

Writing as Problem Solving

John R. Hayes and Linda S. Flower

A top down approach employing protocol analysis can yield valuable data about writing processes. The main features of composition apparent in the data are: writing is goal directed, writing processes are hierarchically organized, some of the processes may interrupt others, recursion is possible, and writing goals may be modified as the result of writing. The first four of these features are embodied in a process model of composition.

It's very exciting to be doing research on composition just now because a great deal is happening. Theorists are exploring new ideas and experimenters are providing us with new empirical results at an unprecedented rate. Currently, there are at least four major theoretical viewpoints guiding the work of researchers who are trying to understand composition:

1. The psycholinguistic viewpoint, represented by Kintsch (1974), Rumelhart (1975), and others;
2. The linguistic viewpoint, represented by Young (1970), Cooper and Odell (1977), de Beaugrande (1979), and others;
3. The developmental viewpoint, represented by Bereiter, Scardamalia, and Bracewell (1979), Graves (1975), and others; and
4. The cognitive processing viewpoint, represented by Collins and Gentner (1979), Nold (in press), and Hayes and Flower (1980).

A casual observer could easily find this multiplicity of approaches confusing. Even researchers active in the field may sometimes find it difficult to characterize their own research. Their attention is likely to be directed at the subject matter they are trying to understand rather than at the assumptions underlying their research methods. In this paper we will stand back from the research we have been doing for the last several years and attempt to say what it is we think we have been doing.

There are two major factors which have shaped our work on writing. First, in conducting our research we have made a number of strategic decisions about what is interesting and about how best to proceed. To put it plainly, these decisions are the incarnation of our scientific biases. Second, our research has been shaped by a refractory world which has insisted that we attend to certain salient facts about the writing processes. In what follows we will describe the strategies which we have chosen to guide our research and then the facts which nature has imposed on us.

Visible Language, XIV 4, pp. 388-399.

Hayes's address: Communications Design Center, Carnegie-Mellon University, Pittsburgh, PA 15213
0022-2224/80/1000-0388\$02.00/0© 1980 Visible Language, Box 1972 CMA, Cleveland, OH 44106.

Strategic Decisions

Our approach proceeds from five strategic decisions which we made about how to conduct our research. Briefly these decisions were:

1. to focus on the act of writing;
2. to try for a process model of writing;
3. to model individual writers;
4. to work wholistically or "top down"; and
5. to divide the writing task into parts for easier analysis.

As we will see below, these decisions are genuine ones in the sense that we could reasonably have made other choices. Alternative approaches to the study of writing do proceed from different decisions on these same issues.

1. Our first and most important decision was to focus on the act of writing—that is, to attend to whatever it is that writers do when they produce a text. Thus, we viewed writing primarily as a process rather than as a product. We felt that by far the richest source of information about writing would be to observe step by step how the writer had actually created the essay. However, we did not intend to ignore the product. Wherever possible, we looked to the writer's essay for evidence to confirm or elaborate the more direct observations of process.

To observe writers in action we have employed process tracing methods borrowed from cognitive psychology. In our studies a typical experiment proceeds as follows: subjects appear at the experimental session knowing that they will be assigned a topic on which to write an essay and that the whole procedure will take about an hour. Further, they know that they will be asked to "think aloud" while writing. The subject is seated in a quiet office with a desk, pencil, and paper, and the tape recorder is turned on. The experimenter then gives the subject an envelope containing the writing assignment—that is, the topic and the intended audience. The subject then busily sets to work writing and commenting roughly as follows: "Well, open up the magic envelope. OK. Whew! This is a killer. Write about abortion pro and con for *Catholic Weekly*. Ok, boy! How am I going to handle this?", etc. This continues for about an hour until the subject says something like, "Well, that's it. Good bye, tape recorder (click)." The data of the study consist of a verbatim transcript of the tape recording (with all the "um's" and pauses and expletives undeleted) together with the essay and all of the notes the writer has generated along the way. The transcript is called a *protocol*. These materials are then examined in considerable detail for evidence which may reveal something of the processes by which the writer has created the essay. In general, the data are very rich in such evidence. Subjects typically give many hints about their plans and goals, e.g., "I'll just jot down ideas as they come to me"; about strategies for dealing with the audience, e.g., "I'll write this as if I were one of them"; about criteria for editing and evaluation, e.g., "For 10-year-olds, we better keep this simple"; and so on. The analysis of this data is called *protocol analysis*.

2. To understand the writing act, we certainly need to identify the processes involved—but this is not enough. We also need to know how these processes are organized to produce a text. That is, we need to know how the processes are sequenced, how one process is terminated and how the one which follows is chosen, how errors are detected, etc. Further, we want to know how simultaneous processes interact. When writers construct sentences, we want to know how they handle such multiple constraints as the requirement for correct grammar, appropriate tone, accuracy of meaning, and smooth transition. In short, we want a model which specifies the processes involved in writing and accurately describes their organization and interaction.

A model is a metaphor for a process: it's a way to describe something, such as the composing process, which refuses to sit still for a portrait. People build models in order to understand how a dynamic system works, and to describe the functional relationships among its parts. In addition, if a model is really to help us understand more, it should speak to some of the critical questions in the field of writing and rhetoric. It should help us see things in a way we didn't see them before.

Our second strategic decision was to direct our research toward the construction of such a model. Ideally, the model should be capable of telling us how writers go about producing a text when they are given a writing assignment. It should tell us what processes are involved, in what order they occur, and at what points the writer will experience difficulty. At present, of course, we must be satisfied with a model which is much less complete than the ideal. The ideal defines where we would like to go, but—alas!—not where we are now.

3. It is apparent that not all writers write in the same way. For example, some writers plan their essays from beginning to end before they write a single word of text, while others never seem to look beyond the next sentence. Further, some writers seem to write with their readers constantly in mind, checking frequently to be sure that they have taken the reader's knowledge and attitudes into account. Others appear serenely unaware that an audience could fail to understand what they, in good faith, have intended to say.

In modeling we can deal with such differences in either of two ways. We can choose to construct a model of the "average" writer and delay until some more propitious time the description of differences among writers. This approach has the merit of simplicity. Further, if things work out well, a model of an average writer might be useful in characterizing individual differences. Thus, models for individual writers might prove to be minor variants of the average model. However, this approach may have the disadvantage that averages sometimes suffer from—the average may be representative of no one. Thus, we sincerely hope that no one has the average number of children—two and a half—nor would we want anyone to have to eat an average course at dinner, which might be a compromise between appetizer and dessert such as oysters with chocolate sauce.

An alternative approach is to construct models which are intended to

describe individuals rather than averages of groups. The disadvantage of this approach is that it may be expensive. In the worst case, each individual may require a separate model. With better luck, models of individual writers will turn out to be variants of a small number of model types. The advantage of this approach is that it is more likely than a model of the average to capture the behavior of actual (rather than idealized) writers.

Our third strategic decision, then, was to model the behavior of individual writers rather than the average behavior of groups of writers.

4. In studying writing, we might well have started with processes which psychologists and psycholinguists have already identified as fundamental ones—processes such as short-term memory, grammatical categorization, and lexical marking. We might then have attempted to synthesize more complex processes using these fundamental processes as building blocks. This synthetic or *bottom up* approach is a very familiar one in science and has frequently been used with great success. Geometry and Newtonian physics are perhaps the best known examples.

However, research often proceeds in the opposite direction; that is, wholistically, or from the top down. Chemistry provides a good example of top down research. Chemical research often starts with a complex compound and then looks for the elementary components and their relations. The top down approach is the one we have chosen to apply in our writing research. We have started from the top with the complete writing act and have attempted to analyze it first into a few relatively complex subprocesses. As the analysis proceeds, the complex subprocesses are analyzed further into progressively simpler subprocesses. Ultimately, we hope that this top down analysis will make contact with the fundamental processes which psychologists and psycholinguists have already identified. Thus, the top down and bottom up approaches may be viewed as complementary.

The advantage of the bottom up approach is that it is rooted in fundamental processes. The advantage of the top down approach is that its results are almost certain to be relevant to real writing situations.

5. Our final strategic decision was to divide the writing task into three parts:

- A. The writer's long-term memory;
- B. The writing processes—that is, the writer excluding the writer's long-term memory; and
- C. The task environment—that is, the world outside the writer's skin. The relevant parts of the task environment are assumed to be: (1) The rhetorical situation—that is, the specifications of topic and audience to which a writer must respond; and (2) The text which the writer has produced so far. This text becomes an increasingly important part of the task environment as writing proceeds.

We chose this division because it is an especially convenient one for psychological analysis and modeling. Transfers of information between the task environment and the writer are usually marked clearly by overt acts of

reading or writing. Further, information retrieval from long-term memory is frequently detectable by examining the verbal protocol. Thus, the boundaries we have chosen divide the writing task into parts whose interactions are relatively easy to observe.

Bitzer's analysis of the rhetorical situation (1968) focuses on the importance of the task environment. Lowes' classic study of Coleridge (1927) focuses on the importance of the writer's long-term memory. Our own research has focused on the writing processes.

Our Model

While we don't want to present our model in great detail (that has been done elsewhere; see Hayes & Flower, 1980), we do want to show enough to illustrate how it had been shaped both by fact and by our strategic decisions.

Figure 1 shows the overall structure of the model. That we have a model of the writing act at all, of course, illustrates our first two strategic decisions: to focus on the writing act and to model it.

The effect of our fifth decision—to divide the writing task into task environment, long-term memory, and writing process—is also evident in Figure 1.

Figure 2 shows the subprocesses of the writing process. Figure 3 shows the inner structure of one of the subprocesses—the monitor. The progression from Figure 1 to Figure 3 parallels the progression of our top down approach. At first we analyze the writing act only into its largest, most evident components. Then, as research proceeds and as data allow, we analyze these major components successively into more refined subparts.

The function of the monitor (see Figure 3) is to control the sequence of writing processes. For example, it determines when idea generation will stop and organizing processes will begin. The middle section of the monitor—rules 3 through 6—is variable in form to allow for differences among writers in the way writing processes are sequenced. Figure 4 shows four alternative forms for rules 3 through 6. The first form represents a writer who polishes each sentence before considering the next. The fourth form represents a writer who plans the entire essay before writing the first sentence.

We know, of course, that there are many more differences among writers than the few we have discussed. Our strategic decision to model individual writers will lead us to search for other variations in the model which will allow us to describe other individual differences. We expect, for example, that some novice writers may entirely omit one or more of the major processes employed by competent writers. We have frequently observed writers who omit the review process—that is perhaps the commonest problem found in papers handed in at the last minute—and we have found at least one writer who showed no trace of an organizing process.

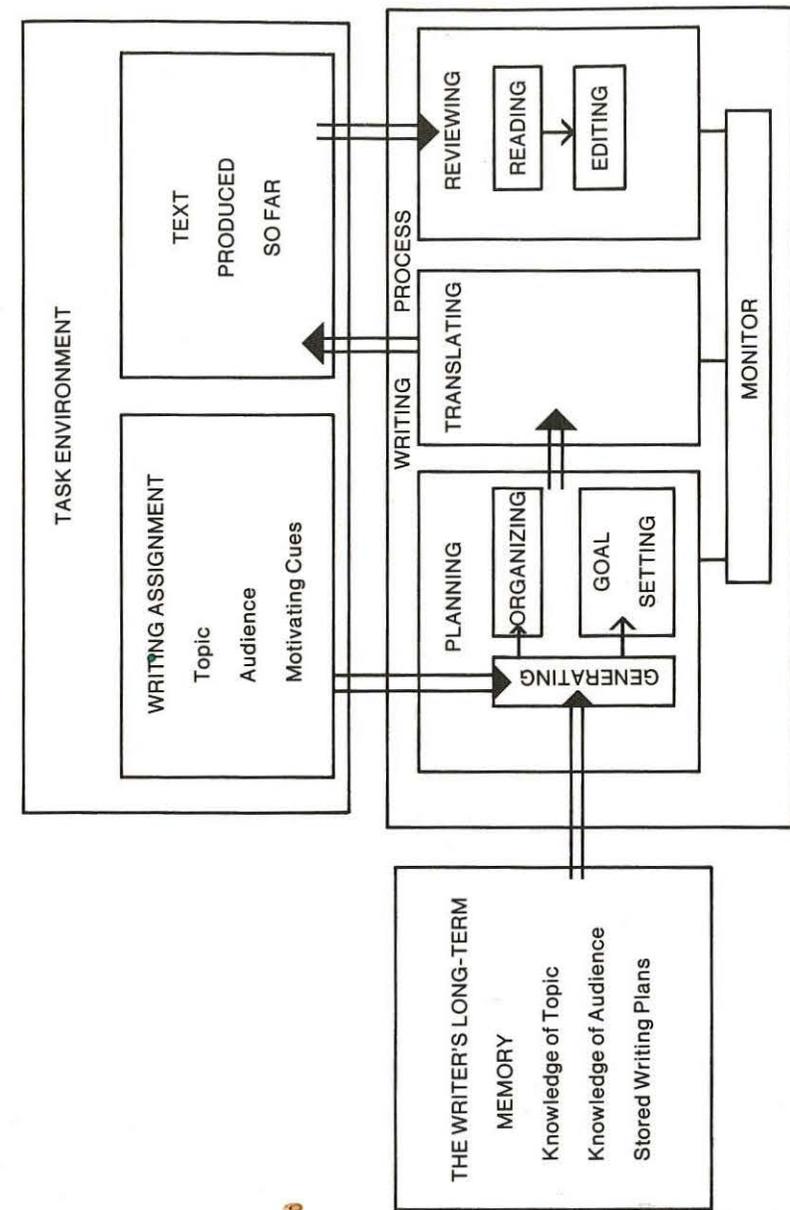


Figure 1. Structure of the writing model

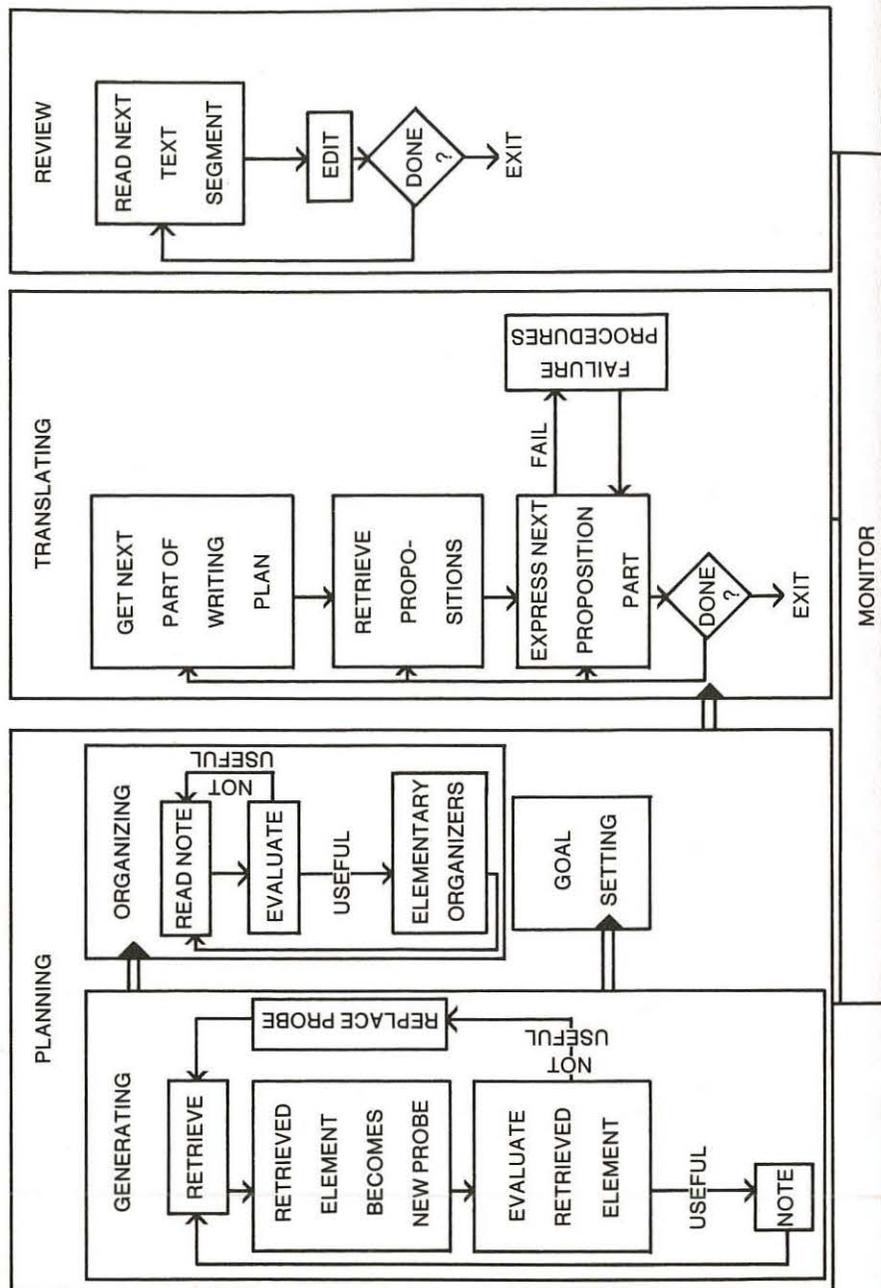


Figure 2. The writing process

1. (Generated language in STM → edit)
2. (New information in STM → generate)
- 3.-6. Goal setting productions
(These vary from writer to writer: see Fig. 4.)
7. [(goal = generate) → generate]
8. [(goal = organize) → organize]
9. [(goal = translate) → translate]
10. [(goal = review) → review]

Figure 3. Monitor.

Configuration 1 (Depth first)

3. [New element from translate → (goal = review)]
4. [New element from organize → (goal = translate)]
5. [New element from generate → (goal = organize)]
6. [Not enough material → (goal = generate)]

Configuration 2 (Get it down as you think of it, then review)

3. [New element from generate → (goal = organize)]
4. [New element from organize → (goal = translate)]
5. [Not enough material → (goal = generate)]
6. [Enough material → (goal = review)]

Configuration 3 (Perfect first draft)

3. [Not enough material → (goal = generate)]
4. [Enough material, plan not complete → (goal = organize)]
5. [New element from translate → (goal = review)]
6. [Plan complete → (goal = translate)]

Configuration 4 (Breadth first)

3. [Not enough material → (goal = generate)]
4. [Enough material, plan not complete → (goal = organize)]
5. [Plan complete → (goal = translate)]
6. [Translation complete → (goal = review)]

Figure 4. Alternate configuration for the monitor.

The Facts

Over several years, as we applied our research method to the analysis of writers in action, we were forced to a number of conclusions about the writing process. Among these conclusions, the most important in shaping our model of the writing process are these:

1. Writing is goal directed.
2. Writing processes are hierarchically organized.
3. Some writing processes may interrupt other processes over which they have priority.
4. Writing processes may be organized recursively.
5. Writing goals may be modified as writing proceeds.

1. *Writing is goal directed.* Evidence that writing is goal directed is easy to find in the protocols. Typically, writers comment on their major goals early in the writing session. For example, one writer who was asked to write about a woman's role for a hostile audience, said: "If an audience were hostile the worst thing to do would be to defend yourself—so I would try to humor them—to make them—uh—more sympathetic maybe. . . ." A second writer assigned this same topic said: "I'm trying to decide whether. . . I want to convince my audience of something specific about—uh—for instance the Equal Rights Amendment or whether something general about women should have the same rights as men. . . and I also need to decide if I want to actively convince my audience or simply state my point of view. . . ." A few lines later she decides: "I'll try to convince them of what it's like to—not to have certain rights. . . ." A third writer said, "I'm not really trying to persuade these people of anything, I'm simply being descriptive. . . . I'm saying this is the way the world is. . . ."

Goals enter into the model at several places. First, the goal-setting part of the planning process (see Figure 2) uses information from the task environment about the rhetorical situation and information from the writer's long-term memory about the topic and the audience to establish goals for the essay. Examples of such goals are seen in one writer's attempt to make her audience more sympathetic to a point of view and another writer's recognition that her essay for 10-year-olds must be made "very interesting."

Second, the monitor (see Figures 3 and 4) sets goals for carrying out writing processes. Such goals are reflected in statements such as, "Let's organize this mess," and, "OK, let's get it down on paper."

2. *Writing is hierarchically organized.* When writers have identified their major goals, e.g., the particular aspect of the topic they want to discuss and their general approach to the audience, they frequently identify subgoals on the route to these major goals. Indeed, the subgoals may in turn have their own subgoals. A writer who had as his main goal to write about the "worries" of a particular group set up subgoals to write about the subtopics "the political issue" and "the philosophical issue." Under each of the subtopics he specified a list of three or four sub-subtopics. Thus, his major goal was expanded into a hierarchy of subgoals.

In the same way, the writer who said that he was "simply being descriptive" elaborated his goal as follows: "I think what I really want is to present maybe one (point) with a lot of illustrations." He then went on to state the point and to develop a list of eight illustrations. In many cases, then, writers tell us in their "thinking aloud" protocols that their goals are hierarchically structured. Even if the writers didn't tell us explicitly though, there would still be plenty of evidence that writing processes are hierarchically organized. For example, many writers start the writing session with a period of planning in which they try to develop an outline to write from. To do this they may first try to generate ideas freely. When they feel they have enough ideas, they try to organize them into an outline. Generating and organizing are part of planning, and planning in turn is part of writing. Clearly these processes are hierarchically organized.

The model reflects the hierarchical organization of writing processes in two ways. First, the structure of processes in the model is intended to match the hierarchical structure of processes observed in writers. Second, operators within the ORGANIZE process allow the model to construct a hierarchical arrangement of goals.

3. *Priority interrupts.* Editing appears to take precedence over all other writing processes in the sense that editing may interrupt the other processes at any time. The generating process appears to be second in order of precedence since it interrupts any process except editing. Here are two examples of edits (in italics) which interrupt the writer while he is generating new ideas: "The problem is to make the uses more general and acceptable—*that's the wrong word—I mean important seeming*"; "Basically the idea is that if one has a special marker in a building that—*which*—means stop. . . ."

Interrupts by the editing process often appear quite abruptly. The writer no sooner has the wrong word out than the editing process leaps on it in the middle of a sentence and changes it. Interrupts by the generating process typically wait for the end of a sentence. However, they also appear rather abruptly. For example, while one writer was busily generating sentences according to his writing plan, he unexpectedly said, ". . . possibility of a pleasantry I suppose at that point. . . ."

The mechanism for priority interrupts in the model is located in the monitor (see Figure 3). The first two rules in the monitor control editing and idea generation. The fact that these two rules come before the goal setting rules (3-6) gives these processes priority over all others.

4. *Recursive processes.* The term "recursive" is used here in the mathematical sense. A recursive process is one which can contain itself as a part. Perhaps the best way to make this term clear is to start by discussing recursive definitions and to proceed by analogy to recursive processes.

A term is said to be defined recursively when its definition contains the term being defined as a part. At first glance, this situation may seem to involve an unacceptable circularity, but as we will see, everything works out all right. Let's take an example from linguistics. Consider the following (incomplete) definition of a sentence:

Sentence = noun phrase + predicate, or
sentence + "and" + sentence.

This is a perfectly workable definition even though the term being defined appears in the definition. Consider the problem of deciding whether or not the following string of words is a sentence:

"Frank has warts and Betty has hiccoughs."

Since the string doesn't have the form "noun phrase + predicate," we test to see if it has the form "sentence + 'and' + sentence." To do this, we have to show that the strings before the "and" and after the "and" are sentences. They turn out to be, since both have the form "noun phrase + predicate."

Just as the idea of a recursive definition contains no intractable circularity, neither does the idea that a recursive process can contain itself as a part contain any intractable circularity. To illustrate a recursive process, we will consider the Wendy protocol. In her first draft, Wendy wrote sentence 1 of the final draft and then followed it directly by sentence 7 of the final draft. When she was editing (a part of the writing process), Wendy decided that readers would have trouble with the transition between sentences 1 and 7. As a result she called on the whole writing process to insert a small essay inside her larger essay. The whole writing process then was used as part of editing, and thus as a part of itself.

In the model, the mechanism underlying recursion is hidden in the fine structure of the editing process. We believe that when the editing process identifies a major fault in the text, e.g., lack of context or poor organization, it may employ the whole writing process in the effort to fix the fault.

5. *Dynamic modification of goals.* When writers choose goals, they are by no means stuck with them. When a goal is difficult to meet, the writer can respond to the difficulty by modifying or abandoning the goal. Consider, for example, the writer who was assigned the task of writing about "abortion: pro and con, for a hostile audience." Early in the session, she decided to include as one of the "pro" topics the point that "a woman should have the right to limit the size of her family." About half-way through the session, she came back to this point and re-evaluated it as follows: "A woman should have the right to limit the size of her family—um—and I think I'm not going to mention that—because that is not a view that is shared by the audience."

The writer mentioned earlier who was considering "a pleasantry," explored the issue in some detail, decided it wouldn't work well, and abandoned it. This same writer had established goals at various times in the session to write about the issues "individual freedom," "the impact of technology," and "attitudes toward scientific analysis." After considerable work, he discovered relations among the three and tied them together under a single heading as "the philosophical issue." Thus, three independent objectives were modified to become subgoals of a single higher level goal.

As yet, the model has no mechanism to account for the modification of goals. We assume that the mechanism must reside in the goal setting processes and that it must use information about processing failures fed

back to it by the translating and organizing processes. As yet we don't have sufficient data to specify how the mechanism works.

References

- Bereiter, C., Scardamalia, M., & Bracewell, R. J. An applied cognitive-developmental approach to writing research. Paper presented at the American Educational Research Association Meeting, San Francisco, April, 1979.
- Bitzer, L. The rhetorical situation. *Philosophy and Rhetoric*, 1968, 1, 1-14.
- Collins, A., & Gentner, D. A framework for a cognitive theory of writing. Unpublished manuscript. Cambridge, Mass.: Bolt Beranek & Newman, 1979.
- Cooper, C. R., & Odell, L. Evaluating writing: describing, measuring, judging. Urbana, Ill.: National Council of Teachers of English, 1977.
- de Beaugrande, R. The processes of invention: association and recombination. *College Composition and Communication*, 1979, 30(3), 1-12.
- Graves, D. H. An examination of the writing processes of seven-year-old children. *Research in the Teaching of English*, 1975, 9(3), 227-241.
- Hayes, J. R., & Flower, L. S. Identifying the organization of writing processes. In L. W. Gregg and E. Steinberg (eds.) *Cognitive processes in writing*. Hillsdale, N.J.: Lawrence Erlbaum, 1980.
- Kintsch, W. *The representation of meaning in memory*. New York: Lawrence Erlbaum, 1974.
- Lowes, J. L. *The road to Xanadu*. Boston: Houghton Mifflin, 1927.
- Nold, E. Revising. In C. H. Frederiksen, M. S. Whiteman, and J. F. Dominic (eds.), *Writing: the nature, development, and teaching of written communication*. Hillsdale, N.J.: Lawrence Erlbaum, in press.
- Rumelhart, D. Notes on a schema for stories. In D. Bobrow and A. Collins (eds.), *Representation and understanding*. New York: Academic Press, 1975.
- Young, R. E., Becker, A. L., & Pike, K. E. *Rhetoric: discovery and change*. New York: Harcourt, Brace, and World, 1970.

Funding for this paper was supported in part by Grant No. NIE-G-78-0195 from the National Institute of Education, Washington, D. C.