a difficult set of procedural steps by discovering them through trial and error is likely to be perceived as easily achievable and, consequently, a candidate for a high probability of success.

Indeed, it seems that a manual structured in such a way as to allow the user to immediately articulate and pursue his or her own goals successfully would provide adequate incentive for involvement in the instructional materials. Forcing a new user to inductively arrive at the definition of the desired procedures through exploration seems neither an efficient nor humane way to teach those procedures to a person who really only wants to "get his [or her] work done."

Another potentially negative consequence of compelling the learner to induce a principle or set of procedures is that—absent varying degrees of instructor intervention—there is no way to ensure that the principles and procedures "discovered" will be correct. It is difficult to see how erroneous procedures would facilitate the minimalist goal of helping the reader to do real work.

The author of exploratory documentation, as a consequence, is necessarily engaged in a murky and risky enterprise. Murky because inference is a poorly understood issue. All discourse is partly inferential (Hirsch, 1977), but there is no metric for degree of inference at either large or small units of discourse, and there are, in fact, various kinds of inference, though the relationships among them are not very clear. These include letting the user infer a particular concept from a generalization, a general concept from a particular one, and a new concept from one or more coordinate concepts.

The enterprise is risky because the author is necessarily attempting to hit a very precise degree of inference in *regard to a given audience segment*. He or she asks: is this a good place to reduce conceptual information or does it demand some sort of fuller treatment? Is this a step to present in "hint" form or should this step be explicit and the next one presented as a "hint"? To the extent that the exploratory nature is diminished, we have standard discourse, with no special qualities. In some instances, very small inferences might not even be noticed by the reader or might be perceived simply as a lively, engaging style of writing. In fact, reading Carroll's studies and looking at samples of his documentation, we are often struck by the very modest degree of inference and the significant similarity between his minimal manuals and typical user's guides.

On the other hand, as conceptual material is reduced and procedures are made less complete and more inferential, the greater the possibility of defeating the user's attempt to learn or pushing the user into one of the other pitfalls of exploratory learning. If there is no safety net beneath this kind of documentation, the user may be in serious difficulty indeed. Brockmann (1990), while he largely accepts the claims for minimalism, also acknowledges the riskiness in exploratory documentation: "Learning by self-discovery [sic] is less predictable and gaps or lack of depth in learning may appear" (p.100).

A corollary to this problem is that each exploratory manual must be written for a highly specific audience and must be tested for that audience very carefully. Indeed a central theme in *The Nurnberg Funnel* (Carroll, 1990) is that minimalist documentation must be developed through continual reiterative testing. We certainly endorse careful adaptation to different audiences and reiterative developmental testing of any documentation, but at the same time we note that it is far safer and easier to strive for completeness and clarity. Both conceptual/descriptive overviews and procedural steps, we maintain, should be explicit, and should include all the information that the author reasonably feels the reader will want or need. Very often, in fact, authors inadvertently leave out key assumptions and bits of necessary information because they themselves have grown too familiar with the software and fail to recognize the user's information needs. But while authors must, of course, resist the impulse to dump excess information on the user, a sound and reasonable goal is a "straightforward" strategy, attempting to present information that will be perceived as explicit and complete.

It is very often true that you cannot meet the needs of novices and experts in a single presentation of information, but within certain limits, this straightforward approach to documentation makes it possible to target a broad segment of the user base. Advanced users have some tolerance for documentation that seems simple for them, and all computer users are free to read, skim, or skip and in so doing exercise the degree of exploratory learning and problem-solving that they choose at any given moment. What is needed is carefully worked out information access and intelligent page or screen design. Both in the print and online worlds, users can be efficiently directed to specific items of information using well-conceived indexes, tables of contents, or electronic information retrieval facilities. Pages and screens can be designed to enable the user to immediately distinguish between conceptual/descriptive overviews, procedural steps, and—when necessary—extra levels of detail or explanation intended for a subsection of the audience.

Carroll (1990), recognizing the very real possibility that users will become untracked when following exploration documentation, places a major emphasis on designing for error recovery. Error recovery sections are included in his print designs (p. 86-87), and his "training wheels interface" (Chapter 7) blocks undesirable user actions in online tutorials. With extensive support for error recovery, the situation improves. Incorrect inferences need not be fatal, and the ongoing process of trial and error may be instructive. But even with error recovery facilities, exploration documentation remains, we believe, a marginal strategy.

The Authoritarian Nature of Guided Exploration

Guided exploration that is compelled by the documentation author takes out of the hands of the user a very important decision: which procedures in this application do I want to learn in depth and which do I need to know only superficially (and can perhaps forget immediately)? A simple point that advocates of exploratory documentation, it seems to us, have failed to acknowledge is that it is sometimes unnecessary to "learn" a set of procedures in order to "execute" them. Carroll himself reports the comments of one of the learners involved in one of his studies that seems to affirm this notion: "I want to do something, not learn how to do everything" (Carroll et al., 1987-88, p. 126).

An example may be appropriate here. The great majority of professional journals have adopted the convention of using endnotes rather than footnotes. But if an author were to prepare a manuscript for a journal that does still follow the footnote convention, he or she would most probably prefer a simple set of procedures that would facilitate the goals of getting the job done with a minimum of time and effort, not a learning environment in which he or she is compelled to induce, or, "learn," the procedures to the point of mastery. "Learning" need not necessarily be a prerequisite to "executing" a set of procedures. What is needed in this particular instance is documentation in a form that will allow the author to create the footnotes he or she needs. Given the improbability that he or she will ever need to create these kinds of footnotes again, there really is no need to compel the author to "learn" or "remember" the procedures that he or she explicitly needs now in order to create the footnotes.

Perhaps an even simpler example might further clarify the issue. We do not need to "learn" a phone number in order to dial it. It may be worth our while, of course, to learn those phone numbers that we anticipate we will need to dial frequently. It is not worth our while, however, to commit to memory, or learn, a phone number that we expect to use only once. For that phone number we will rely on a phone book, and once we have used the information we need in order to dial that number, we will, appropriately, forget it. Indeed, we would view it as quite presumptuous of the phone company to insist that we memorize the number. And, we might find it particularly galling to find that the phone company had structured our efforts to learn so that we were compelled to deduce the number, for example, by discovering the number's prefix through the use of an area map and its last four digits through an arithmetic puzzle (e.g., "The last four digits sum to 33; the first digit is twice the third....). Admittedly, having invested the mental effort to discover the appropriate number, we would likely retain it in memory for quite some time. Indeed, it is not an uncommon view among cognitive scientists that, within limits, the more cognitive effort the learning process demands, the more durable the learning from it.

It is not necessarily the case, however, that this kind of investment of mental energy on my part is an efficient allocation of my mental resources. And, as suggested previously, might it not be considered presumptuous of an application's documenter to decide on my behalf that in order to execute a series of commands that I should have "learned" them? Importantly, it is really only those procedures that I expect to use repeatedly in my work that I need to learn—that are, in other words, worth the investment of my mental effort to learn.

Learning to "know" versus learning to "do"

An enormous body of cognitive research has, during the past two or three decades, focused on how our knowledge is acquired, how it's stored, how it's structured, and how it's activated. In their attempts to define the nature of cognitive architecture, cognitive theorists recognized early the desirability of distinguishing between "declarative knowledge" on the one hand and "procedural knowledge" on the other (Anderson, 1976). The difference, in other words, between "knowing that" and "knowing how." Redish (1988) applies this distinction to computer manuals in terms of their ability to support the user's goal of "reading to do" as opposed to "reading to learn."

Cognitive information processing models, of course, make no claims that they reflect underlying physiological architecture, but it is interesting to note that even physiological studies seem to confirm the distinction. Cohen and Squire (1980) observed that some amnesic patients who had suffered temporal lobe damage, for example, could learn procedural skills as quickly as patients with normal memory, but could not remember having received training in the skills tested nor instances of having been tested on their acquisition of those skills.

Both declarative and procedural knowledge are undoubtedly intricately linked in cognitive structure. In the execution of goal-directed procedures, in fact, it is through the "controlled processes"—our conscious manipulation of declarative knowledge—that our procedural goals are articulated and maintained. Conversely, the execution of well-learned procedures generally requires little conscious monitoring on our part, a property that has led them to be referred to "automatic processes." Indeed, one of the important distinctions between declarative knowledge and procedural knowledge is the degree to which conscious attention needs to be employed in their use. Factual (declarative) knowledge exerts considerable demands on the limited attentional capacity available to us in working memory. Alternatively, we devote very little attention to procedures that we use routinely and frequently in our everyday lives. The difficulty we often encounter in teaching children such simple tasks as tying shoes is that, as adults, we have long since forgotten the declarative knowledge that originally served as the foundation for our creation of the procedural knowledge we employ in tying our own shoes. Teaching a child is consequently difficult— not because tying shoes is a difficult task for a child to learn, but because we no longer have access to the kind of knowledge the child needs in order to learn the procedure. Decomposing the procedure into discrete steps that can be followed by a child seems almost to require that we relearn the task ourselves. Tying shoelaces is something that we do automatically—without devoting much conscious attention to it. In fact, most of us tie our shoes while listening to the radio, conversing with a family member, or planning for the day ahead.

Of course, this is also the kind of facility we develop over time in the use of a software application. Procedures we initially had to look up in a user's guide we learn, with practice, to employ automatically—without much, if any, conscious attention. What has led to this level of "expertise" is the conversion of declarative knowledge into procedural knowledge, a process referred to in cognitive science as "proceduralization" (Gagne, 1985).

How proceduralization takes place requires an understanding of the nature of a procedure. Anderson (1983) actually characterizes two kinds of procedures: pattern recognition procedures and action-sequence procedures. Pattern recognition procedures allow us to classify things—to acquire, for example, the concept of a "cat." While the specific exemplars of the category "cat" can vary widely in terms of their physical attributes, all cats share certain attributes that allow us to classify individual instances of cats as members of the class of cats. Through processes of generalization and discrimination, we learn to recognize patterns of stimuli that we associate with specific concepts.

Action-sequence procedures, on the other hand, are invoked in our attempts to carry out specific sequences of actions. These procedures are represented cognitively as "productions," or "condition-action rules." Productions have two components: an "if clause" and a "then clause." If the condition specified by the "if clause" is met, then an action specified by the "then clause" can be executed. That action, in turn, may satisfy the condition of the second "if clause." When productions are thus related, they are said to be related through "flow of control." Control of the entire procedure, in other words, passes from production to production until the procedure is completed. Complex procedures, of course, may consist of an enormous number of productions.

When we first begin to learn a skill, productions are performed consciously. In fact, at the novice stage of skill acquisition, performance requires that the learner interpret declarative knowledge. At this stage, performance is really a controlled process requiring conscious attention. While the experienced driver, for example, seldom gives a thought to the procedures that go into the process of starting a car and driving away each morning, the novice driver consciously attends to procedures like looking at the shift lever indicator to determine whether the car is in drive. Similarly, the beginning typist has to consciously think about the location of each key on the keyboard. For that reason, this stage of skill acquisition is often referred to as an "interpretive" stage. Like an "interpreted" computer program, each "command" has to be interpreted in order to run. At this stage, procedures are represented as propositions—as declarative knowledge. Each proposition then must be translated into a set of productions. As Gagne (1985, p. 119) notes, "...having a separate production for each step is not...efficient because after each step is carried out a complete propositional representation of the situation must be reinstated before the next production can apply." The demands on working memory, a memory store with severely limited capacity, are considerable at this stage of skill acquisition.

Truly "skilled" performance of procedural tasks requires a modification of the many productions comprising a task and of the relationships existing among them. This stage of performance is often referred to as the "compiled" stage. At this stage, procedures are "chunked," or combined (Anderson, 1983; LaBerge, 1976). Specifically, action, or "then clauses" are combined. Large productions require less of working memory's capacity than do a great many smaller productions. At the "compiled" stage, performance begins to become automatized—begins, in other words, to reach the level of skill at which little conscious thought is required. In fact, Schneider and Fisk (1982) found that even when persons were instructed to consciously attend to an automatized task, they were unable to improve their performance.

The degree to which a compiled procedure is automatized depends to a great extent on whether it is used or not. As Stillings et al. (1987) note, performance of a skill may continue to improve even after compilation. This is sometimes the result of procedures being "tuned" (Anderson, 1983, Rummelhart and Norman, 1978). And this is what is likely to happen when the user repeatedly uses the same procedure.

What, then, are the implications of what we know about the acquisition of procedural skills for the notion of the minimal manual?

First, given the fact that procedural knowledge begins as declarative knowledge, it would seem that efficient proceduralization would be best encouraged by presenting the learner with a simple, straightforward, easy-to-execute listing of the steps comprising the procedure her or she wishes to perform. Discovering those steps through exploration might, of course, assist the learner in his or her attempts to acquire the declarative knowledge from which the procedures will be derived, but at what cost? Gagne (1985, p. 123) comments on this issue:

> ...it seems logical to expect that easy-to-use sequences would be helpful for proceduralization. It also seems logical that requiring students to memorize, in declarative form, the steps in an action sequence, would *not* be useful. The declarative serves as a way station on the road to a procedural representation and so there seems to be little point to fixing this information in long-term memory. Time might be better spent in many practice trials, with lists of steps or other prompts available for reference as long as is necessary. For most procedures the aversiveness of having to look back at the list to reinstate the declarative representation should provide ample motivation for proceduralizing the knowledge.

Minimalism's aversion to having computer documentation provide the user with complete lists of discrete steps, in light of our understanding of the process of proceduralization, seems arguable. To achieve "expertise" in the performance of a set of tasks, the user of a documentation set needs to acquire procedural knowledge. Some researchers even argue that the principal distinction between novices and experts is that while novices may know a principle or rule, they often don't know the conditions of effective application. Alternatively, when experts access knowledge, that knowledge tends to be functional: it is bound to conditions of applicability (Glaser, 1990). Instructional technologies developed to assist learners in the acquisition of procedural skills "...are guided by the assumption that efficient skill is a foundation for subsequent depth of understanding and planning ability. The theoretical implication is that major metacognitive changes are an unconscious byproduct of highly practiced successful performance" (Glaser, 1990, p 32). The user's need to achieve a mastery of the declarative knowledge underlying those procedures that need to be practiced, consequently, seems largely unnecessary if the documentation can provide the user with that declarative knowledge both clearly and unambiguously.

Are we then arguing that we should go back to the "rote" drill and practice approach to teaching procedures? Most emphatically, the answer to that question is no. "Rote learning" is by definition meaningless learning—learning unassimilated by and disconnected from the rest of our knowledge. And we would agree with Carroll that sitting new users down in front of computers and compelling them to learn meaningless commands or keystroke combinations is very likely to result in rote learning—if it results in any learning at all. The assumption that seems to be implicit in criticisms of drill-and-practice approaches to teaching, however, is that the provision of procedures compels the learner to learn those procedures by rote memorization. We would argue, however, that for the learner who has well-defined goals that themselves have been articulated on the basis of the learner's own needs to do real work, a set of commands that can be employed in the attempt to reach those goals is not meaningless.

Finally, those commands that are employed frequently because of the user's own unique goals and activities associated with his or her attempts to achieve those goals (to do his or her real work) will be practiced, and those commands that need not become an integral part of the user's repertoire will not. The user, then, is allowed to decide what procedures need to be "learned" well and practiced, and what procedures do not.

A BRIEF LOOK AT THE RESEARCH ON MINIMALISM

At this point very few studies have attempted to empirically test the precepts of minimalism. A cursory review of this relatively small body of literature, we believe, fails to adequately inform our understanding of the specific effects of guided exploration—in most instances the potential effects of guided exploration are hopelessly confounded with the possible effects of other variables being studied. The results of two of these studies, nevertheless, have been interpreted as offering evidence in support of minimalist claims. It is specifically that evidence, as well as the somewhat more tangential evidence provided by a few other authors, that we wish to examine in this section.

Because "minimalism" embraces several principles, it is important in assessing the results of studies on minimal manuals to ask to what component of minimalism or to what combination of components of minimalism are those results attributable. At the present time, the answer to that question is that it is difficult to tell. The experimental materials used in the studies reported in "The Minimal Manual" (Carroll et al., 1987-88) consist of a commercial manual and a "minimalist" manual that, presumably, embraces all aspects of minimalism. Consequently, the effects of reducing a 180 page manual to a 45-page minimal manual (brevity) are hopelessly confounded with whatever effects guided exploration might have produced. It is possible, in other words, that the experimenters could have inspired remarkable experimental results by virtue of nothing more than simply reducing a cumbersome 180-page manual into a more concise 45-page manual.

Recognizing this possibility, Black, Carroll, and McGuigan (1987; See also Carroll, 1990) attempted to decompose the effects of different aspects of minimalism in a study comparing user learning performance from four different manuals: a "Skeletal Manual" that contained only "essential information about the system"; a "lengthy manual" that added "descriptive, explanatory and summary information to the Skeletal Manual to yield a manual that [was] approximately 50% longer..."; an "Inferential Manual" about the same length as the Skeletal Manual, but which "guides the user to infer some of the information about the system"; and a "Rehearsal Manual" based on the "Inferential Manual," but which, in addition, encourages the user to "think again about [the] same prior information." (It is important to note that the experimenters apparently did not ask the subjects to "practice" what they had learned, but to "think" about what they had learned—to mentally "rehearse," in other words.) The "Inferential Manual" embodies the notion of guided exploration in that it was designed to encourage the user to infer information about the system. The "Skeletal Manual," being both terse and explicit, we assume, resembles a quick reference guide (samples of the experimental materials are not provided in the article).

Before discussing the results of the authors' study, it should be noted that in each of the experimental conditions except the "Lengthy Manual" condition, the experimental manuals were about the same length. As a consequence, despite the intent of these studies to isolate the factors involved in minimal manual design, it is impossible to say, for example, whether superior performance of users of the Inferential Manual to the users of the Lengthy Manual is due to the "guided exploration" component of the Inferential Manual or simply to its comparative brevity. In fact, in each of the experiments conducted, subjects assigned to learn specified tasks using the Lengthy Manual performed much more poorly than subjects assigned to use the other manuals, a fact that would certainly imply that the effects of brevity are influencing performance in all of the other conditions. A balanced experimental design capable of parceling out the effects of brevity would have investigated the effects of manual length across all pedagogical conditions. Additionally, a balanced design would have revealed any potential interaction effects between manual length and pedagogical approach.

This methodological limitation, of course, does not preclude a comparison of the performance of subjects assigned to the other three conditions, whose Inferential, Skeletal, and Rehearsal manuals were of approximately the same length. Black et al. compared users of these manuals on several different measures: (1) the amount of time to learn simple (one-command) tasks, (3) the amount of time required to learn "Command Sequence Tasks" (tasks requiring the use of several commands, and (4) the amount of time required to learn "Realistic Tasks" (creating and revising a letter or a memo).

In all but learning simple tasks, the users of the Inferential Manual outperformed subjects using the other manuals. However, as the authors note, the Skeletal Manual yielded surprisingly good performance—performance not significantly different from the Inferential Manual for learning both simple commands and command sequences. To explain these unexpected results, the authors suggest the possibility—a possibility with which it seems reasonable to concur—that subjects using the Skeletal Manual actually did some inferencing.

But an even simpler explanation of Black et al.'s results seems possible. While subjects in an experiment such as the one reported may have achieved a certain level of competence in the execution of a set of procedures, they have not been afforded the opportunity to truly acquire expertise. It is quite likely, in other words, that the procedural knowledge employed at this early stage of performance is uncompiled. Moreover, a simple, one-key command can't be compiled. Quite simply, there is nothing to compile. Subjects in all conditions, then, would be guided in the execution of these procedures by their declarative knowledge. As a consequence, it would be reasonable to expect that (1) subjects who were compelled to acquire that declarative knowledge inductively—to, in other words, invest a good deal of effort in the acquisition of that knowledge—would outperform those subjects who were not compelled to invest that degree of effort in their attempts to learn, and (2) that given the relative simplicity of the tasks on whose performance the subjects were measured, there would be little difference in performance between those who acquired that knowledge directly and those who acquired that knowledge through inferencing.

If we wish to know how we can most effectively foster expert performance, the important question to be answered is which form of manual is most effective in supporting proceduralization. To test whether proceduralization has taken place, it would be necessary to provide subjects with a distractor task to be performed concurrently with the desired task. Compiled procedures can be performed automatically. If it is possible for the subject to perform a set of commands while counting backward by threes, for example, then we could assume that that set of procedures had been automatized. If the subject couldn't do both tasks simultaneously, the implication would be that the procedures are really being supported by declarative rather than procedural knowledge—that the conversion of declarative to procedural knowledge has not yet taken place.

Nevertheless, alternative experimental design possibilities notwithstanding, the Black et al. study, we believe, offers no compelling evidence that guided exploration is a clearly superior instructional technology in the realm of computer documentation.

Two other studies have addressed the notions of minimalism less directly, but are nonetheless useful in informing our assessment of the efficacy of guided exploration in computer documentation. It is important to point out that in neither of these studies was it the authors' intent to assess directly the efficacy of guided exploration over alternative instructional technologies. Consequently, when we point out methodological problems, we do not mean that the methodologies used by the authors are flawed, but alternatively that those methodological choices have limited our ability to apply the findings of these studies directly to the issue of guided exploration.

Charney, Reder, and Wells (1988) addressed obliquely the topic of discovery learning within the context of a series of experiments they conducted that focused on the role "elaboration" provided by instructional materials plays in fostering subsequent learner efficiency in the application of the knowledge learned.

In a first experiment aimed at assessing the efficacy of textual elaborations in learning and remembering facts, the authors varied both the kind of manual (elaborated or unelaborated) and the subjects' task orientation (subjects either had a specific task to perform or were not provided with a specific task) and found that subjects who had a real task as well as a short manual that nonetheless provided complete and explicit instructions outperformed subjects assigned to the other groups. The authors conclude (p. 63) that "readers with specific tasks in mind need little or no elaboration in the text." To us, this experimental condition would seem to mirror a situation in which a user needs to employ a manual in an effort to accomplish real work, and the fact that readers do, indeed have real work that they want to accomplish is a decidedly minimalist assumption.

In a second experiment, the authors sought to determine what kind of elaborations would be useful for the learner who was attempting to acquire a skill. The authors provided subjects with both conceptual and procedural information and varied the extent to which each was elaborated (rich conceptual and rich procedural, rich conceptual and sparse procedural, etc.). The results of this experiment suggest that learners benefited little from elaborations of general concepts, but benefited significantly "from elaborations on how to apply procedures." Interestingly, however, the authors cite Kieras (1985) who found that "how it works" information (as opposed to "how to do it" information") is useful if the learner must infer operating procedures. The logical conclusion one must draw from this as it relates to minimalism, then, is that in order to support guided exploration, minimalist manuals would benefit from greater conceptual elaboration— precisely the kind of "extraneous" text that minimalism seeks to reduce in its manuals.

In a final experiment the authors' focus was to examine the effects of different kinds of learner interaction with learning materials. Because we believe that the conclusions the authors have drawn on the basis of this study might be inappropriately applied to the issue of discovery learning, we wish to discuss this study in greater depth than we have the two previous studies.

Participants in this study were assigned to one of three conditions, each differing in terms of the instructional materials provided. A third of the participants were given a "Tutorial" that provided them with the specific keystrokes and keystroke procedures to be used to solve a problem. Another third of the subjects were provided an "Exercise" that included a specific goal, but no instructions as to how to achieve the goal. Finally, the remaining subjects were given materials ("Tutorial Plus Exercise") that presented them with three "Tutorial" activities and two "Exercise" activities. Overall performance, assessed two days following training, was superior for subjects in the "Exercise" and "Tutorial Plus Exercise" conditions. The authors conclude on the basis of these results that "...active learning situations in which people remember and apply procedures for themselves are more effective than situations in which people simply learn by studying the procedures."

While at first blush these results might be interpreted as support for guided exploration, two characteristics of the study design militate against that conclusion. First, subjects provided the "Tutorial," unlike those in the other two conditions, were provided with no specific goals. This was not an oversight by the researchers, who argued that this kind of learning environment is typical of that provided by most commercial tutorials. If we are to accept the minimalist claim, however, that a significant problem with current documentation products is their inability to allow the user to pursue his or her goal to do real work immediately, then we must accept the corollary—namely, that our learners do, indeed, have real goals. Absent real goals, the learner has neither much motivation to learn nor a well defined problem space that would provide a meaningful context within which learning could take place. This situation would more probably than not encourage the learner to employ rote learning strategies in his or her attempts to learn, and rote learning strategies are seldom as efficient or as effective as are those strategies employed in learning meaningful material (Ausubel, 1968).

While the authors' claim that the "Tutorial" *materials* they provided to their subjects mirror typical computer tutorial materials may well be true, the inference that, consequently, the learning environment created by the "Tutorial" mirrors a typical "real world" learning environment seems debatable.

A second characteristic of the authors' research design also precludes a direct comparison of step-by-step procedures with discovery learning, namely, that the subjects assigned to the "Exercise" condition, following their own attempts to solve a problem, were allowed to study the experimenter's solution to the same problem. Absent this step, the "Exercise" condition would have constituted an attempt at discovery learning. Again, it should be emphasized that the experimenters were not attempting to create a discovery learning condition in this experiment, so this observation is not intended as a criticism of their experimental design. The point we wish to make is simply that the superior performance of those subjects assigned to the "Exercise" condition might be attributable to the fact that this group (1) had specific goals, (2) were compelled to be actively engaged intellectually in the attempt to perform a set of procedures, and (3) were also provided, prior to assessment of their learning, with those correct procedures. The results of this study cannot, consequently, be interpreted as support for the efficacy of guided exploration.

Again, what we can conclude from the experiments in this study is that (1) readers with specific goals or tasks to accomplish seem to need little elaboration in the documentation, (2) that when elaborations are helpful, they tend to be elaborations on how to apply procedures, and (3) that the existence of a learning environment that encourages the meaningful application of new knowledge bolsters that learning.

Vanderlinden, Cocklin, and McKita (1988) reported the results of an experiment that compared the efficacy of a traditional self-study tutorial with a guided exploration tutorial. Importantly, however, the guided exploration manual contained explicit procedures (GE+P manual). Subjects provided the GE+P manual outperformed those given the self study manual in terms of the number of subsequent problems they encountered completing tutorial tasks, and they completed post-tutorial tests more quickly and more accurately than those given the self study manual.

While conceptually the combination of guided exploration and procedures seems contradictory, in effect this documentation encouraged exploration while providing explicit procedural information on how to accomplish specific tasks. While the results of this experiment, then, do not address differences between guided exploration and traditional procedural manuals, they do suggest the efficacy of the provision of procedures in an environment in which the user is actively engaged in attempts to solve a problem on a computer.

SUMMARY AND CONCLUSIONS

We would like to attempt to summarize, and perhaps expand on, our views of guided exploration by making the following observations.

Despite our attempts to question the advisability of compelling a novice user to inductively arrive at a set of procedures, we do not argue that guided exploration or discovery learning is a barren educational technology. As the studies we have reviewed do suggest, guided exploration involves the learner actively in the learning enterprise and encourages the user to martial the kind of cognitive effort necessary to effectively learn new material. Moreover, implicit in the kinds of guided exploration activities we have seen is the goal of figuring out a set of procedures. Guided exploration, then, can be a directed, goal oriented, meaningful learning activity. For the novice user who has no real work to accomplish, the simple fact that guided exploration provides him or her with a goal and compels the user's active involvement in the learning process surely makes it a preferable learning environment to one in which the user is provided neither with meaningful learning objectives nor with the kinds of learning materials that encourage the learner to adopt anything more than rote learning strategies.

In fact, we would concede (as we, indeed, have in a previous section) that even in learning environments more ecologically valid than those typically created by experimental researchers, computer users spontaneously engage in discovery learning quite frequently. Users of a product will often explore if they believe that exploration will provide a more efficient solution path than will resorting to the procedural details of a manual. Users are always free to skip and skim and jump and, in effect, create their own guided exploration manuals. Any decent procedural documentation becomes guided exploration documentation when the user can find the sections he or she wants quickly, especially when the conceptual information is visually distinct from the procedures.

Whether such exploratory learning succeeds depends on the interplay between both system and user characteristics.

Stated in the simplest terms, exploratory learning is most likely to work when a user finds the software to be reasonably easy (and perhaps enjoyable) to figure out.

Exploratory learning may be more feasible when the software is "small," that is limited in the number of functions. When the number of functions is large, users may not only have extra difficulty comprehending functions or even how they fit together but also may require documentation simply as a supplement to long-term memory. Minimally, they will need to take their own notes.

More important than size is the difficulty of the product as measured by the familiarity and difficulty of the domain concepts for a particular user or class of users. Music composition software, for example, requires considerable domain knowledge and embodies some abstract and subtle concepts. Another very important factor is the quality of the functional interface and the advisory interface, in both their paper and online forms. Assuming identical functionality, the functional interface of many software products is self-disclosing and consistent, while in others cases it is not so and poses major problems for the user. Also, if the user occasionally attempts to consult the print or online documentation as part of an exploratory learning process, and these attempts are unsuccessful, the exploratory process will be seriously impeded. Another factor is the success of the software in motivating certain classes of users or the extrinsic motivation that certain classes of users bring to the process of learning certain software. The best case is when users not only want to learn the software but regard the learning process as a challenge or game.

Unfortunately, these conditions are very often not met. Not only are software products increasingly elaborate, but they tend to embody ever more difficult concepts. Newer word processing products, for example, often offer style sheets, a subtle and non-intuitive feature that gives many users trouble. Also, as people learn ever more applications, they tend to move into less familiar domains. Interfaces, both functional and advisory, still differ greatly in quality, and surely the beneficial effects of motivation can't be counted on; in particular, many users, especially in the midst of a hectic work day, do not regard computer learning as an enjoyable, game-like experience. It is important, then, to consider what happens when these conditions are not met.

The imposition of a guided exploration environment for learning a new system or application is, we believe, counterproductive to the learner who wants to accomplish real work. Implicit in that statement, of course, is the assumption that such a learner comes to the learning task with specific, articulable goals and only needs the documentation to provide the means to achieve those goals.

Compelling a user who has real work to do to figure out through trial and error the means for achieving those goals is at best inefficient, and quite probably would have the effect of frustrating the learner's efforts and lessening his or her motivation to continue to learn. Moreover, guided exploration takes out of the hands of the user very important decisions about what is important to learn and what is not important to learn. Other than, perhaps, some basics that all learners must know to effectively use a program, decisions about what procedures must be learned should be left to the learner, who will make those decisions on the basis of assessments of what must be known in order to do what must be done.

If we accept the fact that "expertise" is mediated by the acquisition of procedural knowledge, then our documentation should be constructed so as to encourage proceduralization. As noted previously, simple, straightforward procedural steps are most likely to be best for this purpose. Frequent use of the same procedures will result in decreased necessity of the learner to refer to the declarative knowledge prompts in order to execute the appropriate commands.

In conclusion, we agree with Carroll that manuals should be succinct, that they should focus on real tasks and activities, and that they should support error recognition and recovery. We find his arguments on behalf of guided exploration far less compelling, however, and can find little in the empirical research literature to support his claims for its efficacy.

REFERENCES

Anderson, J. R. (1976). *Language, Memory, and Thought.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

Anderson, J. R. (1983). *The Architecture of Cognition.* Cambridge, Mass.: Harvard University Press.

Ausubel, D. P. (1968). *Educational Psychology: A Cognitive View.* New York: Holt, Rinehart & Winston, Inc.

Black, J. B., J. M. Carroll, & S. M. McGuigan. (1987). What Kind of Minimal Instruction Manual is the Most Effective. In J. M. Carroll & P. P. Tanner (Eds.), *Human Factors in Computing Systems and Graphics Interface,* CHI + GI Proceedings, 159-162.

Brockmann, R. J. (1990). *Writing Better Computer User Documentation: From Paper to Hypertext.* New York: John Wiley & Sons, Inc.

Bruner, J. S. (1966). *Toward a Theory of Instruction.* Cambridge, MA.: Belknap Press.

Carroll, J. M. (1990). *The Nurnberg Funnel: Designing Minimalist Instruction for Practical Computer Skill.* Cambridge, MA.: MIT Press.

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

Charney, D. H., L. M. Reder, & G. W. Wells (1988). Studies of Elaboration in Instructional Texts. In S. Doheny-Farina (Ed.), Effective documentation: *What We Have Learned from Research.* Cambridge, MA: The MIT Press

Cohen, N. J., & L. R. Squire (1980). Preserved Learning and Retention of Pattern-analyzing Skill in Amnesia: Dissociation of Knowing How and Knowing That. *Science,* 210, 207-210.

Farkas, D. K., & T. R. Williams (1990). John Carroll's The Nurnberg Funnel and Minimalist Documentation. *IEEE Transactions on Professional Communication,* 33, 182-187.

Gagne, E. D. (1985). *The Cognitive Psychology of School Learning.* Boston: Little, Brown and Company.

Glaser, R. (1990). The Reemergence of Learning Theory Within Instructional Research. *American Psychologist, 45,* 29-39.

Hirsch, E. D., Jr. (1977). *The Philosophy of Composition.* Chicago: The University of Chicago Press.

Horton, W. K. (1990). Designing & Writing Online Documentation: Help Files to Hypertext. New York: John Wiley & Sons.

LaBerge, D. (1976). Perceptual Learning and Attention. In W.K. Estes (Ed.), *Handbook of Learning and Cognitive Processes, v. 4.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

Redisch, J. C. (1988). Reading to Learn to Do. *The Technical Writing Teacher, 15,* 223-233.

Rummelhart, D. E. & D. A. Norman (1978). Accretion, tuning and restructuring: Three modes of learning. In J.W. Cotton & R. Klatzky (Eds), *Semantic Factors in Cognition.* Hillsdale, NJ.: Lawrence Erlbaum Associates.

Stillings, N. A., M. H. Feinstein, J. L. Garfield, E. L. Rissland, D. A. Rosenbaum, S. E. Weisler, & L. Baker-Ward. (1987). *Cognitive Science: An Introduction.* Cambridge, MA: The MIT Press.

Schneider, W. & A. D. Fisk. (1982). Concurrent Automatic and Controlled Visual Search: Can Processing Occur without Resource Cost? *Journal of Experimental Psychology: Learning, Memory & Cognition.* 8, 261-278.

Vander Zanden, *J. W., &* A. *J.* Pace (1984). *Educational Psychology in Theory and Practice, second edition,* New York: Random House.

Vanderlinden, G., T. Cocklin, & M. McKita (1988). *Designing Tutorials That Help Users Learn Through Exploration, in On The Edge: A Pacific Rim Conference on Professional Technical Communication,* Proceedings of the IPCC, Institute of Electrical and Electronic Engineers, 295-299.

Weiner, B., R. Nierenberg, & M. Goldstein (1976). Social Learning (locus of control) versus Attributational (causal stability) Interpretations of Expectancy of Success. *Journal of Personality,* 44, 52-68.