

Text Comprehension, Memory, and Learning

Walter Kintsch

People are often able to reproduce a text quite well but are unable to use the information in the text for other purposes. Factors that help people to reproduce a text have been studied for some time. This article explores ways that enable people to learn from texts. Content overlap between a text and the reader's prior knowledge is identified as one factor, and methods are proposed to identify whether a text is suitable for readers with given background knowledge. For readers with low background knowledge, a text should be as coherent and explicit as possible to facilitate learning. However, data are presented to show that for readers with adequate background knowledge, texts with coherence gaps that stimulate constructive activities are in fact better for learning.

In list-learning experiments, there is no distinction between remembering a list and learning a list. A subject who is able to reproduce a list of words is said to remember them or, equivalently, to have learned the words. The situation is different with texts: Remembering a text and learning from it are by no means equivalent. Remembering a text means that one can reproduce it in some form, more or less verbatim and more or less completely, at least its gist. Learning from a text implies that one is able to use the information provided by the text in other ways, not just for reproduction. For instance, one can infer new facts from the information in the text, use it in conjunction with previous knowledge to solve novel problems, and integrate it with what was already known. A typical case of remembering a text would be a reader recounting a story to someone else—or a subject remembering in a free-recall experiment. A typical case of learning from a text would be a newspaper reader calling a stockbroker after reading a news item that seemed relevant to the reader's investments—or a subject answering inference questions after reading a text.

The distinction made here¹ between learning and memory is a matter of the criteria used to define *learning*: Learning requires deep understanding of the subject matter, so that the information acquired can be used productively in novel environments; for mere memory, as assessed by reproduction of the text, a more shallow understanding suffices. Normally reproduction of a text and real understanding are correlated, so that text memory becomes a prerequisite for learning, although that is not necessarily so. For example, Bransford, Barclay, and Franks (1972) have shown that subjects can acquire a good and useful mental model with no memory for the text itself, whereas others (e.g., Moravcsik & Kintsch, 1993) have found that it is sometimes possible to remem-

ber a text quite well without being able to use the information in it productively.

Situation Model and Textbase

The need to distinguish between text memory and learning arises from the impreciseness of the term *comprehension*. As in list learning, what is remembered or learned from a text depends on the nature of the encoding conditions and the retrieval cues present at the time of test. Encoding a text means comprehending it—but there are many ways in which a text can be comprehended, ranging from the most superficial to deep understanding. The common sense notion of *comprehension* is insufficient at this point, but it can be supplemented by theoretical analysis. Thus, in the text comprehension theory of van Dijk and Kintsch (1983; Kintsch, 1992), different levels of comprehension are distinguished. Comprehension always involves, although to different degrees, a surface component. That is, the words and phrases themselves are encoded as are the linguistic relations between them. The semantic and rhetorical structure of the text provides another set of relationships that are important in comprehension and that are frequently encoded. Van Dijk and Kintsch have called this level the *textbase* and have distinguished it from the *situation model*, which corresponds to a deeper level of understanding. In the situation model the information provided by the text is elaborated from prior knowledge and is integrated with it. Figure 1 shows two sentences, taken from a text on heart disease, with their textbase and situation model. The textbase consists of three propositions (the NOT-

Editor's note. Articles based on APA award addresses that appear in the *American Psychologist* are scholarly articles by distinguished contributors to the field. As such, they are given special consideration in the *American Psychologist's* editorial selection process.

This article was originally presented as part of a Distinguished Scientific Contributions award address at the 101st Annual Convention of the American Psychological Association in Toronto, Ontario, Canada, in August 1993.

Author's note. This research was supported by Grant 15872 from the National Institute of Mental Health and a grant from the Mellon Foundation.

Correspondence concerning this article should be addressed to Walter Kintsch, Institute of Cognitive Science, Muenzinger Psychology Building, University of Colorado, Boulder, CO 80309-0344.


¹ The distinction between learning from a text and simply remembering it is not always made. For example, much of school learning is commonly assessed by having students reproduce a text. On the other hand, in the problem-solving domain, an analogous distinction between rote memorization of procedures and problem solving with understanding has traditionally been made (most notably Wertheimer, 1945).

Figure 1
A Two-Sentence Text Fragment With Its Textbase and Situation Model


Text:

When a baby has a septal defect, the blood cannot get rid of enough carbon dioxide through the lungs. Therefore, it looks purple.

Textbase:



Situation Model:



GETRID proposition includes two modifiers), linked by sentence connectives. Of course, not every reader will construct exactly the same textbase: Some reader might, for instance, neglect the modifier ENOUGH. The situation model here is shown in the form of a graph. Most of the graph is based on the reader's knowledge about the circulatory system, rather than on the text directly. Only the fact that there is a gap in the septal wall so that purple blood gets mixed with the red blood is derived from the text itself. Once again, different readers do not necessarily construct the same situation model, especially inasmuch as the (correct) model is rather complex and requires a good understanding of the circulatory system that not every reader possesses.

In the model of Kintsch (1988, 1992), these three levels of analyses—the surface structure, the textbase, and the situation model—are represented as relations in a network of propositions. For that purpose, the schematic used to depict the situation model in Figure 1 has been translated into a set of propositions, as shown in Figure 2. The terms *textbase* and *situation model* are useful to describe the network shown in Figure 2, but functionally there exists only a single network to which both the textbase and situation model contribute links and nodes. Once this network has been constructed, activation is spread throughout it until the pattern of activation settles, with the effect that some of the nodes will become highly activated, whereas others become less strong or, if they are unconnected or connected with inhibitory links, are

actually rejected. In the network shown in Figure 2, the NOTGETRID and MIX nodes become most strongly activated, whereas the two CARRY nodes receive very little activation.


The distinction between memory for a text and learning from a text, in terms of the model sketched above, appears to be a matter of how complete and elaborate a situation model is constructed during comprehension. A text can be recalled even if only a textbase is constructed. It can even be summarized adequately on that basis, for example, by reproducing only the most highly activated propositions. We do not want to imply that someone who can recall and summarize a text does not understand it. However, this may be a very superficial level of understanding. Lacking any situation model, the comprehender knows that blood turns purple because it cannot get rid of carbon dioxide through the lungs—that is what the text states—but not why this is so. The text might as well have asserted that the blood cannot rid itself of carbon dioxide because it turns purple—it would make equally little sense to the reader. In order to understand why there is a “therefore” and not a “because” between these two sentences, the reader must have enough knowledge about the working of the circulatory system to form a situation model that fills in the information necessary for understanding that was not made explicit in the text. This situation model may be more sketchy or more elaborate than the one suggested in Figure 1 or it may be wrong, but something beyond the text itself must be there for a deeper understanding of this text.

Therefore, in writing a science text, our goal should be to allow the reader to create a structure akin to Figure 2 as a mental representation of the text. How is this best achieved?


Quite a bit is known about how to write texts so as to maximize the chances that a reader will construct a good textbase from it. For instance, we know about the role of local coherence (e.g., through common referents)

Figure 2
Constraint Satisfaction Network Corresponding to Figure 1

Textbase:




Situation Model:



Note. The situation model consists of an inference (the MIX proposition) and several previous knowledge nodes.

Figure 3

Proportion of Correct Answers on a Sentence Verification Test for Old Sentences Taken Directly From the Text and for Inferences



and the importance of signaling the macrostructure of the text (e.g., by summarizing the main point of a paragraph in the first sentence—Kintsch & Vipond, 1979; Weaver & Kintsch, 1990). Thus, if the goal of a science textbook writer were to provide easy readability and fluent recall of the textbook, we would be in a good position to give advice. However, that is not our goal. Our concern is not whether students recall the textbook but whether they learn from it (or, at least, this should be the case, although our tests frequently focus on the students' reproduction ability, rather than their learning). So what kind of advice can psychologists provide for textbook writers? Unfortunately, our recommendations here are much less specific, because the research on learning from texts is not nearly as advanced as that on text memory. However, some promising beginnings have been made, and it is at least becoming clear what the significant questions are.

An Advance Organizer That Facilitates Learning But Not Remembering

First of all, we have come to realize that what is good for remembering a text is not necessarily good for learning. To illustrate this point, consider a study by Mannes and Kintsch (1987). The study used the advance organizer paradigm. Subjects were given background knowledge about microbes to study before reading a technical article on the industrial uses of microbes. The background material contained general information about microbes, some of which was directly relevant to the target article and some of which was not. This background information was presented in two forms. In one case, the general in-

formation about microbes was presented in the same order as in the text. In the other case, the organization of the background material had no relation to the text at all and followed the order in which it was presented in the original encyclopedia article from which this material was adapted. The actual content of the two versions was, however, identical; only the organization of the material was different.

When the background material was organized in the same way as the target text, the subjects were better able to remember the target text than when the organization was different. This was true for both free recall and sentence verification. The left half of Figure 3 shows the subjects' percentage correct verification responses for old sentences. One could therefore conclude that students learn better when an advance organizer is structured in the same way as the to-be-learned text. However, exactly the opposite results were obtained for inference statements and a problem-solving task: The best results were obtained when the background material was presented in an order that differed from that of the target text. The right half of Figure 3 shows the percentage correct verification responses for inference statements.

We do not believe these findings to be just another instance of the inconsistent results that are so frequently observed in advance organizer studies but an indication of the importance of the distinction between remembering a text and learning from it. When subjects study the background material, they form a textbase-situation model representation, as in Figure 2. A major feature of that representation is its macrostructure, that is, its global organization into main points and subordinate points.

When subjects read the target text in the same order, its macrostructure fits very well the structure they already have formed for the background material, and hence the text is easy to comprehend, resulting in a well-organized textbase. At the same time, because everything fits so well into preexisting slots, there is little inducement for elaborating a complex situation model. As a result, text memory, which depends on the textbase, is very good, but inferencing and problem solving, which depend on the situation model, are not as good. On the other hand, when there is a discrepancy in the organization of the background material and the target text, readers cannot simply rely on the background material for the organization of the target text but have to form a new macrostructure. There will be interference from the old text structure, some confusion of what was in the background material and what was in the text proper, so that recall will not be as good, and more intrusion errors will occur. At the same time, the mental representation of the target text will be more richly interconnected with other parts of the background text. Thus, when a reader needs to use this information productively for inferencing and problem solving, relevant material is more likely to be accessed, and better performance will be observed.

The Mannes and Kintsch (1987) study has far-reaching implications: One cannot simply assume that text characteristics or reader strategies that are optimal for text memory are also best for learning from text. It is necessary, therefore, to investigate the factors that optimize learning in their own right. The comprehension theory can be a valuable guide in this enterprise in that it can help us to identify the right questions. These questions fall into two broad areas: How is learning affected by the content of texts and by the form of texts?

The Content of the Text

It is well established that background knowledge is important for text comprehension and memory (e.g., Voss, Fincher-Kiefer, Greene, & Post, 1986). The distinction between memory and learning that was introduced above is not usually made in this literature, although more complex behavioral measures show stronger domain knowledge effects than simpler measures, such as recognition (for a review see Schneider, Korkel, & Weinert, 1989), in agreement with that distinction. Sometimes people successfully comprehend texts for which they have no or minimal specific background knowledge, but to do so they need help in the form of well-written texts that allow them to substitute general comprehension strategies for the knowledge they lack (Moravcsik & Kintsch, 1993). In general, however, there can be no doubt that knowledge facilitates text processing.

It is a well-known, general psychological phenomenon that what can be acquired is limited by the current state of the learner. This point has often been made, for example, in developmental psychology (Brown & Reeves, 1987) and skill acquisition (Burton, Brown, & Fischer, 1984). Vygotsky (1934/1986) introduced the concept of proximal zones of development: areas at the borders of

what is already known where future growth or learning can take place successfully. The zone concept may also be useful for understanding learning from texts. Each reader may be characterized by certain limited domain knowledge and certain skills that define a set of texts that can serve as the basis for successful learning. Texts too close to the reader's knowledge are redundant, and texts too far removed are too difficult. Even if they are understood, they may not serve as a basis for successful learning. Of course, how hard a reader tries to achieve a deeper understanding of text is a factor that is not to be neglected, but how much the text overlaps with prior knowledge appears to be a major determinant in learning from texts. Our goal, then, is to specify these zones of learning for particular readers or groups of readers.

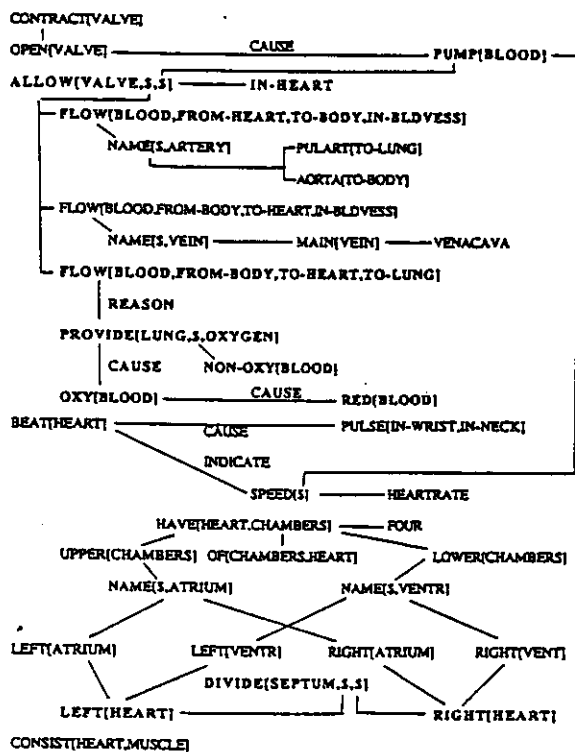
A Simulation of Comprehension for High- and Low-Knowledge Readers

As a starting point, we can make use of the text comprehension model of Kintsch (1988, 1992) to analyze how domain knowledge is involved in learning from a text. To illustrate this procedure, we shall simulate how high- and low-knowledge readers would comprehend a paragraph from a science text about heart disease. Subjects in an experiment conducted by McNamara, Kintsch, Butler-Songer, and Kintsch (1993) were given a pretest designed to assess their knowledge about the functioning of the heart. The correct answers from that pretest can be used to estimate the subject's knowledge about the heart. Obviously, such a procedure will underestimate the subject's knowledge because the subject may have been able to answer questions about the heart that were never asked, but it provides a first approximation. Figure 4 shows knowledge maps for two subjects, one high- and one low-knowledge subject. The figures are based directly on the questions these subjects answered correctly (rephrased as statements). Links are drawn between propositions and their modifiers (e.g., HAVE[HEART, CHAMBERS]—FOUR), whenever a proposition is embedded as an argument of another proposition (e.g., DIVIDE[SEPTUM,\$,\$] is linked to the two propositions embedded in it), as well as for specific semantic relations, like CAUSE.


During comprehension we assume that elements in the paragraph may retrieve elements of the knowledge net. This retrieval process can be described by models of memory retrieval (e.g., the SAM model of Raaijmakers & Shiffrin, 1981). Thus, retrieval cues in the text, including compound retrieval cues, may retrieve associated knowledge elements according to the retrieval mechanisms specified in the SAM model. For our purposes, completely successful retrieval was assumed, so that whenever a text element was associated with a knowledge element, the knowledge element was included in the network. Thus, the phrase "the heart supplies blood to the body" retrieved PUMP[HEART,BLOOD] for the low-knowledge subject and FLOW[BLOOD,FROM-HEART,TO-BODY] for the high-knowledge subject; VALVE retrieved the AL-LOW[VALVE,\$,\$] proposition for the high-knowledge

Figure 4
Knowledge Maps for a High-Knowledge and a Low-Knowledge Subject

Knowledge Map for High-Knowledge Subject:



Knowledge Map for Low-Knowledge Subject:

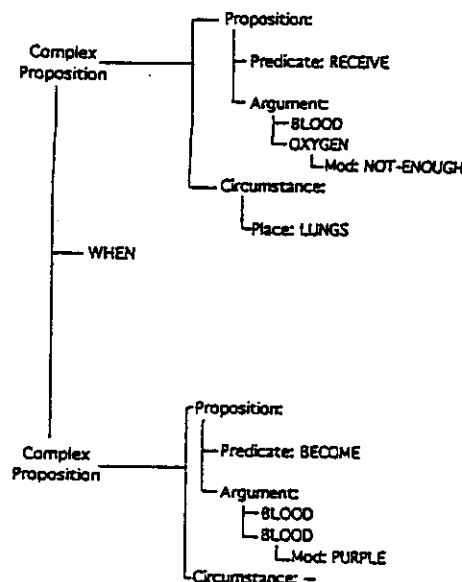


proposition consists of a predicate and arguments characterizing the internal semantic structure of a sentence. Each of these components may be modified, and modifiers may be further modified themselves, as indicated by the syntactic structure of the sentence to be represented (e.g., embedded clauses, relative clauses, adjectivization). Propositions in a discourse may either be indirectly related because they share a circumstance (e.g., time, place) or argument, or they may be directly related, either temporally, conditionally, or causally by an explicit sentence connective.


Comprehension of the texts was simulated according to the Kintsch (1988, 1992) model. In this model, text is processed in cycles roughly corresponding to a sentence. Parameter estimates were obtained through an informal exploration of the parameter space and were strongly constrained by previous results. Working memory was set to 5 elements for the model. Hence, the model reads the first sentence, goes on to the next sentence, and continues until it has read at least 5 elements. The model then retrieves from its knowledge base whatever information is related to the text and adds these elements to the network. For the text as a whole, 12 elements (shown

Figure 5
Propositional and Network Representation of the Sentence "When the blood receives not enough oxygen in the lungs, it becomes purple"

Propositional Representation:



Network Representation:



subject, and so on. In this way, 12 knowledge nodes were included in the situation model for the high-knowledge subject but only one for the low-knowledge subject. These knowledge nodes were added to the text propositions in each cycle, according to the assumption that this kind of knowledge access through long-term working memory (Ericsson & Kintsch, 1991) is not resource and capacity consuming.

The details of the simulation can only be sketched here. The experimental text was propositionalized according to van Dijk and Kintsch (1983, chapter 4). Text propositions in the sense of van Dijk and Kintsch are schematic representations of the meaning of sentences. This schema, of which an example is shown in Figure 5, comprises a proposition slot and a circumstance slot, specifying the time and place of the action or state. The

in boldface in Figure 4) were added in this way to the text for the high-knowledge reader but only one for the low-knowledge reader. Otherwise, the simulation was performed in the same way as described in Kintsch (1992).


Separate simulations were performed for high- and low-knowledge readers. On the whole, there were no major differences between high- and low-knowledge subjects for text elements. The average activation values per text element were almost equal for the two cases. On the other hand, as an obvious consequence of the fact that 12 knowledge elements were included in the simulation for the high-knowledge subject and only 1 element for the low-knowledge subject, the knowledge activation was of an order of magnitude greater for the high-knowledge reader.

Thus, if one considers only the activation of the textbase, background knowledge in this case had no strong effects. However, if one includes the whole situation model—textbase plus activated knowledge—a very different picture emerges. In particular, for the high-knowledge reader, the episodic text memory that has been constructed is solidly connected with the prior knowledge base, whereas the (much smaller) knowledge base of the low-knowledge reader has only a single link with the episodic text structure.

The consequences of this fact for retrievability are noteworthy. For the low-knowledge reader, the episodic text memory is retrievable only through episodic context cues or the single link to "the heart pumps blood." Although the low-knowledge reader remembers the text approximately as well as the high-knowledge reader, in new situations where the context retrieval cue is absent, this knowledge remains largely unavailable. The low-knowledge reader has learned very little. In contrast, the high-knowledge reader has several effective retrieval routes even when the contextual retrieval cue is absent: All of the nodes in the knowledge net that are printed in bold face in the knowledge map for the high-knowledge subject are linked to the text structure, and hence that structure is no longer inert knowledge but is retrievable and usable in novel contexts. It has become part of the reader's knowledge, whereas for the low-knowledge reader it remains largely a separate structure that cannot be accessed in new contexts requiring information about septal defects. To take a concrete instance, compare the retrieval process for the high-knowledge and low-knowledge subject in response to the question, "How does a septal defect affect the flow of blood within the heart?" The high-knowledge subject's more detailed understanding of the structure and functioning of the heart enables him or her to infer that oxygenated and nonoxygenated blood would become mixed through the hole in the septum. The low-knowledge subject cannot make this inference, even though he or she remembers the definition of a septal defect from the text.

The model predicts, therefore, that high- and low-knowledge subjects should be equally able to reproduce the text but that high-knowledge readers should be better

Figure 6
Percentage Correct in Free-Recall and Inference Questions and Sorting Scores After Reading a Text on Heart Diseases for High- and Low-Knowledge Subjects



at reconstructing and elaborating the text. Overall, in a free-recall task, one would expect high-knowledge readers to do somewhat better than low-knowledge readers. However, more pronounced differences between high- and low-knowledge readers should be observable in tasks that depend more strongly on the reader's situation model than on free recall, for example, in answering questions about the text, especially when they involve inferences.

Experimental Evidence

These claims of the model have been verified empirically in a recent study by McNamara et al. (1993). In this study a 683-word text on heart disease (a portion of which was used as for illustration above) was given to 6th- and 8th-grade students who on the basis of a pretest were divided into high- and low-knowledge groups.² Free recall and answers to text-based questions are both performance measures that reflect a mixture between textbase (reproductive recall proper) and situation model (the reconstructive component of recall). The results from these tasks are shown in Figure 6a. High-knowledge readers performed better on both tasks. The textbase was presumably comparable for both groups, but the better situation model of the high-knowledge readers allowed them to reconstruct the text more successfully. Thus, neither free recall nor the text question answering task gives an undistorted picture of the subjects' textbase or situation model because both tasks reflect a complex mixture of the two levels of representation. McNamara et al., therefore, used answers to inference questions and the results of a key word sorting task to assess the readers' situation

² The McNamara et al. (1993) experiment actually contained several additional conditions that will not be discussed here.

model more directly. These results are shown in Figure 6b. Whereas high-knowledge subjects were only approximately 50% better than low-knowledge subjects on tasks that reflect primarily the textbase (Figure 6a), they held an approximate 200% advantage on tasks for which the situation model plays a more significant role (Figure 6b).

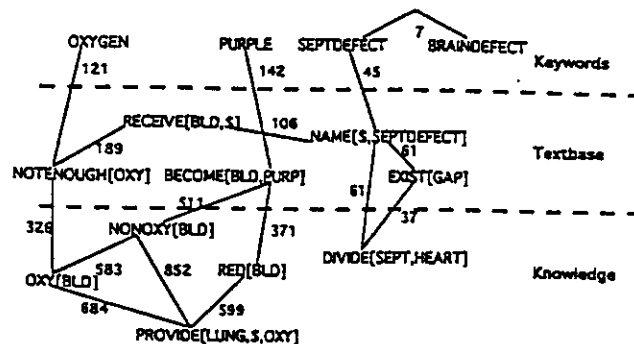
A Simulation of the Sorting Task

Just how the sorting task relates to the situation model, as well as to the textbase, needs to be analyzed in some detail. For this task the students were given 16 key words to sort as they liked, both before and after reading the text. Before reading, the manner in which the cards are sorted depends on the subjects' knowledge structure and their sorting strategies. After reading, the episodic text memory may also influence the sorting, to the extent that it is integrated with the subjects' knowledge. Concepts linked by the text should more likely be sorted together.

Consider four hypothetical key words: PURPLE, OXYGEN, SEPTAL DEFECT, and BRAIN DEFECT. Before reading, there is no particular reason to expect that these key words would be sorted in any systematic fashion, except for the obvious association between SEPTAL DEFECT and BRAIN DEFECT. The situation for high-knowledge subjects after reading is depicted in Figure 7.

Figure 7 shows the four key words, related text propositions, and related knowledge elements for the high-knowledge subject. For the low-knowledge subject, these knowledge elements would be missing. The strength of the links among the four key words after reading can be computed for high- and low-knowledge subjects by including or excluding the knowledge nodes in the network. It is obvious from the figure that for high-knowledge subjects paths exist among all four key words. It is not clear how to compute the strength of a path traversing several nodes, but one possible way would be to set the strength of a multiple-node path as equal to its weakest link.³ In this case, Figure 7 shows that there is a strong pathway of interconnecting links between the key words PURPLE and OXYGEN through the knowledge network. Weaker links of equal strength exist between both of these key words and SEPTAL DEFECT, and a very weak link connects SEPTAL DEFECT and BRAIN DEFECT. Thus, different sorting results would be expected after reading than before reading for high-knowledge subjects. For low-knowledge subjects, somewhat different results are expected. As is obvious from the network shown, the key word PURPLE is unrelated to the rest of the net without the linking knowledge nodes. Hence, low-knowledge subjects should sort together the two DEFECT items, as well as OXYGEN and SEPTAL DEFECT but keep PURPLE separate. This prediction reflects the fact that without the right background knowledge, the relationship between the lack of oxygen and the blood turning purple cannot be understood. The text merely asserts these separate facts, which can be linked only through the reader's background knowledge.

Figure 7
Network Formed by Four Key Words, Related Text Propositions, and Related Knowledge Items



Note. The numbers next to links indicate their activation strengths ($\times 1,000$).

The illustrative example shown in Figure 7 demonstrates that according to the model, the text should induce changes in the way the subjects sort key words for both high- and low-knowledge subjects. However, the changes should be greater and more in the direction of the text for the high- than for the low-knowledge subjects. In order to assess these predictions, McNamara et al. (1993) have scored the way each subject sorted the key words according to how close it was to the text structure. That is, when two words closely related in the text were put into the same category by a subject, a high score was given; a lower score was given when two words that were less strongly related in the text were categorized together, and so on (the scoring scheme is described in detail in the original report). The average sorting scores after reading are also shown in Figure 6b. High-knowledge subjects indeed sorted the key words according to the text more than did low-knowledge subjects.

These results are in complete agreement with the literature on knowledge effects in text comprehension and extend them by showing that domain knowledge effects are also found if one looks at learning from text rather than merely memory for a text. Given the importance of domain knowledge, it would be extremely desirable to be able to assess how close the match is between the domain knowledge of a reader (or a group of readers with common characteristics) and a particular text. If there is insufficient overlap, readers will probably be unable to learn from it; on the other hand, if there is very large overlap, readers already know what the text has to offer. What we need to do is to determine whether a text is in the "zone of learning" for a given reader. The model described above can yield such a measure. However, we have used the model here to analyze only a very small, artificial example. McNamara et al. (1993) have used it on a brief textbook chapter. Using it on long chapters or whole books

³ What students know about the heart is not necessarily correct.

would be extremely cumbersome and is, in fact, not at all feasible. The model is a laboratory tool. A major goal of our future work will be to develop practical ways to determine how a reader's prior knowledge limits his or her ability to learn—that is, the reader's zone of learning.

Learning from text depends not only on a text's content but on the way this content is expressed. Good writing obviously has its virtues—although it depends on just what we mean by “good writing,” which is the topic to be considered in the next section.

The Form of a Text

That good writing can help understanding is a truism. There are, however, limits to what can be achieved with good writing. And there are psychological reasons why inferior writing may have advantages. Good writing smoothes out the difficulties for readers and makes comprehension easy by minimizing what readers need to do for themselves. It is all there, well organized, and ready to be absorbed. However, that is not necessarily the best condition for learning. Making readers participate more actively in the comprehension process can help memory and learning. Just as self-generated items are better retained in a memory task than are experimenter-presented items, inferences that readers generate on their own may be more effective than information stated explicitly in a text. Educators generally tout the benefits of active participation in the learning process over passive knowledge absorption (e.g., Bereiter & Scardamalia, 1989; Brown & Palincsar, 1989; Resnick, 1989). Similarly, we argue that a text that spells everything out and explains everything to the last detail does not leave enough room for constructive activities on the part of the learner. What we need, instead, are texts that provide readers with opportunities to use their knowledge, that is, texts that are not well written in the sense that they spell everything out but that leave gaps for the reader to fill. At the local level that might mean that certain elaborations are not provided in the text but are left for the reader to generate; coherence relations may not always be specified, so that the reader must infer them; necessary supporting information is not provided, in the hope that the readers will be reminded of it on their own. On the global level this might mean that the macrostructure of the text is not clearly signaled but is left for the reader to figure out.

There is, of course, a catch to this proposal: It must indeed be possible for a diligent reader to generate whatever was not made explicit in the text. Thus, not any poorly written, disorganized text will do; instead, such a text will have to be carefully prepared with a certain audience in mind. The audience is important because we must be able to make specific assumptions about what readers can generate on their own and where the text must help them. For readers who know very little in a domain, a fully explicit, totally coherent, well-organized text is undoubtedly beneficial. Britton and Gulgoz (1991) have demonstrated that texts that require background knowledge readers do not have can be greatly improved by making the texts fully coherent and explicit. In their


study, they analyzed the original text for coherence gaps that they then filled in, providing the information necessary for a full understanding of the text. The result of their revision was not only a dramatic improvement in the readers' ability to recall the text but more important, Britton and Gulgoz were able to show that readers correctly understood what the text was telling them. This was not the case for the original version.

An imperfect text can stimulate active processing in knowledgeable readers. McNamara et al. (1993) argued that although such a revision would be helpful to readers with insufficient background knowledge, more knowledgeable readers would be bored by such a text and would learn better from a text that elicited more active processing. The heart disease text discussed above was in fact such a low-coherence text. As has been shown (Figure 5), high-knowledge readers performed much better with it than did low-knowledge readers. McNamara et al. also wrote a fully coherent version of this text, with the same content but with explanations and elaborations added at both the local and global levels. Subjects were given the same kinds of tests as those who read the original text.⁴ The results differed in important ways from those achieved with the original, low-coherence text. For free recall, which depends primarily on a good textbase, the more coherent text led to better performance for both high- and low-knowledge subjects. But interesting reversals occurred for the tasks that depended more on a good situation model, that is, that required an active construction process from the reader. For problem-solving questions, low-knowledge subjects profited from the more coherent text and answered more questions correctly (30%) than when they read the low-coherence text (17%). High-knowledge readers, on the other hand, actually performed better with the low-coherence text (61% vs. 46% for the high-coherence text), in agreement with the expectation that a more demanding text would stimulate greater cognitive activity, so that they would form a better, more elaborate situation model. This interaction was equally pronounced in the postreading sorting scores of high- and low-knowledge readers, shown in Figure 8. The left half of the figure simply duplicates what was already shown in Figure 6: After reading a low-coherence text on heart disease, readers with good background knowledge sorted key words in agreement with the text more than did readers who knew little about the heart. The right half of the figure, however, shows that this relationship was reversed when subjects were given the fully coherent text. Low-knowledge readers did much better when they read a text they could really understand. High-knowledge readers, in contrast, did worse with the well-written, coherent text than with the less well-organized text that did not spell everything out.

⁴ Other possibilities would be the arithmetic or geometric mean of all links, perhaps with some modification for the number of links. In the present case, such schemes yield qualitatively similar results to the one adopted here.

Figure 8

Postreading Sorting Scores for High- and Low-Knowledge Subjects for the Original (Low-Coherence) and Revised (High-Coherence) Text



It is clear why low-knowledge readers do not perform well with the low-coherence text: The text requires that the reader fill in a great deal of what is assumed to be known, and when this knowledge is lacking, even an intelligent and well-motivated reader is at a loss. But why do readers with good background knowledge not perform well with a fully explicit text? The information content of the two texts in this experiment was the same so that, in principle, even readers who knew a lot could obtain all they needed from either text. The answer may be that the text was so easy for them that they felt they understood it well, without actually being sufficiently challenged to work out all the details. A feeling of understanding at the level of the textbase can conceal incomplete understanding at the level of the situation model. By preventing these readers from easily forming a textbase, one can make them work harder. This additional work must take place at the level of the situation model. For example, if two sentences are not explicitly connected, there is no linguistic cue in the text by which one could select a suitable sentence connective—say, *because*, or *therefore*, or *but*. The reader can make such a connection only on the basis of a deep understanding of the situation. Thus, by omitting linguistic surface features, deep situational processing on the part of the reader can be encouraged.

There are, of course, other ways than giving a learner a poorly written text to stimulate his or her activity. One can, for instance, teach the learner strategies for active reading. Particularly successful has been the teaching of self-explanation strategies (Chi, Bassok, Lewis, Reimann, & Glaser, 1989): Readers who are taught to explain the text to themselves achieve a deeper understanding, much like readers who must figure out gaps in a text. Another strategy that has proved very effective is questioning the author (Beck, McKeown, Worthy, Sandora, & Kucan, 1993): The student learns not to accept the author's statements but to argue and test the claims that are made and the manner in which they are presented.

What all of these methods have in common is that they make reading difficult. The problem for learning from text is that reading can be too easy: It takes very little effort to read most texts with a good sense of understanding. All too often, however, that understanding is only superficial. Where no adequate situation model has been formed and the new information is not linked to the reader's knowledge base, the text can be remembered for a while, but learning—in the sense that this term has been used here—has not occurred. Learning requires the active construction of a situation model, integrating text information with the reader's prior knowledge.

Conclusion

Content overlap between text and knowledge appears to be a necessary condition for learning from text. Our research is concerned with determining the "learnability" zones for people with a given knowledge base: Texts that are optimal for learning should overlap in content sufficiently, but not totally, with what readers already know. Texts outside the zone may be remembered, but they are very likely to remain inert knowledge, not to be well integrated with prior knowledge. We have not proved this hypothesis conclusively yet, but the evidence presented here is suggestive.

If there is insufficient content overlap, general language signals become crucial for the reader. Good writing for this reader means being coherent, explicit, and clearly structured. On the other hand, if there is the right amount of content overlap, texts with coherence gaps that stimulate the reader's active inferencing are optimal. Good writing here means letting the reader construct his or her own memory structure. Thus, even though there can be no text that is optimal for everyone, the research reported here suggests that methods can be developed to match readers and texts in a principled and practical manner.

REFERENCES

- Beck, I. L., McKeown, M. G., Worthy, J., Sandora, C. A., & Kucan, L. (1993). *Questioning the author: A year long classroom implementation to engage students with text*. Unpublished manuscript, University of Pittsburgh.
- Bereiter, C., & Scardamalia, M. (1989). Intentional learning as a goal of instruction. In L. B. Resnick (Ed.), *Knowing, learning, and instruction* (pp. 361-392). Hillsdale, NJ: Erlbaum.
- Bransford, J. D., Barclay, J. R., & Franks, J. J. (1972). Sentence memory: A constructive versus interpretive approach. *Cognitive Psychology*, 3, 193-209.
- Britton, B. K., & Gulgoz, S. (1991). Using Kintsch's model to improve instructional text: Effects of inference calls on recall and cognitive structures. *Journal of Educational Psychology*, 83, 329-345.
- Brown, A. L., & Palincsar, A. S. (1989). Guided, cooperative learning and individual knowledge acquisition. In L. B. Resnick (Ed.), *Knowing, learning, and instruction* (pp. 393-452). Hillsdale, NJ: Erlbaum.
- Brown, A. L., & Reeves, R. A. (1987). Bandwidths of competence: The role of supportive contexts in learning and development. In L. S. Liben (Ed.), *Development and learning: Conflict or congruence?* (pp. 173-223). Hillsdale, NJ: Erlbaum.
- Burton, R. R., Brown, J. S., & Fischer, G. (1984). Skiing as a model of instruction. In B. Rogoff & J. Lave (Eds.), *Everyday cognition: Its*

- development and social context (pp. 139-150). Cambridge, MA: Harvard University Press.
- Chi, M. T. H., Bassok, M., Lewis, M., Reimann, P., & Glaser, R. (1989). Self-explanation: How students study and use examples in learning to solve problems. *Cognitive Science*, *13*, 145-182.
- Ericsson, K. A., & Kintsch, W. (1991). *Memory in comprehension and problem solving: A long-term working memory* (Tech. Rep. ICS No. 91-13). Boulder: University of Colorado.
- Kintsch, W. (1988). The use of knowledge in discourse processing: A construction-integration model. *Psychological Review*, *95*, 163-182.
- Kintsch, W. (1992). A cognitive architecture for comprehension. In H. L. Pick, P. van den Broek, & D. C. Knill (Eds.), *The study of cognition: Conceptual and methodological issues* (pp. 143-164). Washington, DC: American Psychological Association.
- Kintsch, W., & Vipond, D. (1979). Reading comprehension and readability in educational practice and psychological theory. In L. G. Nilsson (Ed.), *Perspectives of memory research* (pp. 325-366). Hillsdale, NJ: Erlbaum.
- Mannes, S. M., & Kintsch, W. (1987). Knowledge organization and text organization. *Cognition and Instruction*, *4*, 91-115.
- McNamara, D. S., Kintsch, E., Butler-Songer, N., & Kintsch, W. (1993). *Text coherence, background knowledge, and levels of understanding in learning from text*. Manuscript submitted for publication.
- Moravcsik, J. E., & Kintsch, W. (1993). Writing quality, reading skills, and domain knowledge as factors in text comprehension. *Canadian Journal of Experimental Psychology*, *47*, 360-374.
- Raajmaker, J. G., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, *88*, 93-134.
- Resnick, L. B. (Ed.). (1989). *Knowing, learning, and instruction: Essays in honor of Robert Glaser*. Hillsdale, NJ: Erlbaum.
- Schneider, W., Körkel, J., & Weinert, F. E. (1989). Domain-specific knowledge and memory performance: A comparison of high- and low-aptitude children. *Journal of Educational Psychology*, *81*, 306-312.
- van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension*. San Diego, CA: Academic Press.
- Voss, J. F., Fincher-Kiefer, R. H., Greene, T. R., & Post, T. A. (1986). Individual differences in performance: The contrastive approach to knowledge. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (pp. 297-334). Hillsdale, NJ: Erlbaum.
- Vygotsky, L. S. (1986). *Thought and language*. Cambridge, MA: MIT Press. (Original work published 1934)
- Weaver, C. A., III, & Kintsch, W. (1990). Expository text. In R. Barr, M. L. Kamil, P. B. Rosenthal, & P. D. Pearson (Eds.), *Handbook of reading* (pp. 230-245). New York: Longman.
- Wertheimer, M. (1945). *Productive thinking*. Chicago: University of Chicago Press.