ACTIVELY LEARNING TO USE A WORD PROCESSOR

John M. Carroll and Robert Mack
IBM Thomas Watson Research Center
Yorktown Heights, New York 10598

Learning to use a word processor provides a study of real complex human learning that is fundamentally "active", driven by the initiatives of the learner. People learn by actively trying things out, by reasoning, and by referring to prior knowledge. Our view is that these are natural -- albeit demanding -- strategies for people to adopt when confronted by a learning task of non-trivial complexity. What is especially noteworthy in the present case is that the learners we have studied are almost entirely innocent with respect to computer technology. In the context of learner innocence, we argue, these "natural" strategies entrain severe and wide ranging learning problems. Analysis of these problems, in turn, suggests research directions for the analysis of real human learning within Cognitive Science and practical directions in which computer word processing systems, and the educational technologies that support their training and use, might evolve.

In this research project, ten office temporaries spent four half-days learning to use one of two possible word processing systems in our laboratory. These people were highly experienced in routine office work, but quite naive with respect to computers in general and word processing in particular. We asked them to imagine a scenario in which a word processing system had recently been introduced to their office and they had been asked to be the first to learn it (to then pass this knowledge on to colleagues). The point was that they were to learn to use the system using the training materials that accompany it as their only resource.

Our method involved prompting learners to "think aloud" as they worked through the training materials. They were to report questions that were raised in their minds, plans and strategies they felt they might be considering or following out, and inferences and knowledge that might have been brought to awareness by on-going experiences. We remained with the learners, to keep them talking and to intervene if at any time it appeared that a problem was so grave that a learner might leave the experiment if we did not help out. Our prompting remained non-directive, and indeed once learners got going we needed to prompt very infrequently. Our analysis consisted first of an enumeration of "critical incidents", constrained by the consensus of the experimenters, which were cataloged and classified in various ways. The chief goal of this was to form a picture of the typical experience of a learner, and it is this induced "prototype" learning experience to which we will refer in what follows.

Learning by doing.

Our learners relentlessly wanted to learn by trying things out rather than by reading about how to do them. Half of our learners tried to sign on to the word processor before reading how to do so. In part this was impatience: they were reluctant to read a lot of explanation or get bogged down following meticulous directions. But it also devolved from mismatched goals: Learners wanted to discover how to do specific things at particular times, and this did not always accord with the sequence in which topics were treated in the manual.

Learning by trying things out according to a personal agenda of needs and goals is not merely a preference. Learners who try to follow out manual instructions are often unable to do so. The instruction sequences are fragile in the sense that it is easy to get side-tracked and there is no provision in them for recovery. One example is a learner who inadvertently paginated (reformatted) a document at the beginning of an exercise on revising documents. This not only rearranged the lines in the file to make right margins even, it also stored the document away. The learner had not yet learned how to retrieve documents and the manual itself provided no recovery information for this (or any other) type of error. Accordingly, she was forced to try to discover how to retrieve the document on her own.

Once the document was restored, she was faced with an equally staggering problem: the pagination operation had rearranged the lines of her file so that the revising instructions did not refer to the same document. An experienced user who understood reformatting could have reinterpreted the instructions and adapted them to this rearranged text. But this learner had no idea what she had done, and thus was puzzled by the fact that the instructions seemed to be wrong. The fragility of instruction sequences, coupled with the propensity of learners to try to recover by initiating exploratory forays, can result in problem tangles: Learners, who may not even fully understand the individual operations, have little basis for appreciating the subtle interdependence of clusters of word processor operations. They find themselves in distorted or even unrecognizable problem situations.

When learners do not, or cannot, follow directions the problems that arise can result in their losing track of what they are trying to do. It is likely, of course, that this loss of task orientation contributes to the overall failure of learning -- as indicated by the trouble all learners had applying their learning experiences to the routine typing "transfer task" after training. None of the learners were able to type, revise, and print a simple one page letter without some trouble with each of these basic skills.

What is more surprising perhaps is that even when learners were able to successfully follow instruction sequences out, they still seemed to experience a loss of task orientation, as evidenced by comments like: "What did we do?", "I know I did something, but I don't know what it is!" or "I'm getting confused because I'm not actually doing anything except following these directions." For these subjects, the overall orientation toward accomplishing meaningful tasks (e.g., type a letter, print something out) has been subverted by a narrower orientation toward following out a sequence of instructions.

Learning by thinking.

Just as learners take the initiative to try things on their own, so also are they active in trying to make sense of their experience with the word processor. Learning passively by rote assimilation of information is atypical. Rather, learners actively try to develop hypotheses about why it operates the way it does. These quests after meaning can be triggered by new and salient facts. They can be forced by discrepancies between what is expected and what actually happens. They can be structured by the learner's personal agenda of goals and queries, referred to as new problems arise. In each case, learners' lack of knowledge about word processing makes it difficult for them to reason out coherent solutions that accurately represent the objective operation of the system.

For example, learners have no basis for recognizing and ruling out irrelevant connections; their interpretations of word processing systems are often influenced by spurious connections between what they think they need and what they perceive. In one case, a learner tried to decide if a "File" command had stored a document file away. It was not stored because the command was entered in a text input mode where all typed strings are interpreted as text, and not executed as commands. But she assumed that the file had been stored, and adduced evidence to confirm this premise. For example, at one point she notices a status message "INPUT MODE 1 FILE" which indicates that she is in the text input mode. However, the word "file" matched her file command, and this was enough to suggest some kind of feedback that her "File" (as in store document) command had worked.

In such cases, reasoning appears to consist in adducing factual support to a premise the learner would like to hold as true. The learner above began with the hypothesis that she had stored the document file away, and sought evidence to confirm that this was the case. Her adduction here was incorrect because she did not know which facts were relevant to verifying the premise. In other cases, reasoning appears to consist in abducing a hypothesis when it, together with other assumptions the learner may already hold, is consistent with some fact or observation. One learner tried to move the cursor in a protected area of the display. When this locked the keyboard, she hypothesized that this fact meant that she was at the right place on the screen to do what she set out to do.

Learners also set goals which they actively pursue by trying to solve problems. They are hampered in this by their innocence of the appropriate problem space, or domain of possible actions and interpretations relevant to accomplishing goals and addressing queries. Accordingly, their strategies are often local and fragmentary; they have difficulty integrating information or other experiences, and in formulating their concerns in ways that map transparently onto system functions. When learners cannot solve problems or answer questions, they add them to a personal agenda of goals and queries as they go along. As new opportunities arise, learners return to these standing queries and try to resolve them.

Learning by knowing.

To this point, we have argued that a new user of a word processing system relies on active exploration and ad hoc reasoning as learning strategies. However, not all possibilities are explored and not all hypotheses that could be reached are reached. What constrains these strategies is a sense of what could be appropriate -- and this devolves from prior knowledge on the part of the learner: knowledge about devices "like" word processors (e.g., typewriters), knowledge about office routine and work in general, even knowledge culled from interacting with the word processor up to that point in time.

Our learners were unable to resist referring to their prior knowledge about typewriters as a basis for interpreting and predicting experience with word processors. One came to a halt as she read an instruction in the manual which said "Backspace to erase." It seemed that she could not interpret this instruction for, as she pointed out, BACKSPACE does not erase anything. She had irresistibly availed herself of her knowledge of how backspacing works on a typewriter, unable to even consider that this knowledge might be inappropriate for the present case. Other learners tried to use SPACE and RETURN keys to move the cursor -- which insert spaces and blank lines -- but merely move the typing point on a typewriter.

Our learners were experienced with conventional office work: typing letters, filing, etc. Their knowledge about how these routine tasks are organized in the office creates expectations in them about how analogous tasks ought to be performed in the "office of the future" (as represented by the word processor in our laboratory). Thus, one response to revising a letter task is to retype. This is striking since it is the capability of the word processor to store and retrieve documents -- for revision, among other things -- that is its fundamental advance over previous office technologies.

As a learning experience progresses, the learner is acquiring and organizing new bits of knowledge. The ultimate goal -- and the final measure of success in the learning situation -- is that of assembling these pieces into a coherent fabric, an understanding of the word processor. Along the way, any prior bit of knowledge is available for use as a basis for expectations concerning successive interactions with the system. One system we studied seemed to flaunt inconsistency in similar operations. Thus, to delete a word, one positions the cursor under the word's initial character and keypresses WORD DELETE. However, to underscore a word, one positions the cursor under the final character of a word and keypresses WORD UND. This inconsistency caused one learner to misexecute one and then the other of these two operations in a dismal cycle of negative transfer.

Summary.

Perhaps the most apt discussion of the world of the new user of a word processing system is that often quoted phrase of William James: "a bloomin' buzzin' confusion" People in this situation see many things going on, but they do not know which of these are relevant to their current concerns. Indeed, they do not know if their current concerns are the appropriate concerns for them to have. The learner reads something in the manual; sees something on the display; and must try to connect the two, to integrate, to interpret. It would be unsurprising to find that people in such a situation suffer conceptual -- or even physical -- paralysis. They have so little basis on which to act.

And yet people do act. Indeed, perhaps the most pervasive tendency we have observed is that people simply strike out into the unknown. If the rich and diverse sources of available information cannot be interpreted, then some of these will be ignored. If something can be interpreted (no matter how specious the basis for this interpretation), then it will be interpreted. Ad hoc theories are hastily assembled out these odds and ends of partially relevant and partially extraneous generalization. And these "theories" are used for further prediction. Whatever initial confusions get into such a process, it is easy to see that they are at the mercy of an at least partially negative feedback loop: things quite often get worse before they get better.

What's wrong? We would argue that the learning practices people adopt here are typical, and in many situations adaptive. The problem in this particular learning situation is that new learners of word processors are innocent in the extreme. "Word processor", so far as we know, is not a natural concept. People who do not know about word processors have little, possibly nothing, to refer to in trying to actively learn to use such things. Innocence turns reasonable learning strategies into learning problems.